

Forecasting Default with the KMV-Merton Model

SREEDHAR T BHARATH and TYLER SHUMWAY*
University of Michigan

April 26, 2005

Abstract

We examine the accuracy and contribution of the default forecasting model based on Merton's (1974) bond pricing model and developed by the KMV corporation. Comparing the KMV-Merton model to a similar but much simpler alternative, we find that it performs slightly worse as a predictor in hazard models and in out of sample forecasts. Moreover, several other forecasting variables are also important predictors, and fitted hazard model values outperform KMV-Merton default probabilities out of sample. Implied default probabilities from credit default swaps and corporate bond yield spreads are only weakly correlated with KMV-Merton default probabilities after adjusting for agency ratings, bond characteristics, and our alternative predictor. We conclude that the KMV-Merton model does not produce a sufficient statistic for the probability of default, and it appears to be possible to construct such a sufficient statistic without solving the nonlinear equations required by the KMV-Merton model.

*Department of Finance, University of Michigan Business School, 701 Tappan Street, Ann Arbor, MI 48109. Bharath can be reached at 734-763-0485 or sbharath@umich.edu. Shumway is visiting Stanford GSB for 2004-05, and can be reached at 650-725-9265 or shumway@umich.edu. We thank seminar participants at Michigan, Boston College, and Stanford. We also thank Bill Beaver, Darrell Duffie, Wayne Ferson, Kyle Lundstedt, Ken Singleton, and Jorge Sobehart for their comments.

Due to the advent of innovative corporate debt products and credit derivatives, academics and practitioners have recently shown renewed interest in models that forecast corporate defaults. One innovative forecasting model which has been widely applied in both practice and academic research¹ is a particular application of Merton’s model (Merton, 1974) that was developed by the KMV corporation, which we refer to as the KMV-Merton model². This paper assesses the accuracy and the contribution of the KMV-Merton model.

The KMV-Merton model applies the framework of Merton (1974), in which the equity of the firm is a call option on the underlying value of the firm with a strike price equal to the face value of the firm’s debt. The model recognizes that neither the underlying value of the firm nor its volatility are directly observable. Under the model’s assumptions both can be inferred from the value of equity, the volatility of equity and several other observable variables by solving two nonlinear simultaneous equations. After inferring these values, the model specifies that the probability of default is the normal cumulative density function of a z-score depending on the firm’s underlying value, the firm’s volatility and the face value of the firm’s debt.

The KMV-Merton model is a clever application of classic finance theory, but how well it performs in forecasting depends on how realistic its assumptions are. The model is a somewhat stylized structural model that requires a number of assumptions. Among other things, the model assumes that the underlying value of each firm follows geometric Brownian motion and that each firm has issued just one zero-coupon bond. If the model’s strong assumptions are violated, it should be possible to construct a reduced form model with more accuracy.

We examine two hypotheses in this paper. First, we ask whether the probability of default implied by the Merton model is a sufficient statistic for forecasting bankruptcy. If the Merton model is literally true, it should be impossible to improve on the model’s implied probability for forecasting. If it is possible to construct a reduced form model with better predictive properties, we can conclude that the KMV-Merton probability (π_{KMV}) is not a sufficient statistic for forecasting default.

¹The model is discussed in Duffie and Singleton (2003) and Saunders and Allen (2002). It is applied by Vassalou and Xing (2003), among others.

²While others refer to this model simply as a Merton model, we prefer to call it the KMV-Merton model because (1) deriving the KMV-Merton default probability from observed equity data is a nontrivial extension of the ideas in the classic Merton model and (2) the proprietors of KMV developed this clever extension of the Merton model and we believe they deserve some credit for its development. We do not intend to imply that we are using exactly the same algorithm that Moody’s KMV uses to calculate distance to default. Differences between our method and that of Moody’s KMV are discussed in Section I B and in Table 2.

Our second hypothesis is that the Merton model is an important quantity to consider when predicting default. We hypothesize that the information in π_{KMV} cannot be completely replaced by a reasonable set of simple variables, or that a sufficient statistic for default probability cannot neglect π_{KMV} . We actually separate the KMV-Merton technique into two potentially important components: the functional form for default probability implied by the Merton model and the solution of nonlinear equations required by the model. It is possible that one of these components is important while the other is not.

We test these two hypotheses in five ways. First, we incorporate π_{KMV} into a hazard model that forecasts defaults from 1980 through 2003. With the hazard model, we compare π_{KMV} to a naive alternative (π_{Naive}) which is much simpler to calculate, but retains some of the functional form of π_{KMV} . We also compare it to several other default forecasting variables. Second, we compare the short term, out of sample forecasting ability of π_{KMV} to that of π_{Naive} . Third, we examine the forecasting ability of several alternative predictors, each of which calculates KMV-Merton probabilities in a slightly different way. Fourth, we examine the ability of KMV-Merton probabilities to explain the probability of default implied by credit default swaps, and fifth we regress corporate bond yield spreads on π_{KMV} , π_{Naive} and other variables.

Assessing the KMV-Merton model's value is of importance for two reasons. Perhaps the most important reason is that many researchers and practitioners are applying the model and yet we do not know very much about its statistical properties. For example, Vassalou and Xing (2003) use π_{KMV} to examine whether default risk is priced in equity returns. As a second example, the Basel Committee on Banking Supervision (1999) considers exploiting the KMV-Merton model a viable practice currently employed by numerous banks. To have confidence in both the risk management of the banking sector and the accuracy of academic research, the power of the KMV-Merton model must be examined.

A second reason to assess the KMV-Merton model is to test the Merton (1974) model in a new way. If the Merton model is literally true, π_{KMV} should be the best default predictor available. The Merton model has been rejected previously for failing to fit observed bond yield spreads.³ Comparing the model to reduced form alternatives gives us a fresh perspective about how realistic the model's assumptions are.

Over the past several years, a number of reserchers have examined the contribution of the KMV-

³see Jones et al. (1984).

Merton Model. The first authors to examine the model carefully were practitioners employed by either KMV or Moody's. A couple of years ago, several papers addressing the accuracy of the KMV-Merton model were available on the internet. Some papers, including Stein (2000), Sobehart and Stein (2000) and Sobehart and Keenan (1999) argued that KMV-Merton models can easily be improved upon. Other papers, including Kealhofer and Kurbat (2001), argued that KMV-Merton models capture all of the information in traditional agency ratings and well known accounting variables. Curiously, while some practitioner papers can now be found in print, including Sobehart and Keenan (2002a) and (2002b) and Falkenstein and Boral (2001), it has become very difficult to find electronic copies of some of the papers cited above since Moody's acquired KMV in April 2002.

Perhaps in response, an academic literature has recently developed that critically assesses the model. Both Hillegeist, Keating, Cram and Lundstedt (2004) and Du and Suo (2004) examine the model's predictive power in ways that are similar to some of our analyses. Duffie and Wang (2004) show that KMV-Merton probabilities have significant predictive power in a model of default probabilities over time, which can generate a term structure of default probabilities. Campbell, Hilscher and Szilagyi (2004) estimate hazard models that incorporate both π_{KMV} and other variables for bankruptcy, finding that π_{KMV} seems to have relatively little forecasting power after conditioning on other variables. While our findings are consistent with the findings of all of these papers, we analyze the performance of π_{KMV} in several novel ways. In particular, we introduce and assess our naive predictor and we examine the ability of π_{KMV} to explain credit default swap premia and bond yield spreads. Like all of these researchers, we have no particular interest in finding evidence for or against the KMV-Merton model. Therefore, we hope to help resolve confusion about some of the issues raised in the practitioner literature described above.

We find that it is fairly easy to reject hypothesis one, or that π_{KMV} is not a sufficient statistic for default probability. We also find that after conditioning on π_{Naive} , it appears to be possible to construct a reduced form model that does not benefit by conditioning on π_{KMV} (or in which π_{KMV} is not statistically significant). We therefore conclude that while π_{KMV} has some predictive power for default, most of the marginal benefit of π_{KMV} comes from its functional form rather than from the solution of the two nonlinear equations on which it is based. The contribution of π_{KMV} to a well-specified reduced form model is fairly low.

The paper proceeds as follows. The next section details the KMV-Merton model, our naive

alternative default probability, and the hazard models that we use to build reduced form models. Section I also lists several ways in which our KMV-Merton model differs from the model that Moodys KMV actually sells. Section II discusses the data that we use for our tests and Section III outlines our results. We conclude in Section IV.

I. Default Forecasting Models

As discussed above, we examine our hypotheses by examining the statistical and economic significance of the KMV-Merton default probabilities (π_{KMV}) and a simple, naive alternative (π_{Naive}). Before examining the empirical value of these variables, we need to describe them carefully. The KMV-Merton model was developed by the KMV corporation in the late 1980s. It was successfully marketed by KMV until KMV was acquired by Moodys in April 2002. The model is now sold to subscribers by Moody's KMV.

A. The KMV-Merton Model

The KMV-Merton default forecasting model produces a probability of default for each firm in the sample at any given point in time. To calculate the probability, the model subtracts the face value of the firm's debt from from an estimate of the market value of the firm and then divides this difference by an estimate of the volatility of the firm (scaled to reflect the horizon of the forecast). The resulting z-score, which is referred to as the distance to default, is then substituted into a cumulative density function to calculate the probability that the value of the firm will be less than the face value of debt at the forecasting horizon. The market value of the firm is simply the sum of the market values of the firm's debt and the value of its equity. If both these quantities were readily observable, calculating default probabilities would be simple. While equity values are readily available, reliable data on the market value of firm debt is generally unavailable.

The KMV-Merton model estimates the market value of debt by applying the Merton (1974) bond pricing model. The Merton model makes two particularly important assumptions. The first is that the total value of a firm is assumed to follow geometric Brownian motion,

$$dV = \mu V dt + \sigma_V V dW \tag{1}$$

where V is the total value of the firm, μ is the expected continuously compounded return on V , σ_V is the volatility of firm value and dW is a standard Weiner process. The second critical assumption of the Merton model is that the firm has issued just one discount bond maturing in T periods. Under these assumptions, the equity of the firm is a call option on the underlying value of the firm with a strike price equal to the face value of the firm's debt and a time-to-maturity of T . Moreover, the value of equity as a function of the total value of the firm can be described by the Black-Scholes-Merton Formula. By put-call parity, the value of the firm's debt is equal to the value of a risk-free discount bond minus the value of a put option written on the firm, again with a strike price equal to the face value of debt and a time-to-maturity of T .

Symbolically, the Merton model stipulates that the equity value of a firm satisfies

$$E = V\mathcal{N}(d_1) - e^{-rT}F\mathcal{N}(d_2), \quad (2)$$

where E is the market value of the firm's equity, F is the face value of the firm's debt, r is the instantaneous risk-free rate, $\mathcal{N}(\cdot)$ is the cumulative standard normal distribution function, d_1 is given by

$$d_1 = \frac{\ln(V/F) + (r + 0.5\sigma_V^2)T}{\sigma_V\sqrt{T}}, \quad (3)$$

and d_2 is just $d_2 = d_1 - \sigma_V\sqrt{T}$. While this is a fairly complicated equation, most financial economists are familiar with this formula as the Black-Scholes-Merton option valuation equation.

The KMV-Merton model makes use of two important equations. The first is the Black-Scholes-Merton equation (2), expressing the value of a firm's equity as a function of the value of the firm. The second relates the volatility of the firm's value to the volatility of its equity. Under Merton's assumptions the value of equity is a function of the value of the firm and time, so it follows directly from Ito's lemma that

$$\sigma_E = \left(\frac{V}{E}\right) \frac{\partial E}{\partial V} \sigma_V. \quad (4)$$

In the Black-Scholes-Merton model, it can be shown that $\frac{\partial E}{\partial V} = \mathcal{N}(d_1)$, so that under the Merton model's assumptions, the volatilities of the firm and its equity are related by

$$\sigma_E = \left(\frac{V}{E}\right) \mathcal{N}(d_1) \sigma_V, \quad (5)$$

where d_1 is defined in equation (3).

The KMV-Merton model basically uses these two nonlinear equations, (2) and (5), to translate the value and volatility of a firm's equity into an implied probability of default. In most applications, the Black-Scholes-Merton model describes the unobserved value of an option as a function of four variables that are easily observed (strike price, time-to-maturity, underlying asset price, and the risk-free rate) and one variable that can be estimated (volatility).⁴ In the KMV-Merton model, however, the value of the option is observed as the total value of the firm's equity, while the value of the underlying asset (the value of the firm) is not directly observable. Thus, while V must be inferred, E is easy to observe in the marketplace by multiplying the firm's shares outstanding by its current stock price. Similarly, in the KMV-Merton model, the volatility of equity, σ_E , can be estimated but the volatility of the underlying firm, σ_V must be inferred.

The first step in implementing the KMV-Merton model is to estimate σ_E from either historical stock returns data or from option implied volatility data. The second step is to choose a forecasting horizon and a measure of the face value of the firm's debt. For example, it is common to use historical returns data to estimate σ_E , assume a forecasting horizon of one year ($T = 1$), and take the book value of the firm's total liabilities to be the face value of the firm's debt. The third step is to collect values of the risk-free rate and the market equity of the firm. After performing these three steps, we have values for each of the variables in equations (2) and (5) except for V and σ_V , the total value of the firm and the volatility of firm value respectively.

The fourth, and perhaps most significant step in implementing the model is to simultaneously solve equations (2) and (5) numerically for values of V and σ_V . Once this numerical solution is obtained, the distance to default can be calculated as

$$DD = \frac{\ln(V/F) + (\mu - 0.5\sigma_V^2)T}{\sigma_V\sqrt{T}}, \quad (6)$$

where μ is an estimate of the expected annual return of the firm's assets. The corresponding implied probability of default, sometimes called the expected default frequency (or EDF), is

$$\pi_{\text{KMV}} = \mathcal{N}\left(-\left(\frac{\ln(V/F) + (\mu - 0.5\sigma_V^2)T}{\sigma_V\sqrt{T}}\right)\right) = \mathcal{N}(-DD). \quad (7)$$

If the assumptions of the Merton model really hold, the KMV-Merton model should give very accurate default forecasts. In fact, if the Merton model holds completely, the implied probability

⁴Of course, it is common to infer an implied volatility from an observed option price.

of default defined above, π_{KMV} , should be a sufficient statistic for default forecasts. Testing this hypothesis is one of the central task of this paper.

Simultaneously solving equations (2) and (5) is a reasonably straightforward thing to do. However, KMV does not simply solve these equations numerically. Crosbie and Bohn (2001) explain that “In practice the market leverage moves around far too much for [equation (5)] to provide reasonable results.” To resolve this problem, we follow KMV by implementing a complicated iterative procedure. First, we propose an initial value of $\sigma_V = \sigma_E[E/(E + F)]$ and we use this value of σ_V and equation (2) to infer the market value of each firm’s assets every day for the previous year. We then calculate the implied log return on assets each day and use that returns series to generate new estimates of σ_V and μ . We iterate on σ_V in this manner until it converges (so the absolute difference in adjacent σ_V s is less than 10^{-3}). Unless specified otherwise, in the rest of the paper values of π_{KMV} are calculated by following this iterative procedure and calculating the corresponding implied default probability using equation (7).

Before describing alternative models, it is useful to interpret the KMV-Merton model a little. The most critical inputs to the model are clearly the market value of equity, the face value of debt, and the volatility of equity. As the market value of equity declines, the probability of default increases. This is both a strength and weakness of the model. For the model to work well, both the Merton model assumptions must be met and markets must be efficient and well informed.

In its promotional material, KMV points to the Enron case as an example of how their method is superior to that of traditional agency ratings. When Enron’s stock price began to fall, its distance to default immediately decreased. The ratings agencies took several days to downgrade Enron’s debt. Clearly, using equity values to infer default probabilities allows the KMV-Merton model to reflect information faster than traditional agency ratings. However, when Enron’s stock price was unsustainably high, KMV’s expected default frequency for Enron was actually significantly lower than the default probability assigned to Enron by standard ratings. If markets are not perfectly efficient, then conditioning on information not captured by π_{KMV} probably makes sense.

B. Our Method versus Moody’s KMV

We should point out that there a number of things which differentiate the KMV-Merton model which we test from that actually employed by Moody’s KMV. One important difference is that we use Merton’s model while Moody’s KMV uses a proprietary model that they call the KV model.

Apparently the KV model is a generalization of the Merton model that allows for various classes and maturities of debt. Another difference is that we use the cumulative normal distribution to convert distances to default into default probabilities. Moody’s KMV uses its large historical database to estimate the empirical distribution of changes in distances to default and it calculates default probabilities based on that distribution. The distribution of distances to default is an important input to default probabilities, but it is not required for ranking firms by their relative probability. Therefore, several of our results will emphasize the model’s ability to rank firms by default risk rather than its ability to calculate accurate probabilities⁵. Finally, KMV may also make proprietary adjustments to the accounting information that they use to calculate the face value of debt. We cannot perfectly replicate the methods of Moody’s KMV because several of the modeling choices made by Moody’s KMV are proprietary information, and subscribing to their database is prohibitively expensive for us.

While our method does not match that of Moody’s KMV exactly, it is the same method employed by Vassalou and Xing (2003) and other researchers. Our results can be considered relevant for a “feasible” KMV-Merton model, which can be estimated and implemented by academic researchers or practitioners that do not want to subscribe to Moody’s KMV. We should note that it is entirely possible that the proprietary features of KMV’s model make its performance superior to what we document here. In order to compare our method with that of Moody’s KMV, in Section III we compare our estimates with the estimates produced by Moody’s KMV for a sample of large firms in the US.

C. A Naive Alternative

To test whether π_{KMV} adds value to reduced form models, we construct a simple alternative “probability” that does not require simultaneously solving equations (2) and (5) or implementing the iterative procedure described above. We construct our naive predictor with two objectives. First, we want our naive predictor to have a reasonable chance of performing as well as the KMV-Merton predictor, so we want it to capture the same information that the KMV-Merton predictor uses. We also want our naive probability to approximate the functional form of the KMV-Merton probability. Second, we want our naive probability to be simple, so we avoid simultaneously solving any

⁵If the model ranks firms accurately then using historical data to map relative rankings into accurate probabilities is a straightforward task.

equations or estimating any difficult quantities in its construction. We wrote down the form for our naive probability after studying the KMV-Merton model for a little while. None of the numerical choices below is the result of any type of estimation or optimization.

To begin constructing our naive probability, we approximate the market value of each firm's debt with the face value of its debt,

$$\text{Naive } D = F, \quad (8)$$

Since firms that are close to default have very risky debt, and the risk of their debt is correlated with their equity risk, we approximate the volatility of each firm's debt as

$$\text{Naive } \sigma_D = 0.05 + 0.25 * \sigma_E. \quad (9)$$

We include the five percentage points in this term to represent term structure volatility, and we include the twenty-five percent times equity volatility to allow for volatility associated with default risk. This gives us an approximation to the total volatility of the firm of

$$\text{Naive } \sigma_V = \frac{E}{E + \text{Naive } D} \sigma_E + \frac{\text{Naive } D}{E + \text{Naive } D} \text{Naive } \sigma_D = \frac{E}{E + F} \sigma_E + \frac{F}{E + F} (0.05 + 0.25 * \sigma_E). \quad (10)$$

Next, We set the expected return on the firm's assets equal to the firm's stock return over the previous year,

$$\text{Naive } \mu = r_{it-1}. \quad (11)$$

This allows us to capture some of the same information that is captured by the KMV-Merton iterative procedure described above. The iterative procedure is able to condition on an entire year of equity return data. By allowing our naive estimate of μ to depend on past returns, we incorporate the same information. The naive distance to default is then

$$\text{Naive } DD = \frac{\ln[(E + F)/F] + (r_{it-1} - 0.5 \text{Naive } \sigma_V^2)T}{\text{Naive } \sigma_V \sqrt{T}}. \quad (12)$$

This naive alternative model is easy to compute – it does not require solving the equations simultaneously. However, it retains the structure of the KMV-Merton distance to default and expected default frequency. It also captures approximately the same quantity of information as the KMV-Merton probability. Thus, examining the forecasting ability of this quantity will help us separate

the value of simultaneously solving the equations and the value of the functional form of π_{KMV} . We define our naive probability estimate as

$$\pi_{\text{Naive}} = \mathcal{N}(-\text{Naive } DD). \quad (13)$$

It is fairly easy to criticize our naive probability. Our choices for modeling firm volatility are not particularly well motivated and our decision to use past returns for μ is arbitrary at best. However, to quibble with our naive probability is to miss the point of our exercise. We have constructed a predictor that is extremely easy to calculate, and it may have significant predictive power. If the predictive power of our naive probability is comparable to that of π_{KMV} , then presumably a more carefully constructed probability that captures the same information should have superior power.

D. Alternative Predictors

One purpose of our paper is to examine the relative importance of several of the components of the KMV-Merton calculation. Comparing the predictive performance of our naive probability to that of π_{KMV} is one way to accomplish this purpose. Another way we accomplish this purpose is by examining the predictive performance of several alternative predictors, or predictors that calculate KMV-Merton default probabilities in alternative, somewhat simpler ways.

One predictor, $\pi_{\text{KMV}}^{\mu=r}$, is calculated in exactly the same manner as π_{KMV} , except that the expected return on assets used for π_{KMV} is replaced by the risk-free rate, r . Considering this predictor helps us gauge the importance of estimating the expected return on assets for the distance to default. A second alternative predictor, $\pi_{\text{KMV}}^{\text{simul}}$, is calculated by simultaneously solving equations (2) and (5). This predictor avoids the iterative procedure in the text, estimating equity volatility with one year of historical returns data and using r as the expected return on assets. The third alternative predictor, $\pi_{\text{KMV}}^{\text{imp}\sigma}$, uses the option-implied volatility of firm equity (implied σ_E) to simultaneously solve equations (2) and (5).

E. Hazard Models

In order to assess the KMV-Merton model's accuracy, we need a method to compare π_{KMV} to alternative predictor variables. We employ a Cox proportional hazard model to test our two hypotheses. Hazard models have recently been applied by a number of authors and probably

represent the state of the art in default forecasting with reduced form models⁶. Proportional hazard models make the assumption that the hazard rate, $\lambda(t)$, or the probability of default at time t conditional on survival (lack of default) until time t is,

$$\lambda(t) = \phi(t)[\exp(x(t)' \beta)], \tag{14}$$

where $\phi(t)$ is referred to as the “baseline” hazard rate and the term $\exp(x(t)' \beta)$ allows the expected time to default to vary across firms according to their covariates, $x(t)$. The baseline hazard rate is common to all firms. Note that in this model the covariates may vary with time. Most of our default predictors, including the KMV-Merton probability, vary with time. The Cox proportional hazard model does not impose any structure on the baseline hazard $\phi(t)$. Cox’s partial likelihood estimator provides a way of estimating β without requiring estimates of $\phi(t)$. It can also handle censoring of observations, which is one of the features of the data. Details about estimating the proportional hazard model can be found in many places, including in Cox and Oakes (1984).

Our first hypothesis, that the KMV-Merton probability is a sufficient statistic for forecasting default, implies that no other variable in a hazard model should be a stastically significant covariate. Our second hypothesis, that the KMV-Merton probability is a useful quantity, implies that no other set of variables should be able to make the probability insignificant. As a robustness check, we will also sort firms each year by their probabilities from each model and find the number of defaults in several bins.

F. Implied Probabilities of Default

Besides examining the default prediction ability of the KMV-Merton model, we examine its ability to explain the variation in two market-based default probability variables. We regress both the implied probability of default from credit default swaps (CDS) and the yield spread on corporate bonds on π_{KMV} and π_{Naive} . While there is a large literature on explaining bond yield spreads, using CDS data to assess default probabilities is relatively new. Other recent papers that use CDS data include Longstaff, Mithay and Neis (2004) and Berndt, Douglas, Duffie, Ferguson and Schranz (2004).

Credit default swaps are one example of credit derivatives, and credit derivative markets have

⁶Shumway (2001) and Chava and Jarrow (2004) argue that hazard models are superior to other types of models.

experienced explosive growth in recent years. According to the British Bankers' Association the total notional principal for outstanding credit derivatives increased from \$180 billion in 1997 to more than \$2 trillion by the end of 2002 and it is expected to reach \$4.8 trillion by the end of 2004. Popular credit derivatives such as the credit default swap allow market participants to trade credit risks with each other. We use the information in credit default swap premia to extract a direct measure of default probabilities and compare it with the estimates obtained from our methods.

In a credit default swap, the party buying credit protection pays the seller a fixed premium until either default occurs or the swap contract matures. In the event of a default, because these payments are made in arrears, a final accrual payment by the buyer is required. In return, if the underlying firm (the reference entity) defaults on its debt, the protection seller is obligated to buy back from the buyer the defaulted bond at its par value. The pay off from a credit default swap is simply one minus the recovery rate, which is the loss given default for every dollar of notional principal. Thus a CDS is similar to an insurance contract compensating the buyer for losses arising from a default.

Let s be the CDS spread, which is the amount paid per year as a percentage of the notional principal. Most CDS contracts have a maturity of five years. Let T determine the life of the CDS contract. Further assume that the probability of a reference entity defaulting during a year conditional on no earlier default is π_{CDS} . For simplicity we assume that defaults always happen halfway through the year and the payments on the CDS are made once a year, at the end of each year. Thus the final accrual payment will be made halfway through the year and will be equal to $0.5s$. We also assume that the risk-free (LIBOR) rate is r with continuous compounding and the recovery is δ .

Thus the expected present value of the payments made on the CDS (assuming a notional principal of \$1) is given by

$$\sum_{t=1}^{t=T} (1 - \pi_{\text{CDS}})^t e^{-rt} s + (1 - \pi_{\text{CDS}})^{t-1} \pi_{\text{CDS}} e^{-r(t-0.5)} 0.5s \quad (15)$$

The first term represents the discounted present value of the expected payments made at the end of each year provided the reference entity survives until period t and the second term represents the present value of the accrual payments made in the case of a default assuming default happens midway through the year.

Similarly the expected present value of the payoff is given by

$$\sum_{t=1}^{t=T} (1 - \pi_{\text{CDS}})^{t-1} \pi_{\text{CDS}} (1 - \delta) e^{-r(t-0.5)} \quad (16)$$

We need an implied estimate of recovery rate δ in order to value the payoff. The same recovery rate is typically used to (a) estimate implied default probabilities and (b) value the CDS. The net result of this is that the value of a CDS (or the estimate of a CDS spread) is not very sensitive to the recovery rate. This is because implied probabilities of default are approximately proportional to $1/(1-\delta)$ and the payoffs from a CDS are proportional to $(1-\delta)$, so that the expected payoff is almost independent of δ (Hull, Predescu and White (2004)).

Setting the present value of the expected payments equal to the expected payoffs we can solve for π_{CDS} from the resulting non linear equation. Since this calculation assumes that there is no risk premium associated with default, the resulting implied probability should be considered a risk-neutral default probability. As described above, we solve for π_{CDS} for a sample of CDS spreads, and regress π_{CDS} on π_{KMV} , π_{Naive} , and other variables. The results of our regressions are described in Section III E and Table 6.

II. Data

We begin by examining all firms in the intersection of the Compustat Industrial file - Quarterly data and CRSP daily stock return for NYSE, AMEX and NASDAQ stocks between 1980 and 2003. We exclude financial firms (SIC codes 6021,6022,6029,6035,6036) from the sample.

We obtain default data for the period 1980-2000 from the data base of firm default maintained by Edward Altman (The Altman default database). We supplement this information for 2001 through 2003 by using the list of defaults published by Moody's at their website www.moodys.com. In all we obtain a total of 1,449 firm defaults covering the period 1980-2003.

The inputs to the KMV-Merton model include σ_E the volatility of stock returns, F the face value of debt, r the risk free rate and T the time period. σ_E is the annualized percent standard deviation of returns and is estimated from the prior year stock return data for each month. For r , the risk free rate, we use the one-year Treasury Constant Maturity Rate obtained from the Board of Governors of the Federal Reserve system⁷. E , the market value of each firm's equity (in millions

⁷Available at <http://research.stlouisfed.org/fred/data/irates/gsl> (H.15 Release)

of dollars), is calculated from the CRSP database as the product of share price at the end of the month and the number of shares outstanding. Following Vassalou and Xing (2003), we take F , the face value of debt, to be debt in current liabilities (COMPUSTAT data item 45) plus one half of long term debt (COMPUSTAT data item 51). In addition to the above variables, following Shumway (2001), we measure each firm's past excess return in year t as the return of the firm in year $t-1$ minus the value-weighted CRSP NYSE/AMEX index return in year $t-1$ ($r_{it-1} - r_{mt-1}$). Each firm's annual returns are calculated by cumulating monthly returns. We also collect each firm's ratio of net income to total assets. These variables, though not required for the KMV-Merton model, will augment the information set for the alternative models we consider later in the paper.

There are a number of extreme values among the observations of each variable constructed from raw COMPUSTAT data. To ensure that statistical results are not heavily influenced by outliers, we set all observations higher than the 99th percentile of each variable to that value. All values lower than the first percentile of each variable are winsorized in the same manner. The minimum and maximum numbers reported in Table 1 are calculated after winsorization.⁸ Table 1 provides summary statistics for all the variables described above.

Looking at the summary statistics in Table 1, it is slightly odd that the average firm's past excess return is -8.7 percent. This value is negative because of the winsorization of the upper tail extreme values at the 99th percentile level. More significantly, the distribution of the expected default frequency obtained from the Merton model, π_{KMV} , is very similar to the naive alternative, π_{Naive} . Our point estimate of 10.95% for the mean value of π_{KMV} in 1980-2003 is a bit higher than the estimate of 4.21% for the period 1971-1999 reported in Vassalou and Xing (2003). The correlation between the naive and Merton model expected default frequencies is very high at 86 percent, and it is significant at the 1% level. The similarity in distributions is also evident between the naive and Merton model estimates of asset volatility. The correlation between the 2 asset volatilities is 87 percent, and it is also significant at the 1% level.

Given that the naive counterparts (π_{Naive} and Naive σ_V) of the output from the Merton model (π_{KMV} and σ_V) are quite similar, what is the incremental value of solving the KMV-Merton model? The next section addresses this question.

⁸We do not winsorize the expected default frequency measures from the Merton Model and the naive alternative, since these are naturally bounded between 0 and 1.

III. Results

We present a number of empirical results, including correlations of our probability estimates with those published by Moody’s KMV, estimates of hazard models for time to default, out of sample forecast assessments, CDS implied default probability regressions and bond yield spread regressions. We discuss each type of result in turn.

A. Comparing Moody’s KMV Probabilities to Ours

As mentioned above, our method for calculating π_{KMV} and that employed by Moody’s KMV differ in several potentially important respects. In order to gauge how close our methods are, we would like to compare our probability estimates to those calculated by Moody’s KMV. It would be natural to acquire data directly from Moody’s KMV for this purpose, but Moody’s KMV data are prohibitively expensive for us. Fortunately, in November of 2003, Ronald Fink of Moody’s KMV published an article in *CFO Magazine* titled “Ranking America’s top debt issuers by Moody’s KMV Expected Default Frequency.” This magazine article included a table with Moody’s KMV EDF data for one hundred firms. We are able to calculate default probabilities for eighty of the firms listed in the article. We include a comparison of our probability estimates and those of Moody’s KMV in Table 2. Each default probability is computed as of August 2000.

Among the 80 firms for which we have data, the rank correlation between our calculated π_{KMV} and that calculated by Moody’s KMV is 79 percent. The rank correlation between our naive probability and the Moody’s KMV probability is also 79 percent. These high correlations indicate that both of our probability measures do a good job of capturing the information in the probability estimates published by Moody’s KMV. Table 2 also shows that the rank correlation between our (iterated) estimate of firm volatility and that of Moody’s KMV is only 57 percent, while the rank correlation of our naive estimate of firm volatility and the firm volatility published by Moody’s KMV is much higher, at 85 percent. Given the nature of our naive probability estimate, this high correlation is remarkable. Again, this demonstrates that we are able to capture much of the information in Moody’s KMV estimates with our measures.

B. Hazard Model Results

Table 3 contains the results of estimating several Cox proportional hazard models. Models 1 through 3 are univariate hazard models, which explain time-to-default as a function of the KMV-Merton probability, the naive probability, and the log of market equity. While these are relatively simple univariate models, the fact that their explanatory variables vary with time means they are more complicated than they might at first appear. Models 1 through 3 confirm that the KMV-Merton probability, the naive probability and market equity are all extremely significant default predictors. Interestingly, the naive predictor and the KMV-Merton probability, which have similar magnitudes, also have similar coefficients and standard errors. The market value of equity is less significant than either the naive or the KMV-Merton probability. Unreported models that use the log of KMV-Merton distance to default rather than the KMV-Merton probability perform uniformly worse than the results reported.

Model 4 in Table 3 combines the KMV-Merton and the naive probability in one hazard model. Both covariates are very statistically significant, allowing us to conclude that the KMV-Merton is not a sufficient statistic for default probability, or allowing us to reject hypothesis one. Interestingly, both coefficients have similar magnitudes and similar statistical significance, but their significance and magnitude is much smaller in Model 4 than in Models 1 and 2. This reflects the fact that the KMV-Merton and naive probabilities are highly correlated. In fact, in our sample, their correlation coefficient is 0.86. Model 5 similarly combines the KMV-Merton probability and market equity, showing again that we can reject hypothesis one.

Models 6 and 7 include a number of other covariates: the firm's returns over the past year, the log of the firm's debt, the inverse of the firm's equity volatility, and the firm's ratio of net income to total assets. Each of these predictors is statistically significant, making our rejection of hypothesis one quite robust. Interestingly, with all of these predictors included in the hazard model, the KMV-Merton probability is no longer statistically significant, implying that we can reject hypothesis two. The magnitude of the KMV-Merton coefficient is much smaller in Model 6 than it is in Model 4, while its standard error is quite similar. The naive probability retains its statistical significance even though its coefficient drops by approximately one half.

C. Out of Sample Results

Table 4 contains our assessment of the out of sample predictive ability of several variables. To create the table, firms are sorted into deciles each quarter based on a particular forecasting variable. Then the number of defaults that occur in each of the decile groups is tabulated, with the percentage of defaults in the highest probability deciles reported in the table. One advantage of this approach is that the rankings of firms into default probability deciles can be done without estimating actual default probabilities. If our model for translating distances into default into default probabilities is slightly misspecified (in particular, if the normal CDF is not the most appropriate choice), our out of sample results will be unaffected.⁹

Panel A compares the predictions of the KMV-Merton model to the naive model, market equity, and past returns. While the KMV-Merton model probability is able to classify almost 65 percent of defaulting firms in the highest probability decile at the beginning of the quarter in which they default, the naive model is able to classify 65.8 percent of defaulting firms in the top decile. Fully 80.0 percent of defaults occur in the highest π_{KMV} quintile, while 80.1 percent occur in the highest π_{Naive} quintile. It is remarkable that the out of sample performance of π_{KMV} is worse than that of π_{Naive} .

The out of sample performance of both π_{KMV} and π_{Naive} is quite a bit better than simply sorting firms on their market equity. This is consistent with the results of Vassalou and Xing (2003) and indicates that the success of π_{KMV} is not simply reflecting the predictive value of market equity. Apparently, it is quite useful to form a probability measure, by creating a z-score and using a cumulative distribution to calculate the corresponding probability. Given that π_{KMV} does not perform better than π_{Naive} in either hazard models or out of sample forecasts, the probability measure idea behind the KMV-Merton model may be a more valuable innovation than the simultaneous solution of equations (2) and (5).

Simply sorting firms on their excess equity return over the last year has surprisingly good forecasting power, as does sorting firms by their value of net income over total assets. This is consistent with the economic and statistical significance of both of these variables in the hazard model results reported in Table 3. Since the KMV-Merton model has no simple way to capture

⁹A rough calibration of probabilities associated with distance to default rankings can be inferred from the data in Table 4. For example, the probability that firms in the top decile of π_{KMV} will default in the next quarter is equal to the number of defaults occurring in the top decile (1449 * 0.649) divided by one tenth of the number of firm-quarter observations used to create the table (350,662 * 0.1), giving a probability of 2.7 percent.

innovations in past returns or income, it is difficult to believe that π_{KMV} can be a sufficient statistic for default. Any reasonable default prediction model probably needs to include some measure of past returns and net income.

Panel B reports similar forecasting assessments for a shorter time period, from 1991 to 2003. Looking at defaults in this shorter period allows us to examine the out of sample performance of the hazard models reported in Table 3. We estimate Models 6 and 7 from Table 3 each quarter, using data available in that quarter, to define our decile groups. For example, we sort firms in the second quarter of 1995 based on the fitted values of a hazard model estimated with data from 1980 through the first quarter of 1995. The forecasting success of hazard Models 6 and 7 appear in Columns 4 and 5 of the table.

The hazard models assessed in Panel B clearly outperform both π_{KMV} and π_{Naive} . This is not surprising, given that they employ more information in making their forecasts. This again implies that we can easily reject hypothesis one – π_{KMV} is not a sufficient statistic to forecast default. Interestingly, the hazard model that does not include π_{KMV} as a covariate (Model 7) performs substantially better than π_{KMV} , categorizing 76.8 percent of defaults in the highest hazard decile when they default versus 68.8 percent for π_{KMV} . However, the hazard model that includes π_{KMV} (Model 6) performs slightly better than Model 7, categorizing 77.1 percent of defaulting firms in the top hazard decile versus 76.8 percent. While π_{KMV} appears to be making only a marginal contribution to a well-specified hazard model, it appears to be making a small positive contribution. This makes it more difficult to unambiguously reject hypothesis two.

D. Alternative Predictors

In order to assess the importance of various calculations required to generate the KMV-Merton probabilities, we examine the forecasting performance of three alternative probabilities in Table 5. The first of our three measures is a KMV-Merton probability that is calculated under the assumption that the expected return of each firm is the risk-free rate. We examine this predictor, which we denote $\pi_{\text{KMV}}^{\mu=r}$, to determine how important the calculation of μ is for the distance to default in equation (6). Our second predictor, which we denote $\pi_{\text{KMV}}^{\text{simul}}$, is a KMV-Merton probability that is calculated by simultaneously solving equations (2) and (5) rather than following the more complicated iterative procedure described in the text. Our third alternative predictor is a KMV-Merton probability that is calculated with option implied volatility rather than historical equity

volatility. The implied volatility predictor, which we denote $\pi_{\text{KMV}}^{\text{imp}\sigma}$, simultaneously solves equations (2) and (5) instead of using the iterative procedure. Each of our three alternative predictors can be thought of as KMV-Merton probabilities that are calculated with some strong simplifying assumptions. If these predictors perform as well as π_{KMV} then we can conclude that the simplifying assumptions are valid.

It is important to point out that while our sample for $\pi_{\text{KMV}}^{\mu=r}$ and $\pi_{\text{KMV}}^{\text{simul}}$ is the same as the samples described in the rest of the paper, our sample for $\pi_{\text{KMV}}^{\text{imp}\sigma}$ is much smaller, spanning 1996 through 2003 and containing 101,201 firm-months with complete data. We obtain the implied volatility of 30-day at-the-money call options from the IVY Database of Optionmetrics LLC. IVY is a comprehensive database of historical price, implied volatility and sensitivity information for the entire U.S. listed index and equity options market and contains historical data beginning in January 1996. The implied volatilities are calculated by Optionmetrics in accordance with the standard conventions used by participants in the equity option market, using a Cox-Ross-Rubinstein binomial tree model which is iterated until convergence of the model price to the market price of the option.

Table 5 reports summary statistics for each variable, correlations between each of the variables and π_{KMV} , and measures of out of sample prediction accuracy that correspond to those in Table 4. Looking at the correlations in Panel B, each of our alternative predictors is highly correlated with π_{KMV} and with the other predictors in the table. Interestingly, the simultaneously solved $\pi_{\text{KMV}}^{\text{simul}}$ is more correlated with π_{Naive} than π_{KMV} . The probability calculated with implied volatility is less correlated with π_{KMV} than most of the other probabilities.

The out of sample forecast accuracy in Panel C allows us to gauge the relative importance of the iterative procedure, the estimation of μ , and the estimation of equity volatility. As in Table 4, Panel C sorts all firm-quarters by each predictor and then counts the number of defaults that occur among firms in each decile of the predictor. In Panel C, the results for $\pi_{\text{KMV}}^{\mu=r}$ and $\pi_{\text{KMV}}^{\text{simul}}$ in the second and third columns are directly comparable to the results in Panel A of Table 4. However, because of the different sample size, the result for $\pi_{\text{KMV}}^{\text{imp}\sigma}$ are not comparable to any results in Table 4. To provide a performance benchmark for the $\pi_{\text{KMV}}^{\text{imp}\sigma}$ results, the fifth and sixth columns of Table 5 report on the success of π_{KMV} and π_{Naive} using the subset of firms for which implied volatility is available.

The estimation of μ for distance to default (equation 6) is apparently quite important. The probability that sets μ equal to the risk free rate performs substantially worse than π_{KMV} out

of sample, classifying only 60 percent of defaulting firms in the highest probability decile at the beginning of the quarter in which the firms default. π_{KMV} is able to classify almost 65 percent of defaulting firms in the highest decile. Calculating KMV-Merton probabilities with the iterative procedure described in the text is apparently less important. The simultaneously solved $\pi_{\text{KMV}}^{\text{simul}}$ actually has better out of sample predictive performance than the iteratively solved π_{KMV} . This is consistent with the relative success of our naive probability in Tables 3 and 4. Finally, using implied equity volatility rather than estimated equity volatility in our probability improves out of sample performance substantially. However, given that there are only 88 defaults to forecast in our sample of firms with corresponding options contracts, it is difficult to apply this finding to the broader sample of firms.

E. CDS Spread Regressions

Our previous results demonstrate that while π_{KMV} appears to be a useful quantity for forecasting defaults, it is not a sufficient statistic for the purpose of forecasting. Our next two sets of results examine whether π_{KMV} is an important explanatory variable for pricing credit-sensitive securities. First we examine regressions of the implied probability of default from credit default swaps on π_{KMV} and several alternatives. Bond yield spread regressions are our final set of results.

CDS default probability regressions are reported in Table 6. We obtain the data on CDS spreads from www.credittrade.com for the period December 1998 to July 2003. From this source, we are able to collect 3833 firm-months of CDS spread observations. We calculate the probability that a firm defaults in the next year, π_{CDS} , according to the algorithm described in section I.F. We then regress π_{CDS} on π_{KMV} , π_{Naive} and all the other variables in the hazard models described in Table 3. Since π_{CDS} is a risk-neutral probability of default, we also regress π_{CDS} on $\pi_{\text{KMV}}^{\mu=r}$, which is the risk-neutral probability of default implied by the KMV-Merton model. We do not report these regression results because they are qualitatively similar but uniformly weaker than our results with π_{KMV} . Instead, we report the correlation coefficient of π_{CDS} and $\pi_{\text{KMV}}^{\mu=r}$ in Panel B of Table 3.

If π_{KMV} is a well-calibrated and accurate probability of default, then the coefficient on π_{KMV} in these regressions should be greater than or equal to one. If there is no risk premium for default risk then the coefficient should be exactly one. If there is a risk premium for default risk, then the risk-neutral default probability will generally be larger than the true probability, and the coefficient on the true probability will be greater than one. If the coefficient on π_{KMV} is less than one, there

must be some information that is captured in π_{CDS} but absent from π_{KMV} .

Looking first at the correlations between π_{CDS} and our probability measures, we see that π_{Naive} is much more correlated with π_{CDS} (at 51 percent) than π_{KMV} (at 32 percent). π_{CDS} and $\pi_{\text{KMV}}^{\mu=r}$ have a surprisingly low correlation (16 percent). Turning to the regressions, the coefficient on π_{KMV} in the univariate regression of Model 1 is just 0.05, much lower than the predicted lower bound of one. The same univariate regression with π_{Naive} replacing π_{KMV} yields a slightly more reasonable coefficient of 0.13. Combining π_{KMV} and π_{Naive} in one model (Model 3) results in π_{KMV} losing all significance and the coefficient and significance of π_{Naive} changing very little. Including all the other predictive variables makes π_{KMV} statistically significant again, though the coefficient on π_{KMV} remains much smaller than either one or the coefficient on π_{Naive} . The CDS spread regression results in Table 6 show that, for the purpose of pricing credit-sensitive securities, the naive probability estimate performs at least as well as the KMV-Merton probability.

F. Bond Spread Results

Our final set of results are regressions of bond yields on our default probabilities. Before discussing our regression results, we describe the sample used to estimate the regressions. Summary statistics for the bond yield sample appear in Panel A of Table 7.

Our bond data are extracted from the Lehman Brothers Fixed Income Database distributed by Warga (1998). This database contains monthly price, accrued interest, and return data on all corporate and government bonds from 1971-1997. We use the data from the 1980-1997 period to be consistent with the default prediction sample. This is the same database used by Elton et al.(2001) to explain the rate spread on corporate bonds. In addition, the database contains descriptive data on bonds, including coupons, ratings, and callability. A subset of the data in the Warga database is used in this study. First, all bonds that were matrix priced rather than trader priced are eliminated from the sample¹⁰. Employing matrix prices might mean that all our analysis uncovers is the rule used to matrix-price bonds rather than the economic influences at work in the market. Eliminating matrix-priced bonds leaves us with a set of prices based on dealer quotes. This is the same type of data as that contained in the standard academic source of government bond data: the CRSP

¹⁰For actively traded bonds, dealers quote a price based on recent trades of the bond. Bonds for which a dealer did not supply a price have prices determined by a rule of thumb relating the characteristics of the bond to dealer-priced bonds. These rules of thumb tend to change very slowly over time and do not respond to changes in market conditions.

government bond file. Next, we eliminate all bonds with special features that would result in their being priced differently. This means we eliminate all bonds with options (e.g. callable bonds or bonds with a sinking fund), all corporate floating rate debt, bonds with an odd frequency of coupon payments, and inflation-indexed bonds. In addition, we eliminate all bonds not included in the Lehman Brothers bond indexes, because researchers in charge of the database at Lehman Brothers indicate that the care in preparing the data was much less for bonds not included in their indexes. This also results in eliminating data for all bonds with a maturity of less than one year. This exclusion is also consistent with our estimates of π_{KMV} and π_{Naive} , which are based on a one year forecasting horizon. Finally, we also remove AAA (Moody’s rating Aaa) bonds because the data for these bonds appear problematic. Both Elton et al. (2001) and Campbell and Taksler (2003) exclude AAA bonds from their analysis for this reason. We are finally left with 62,584 bond-months with complete data in our sample.

There are a number of extreme values among the observations of each variable constructed from the Warga data. To ensure that statistical results are not heavily influenced by outliers, we set all observations higher than the 99th percentile of each variable to that value. All values lower than the first percentile of each variable are winsorized in the same manner. The minimum and maximum numbers reported in Panel A for the bond are calculated after winsorization.¹¹

We compute the spread on the corporate bond as the difference between the yield to maturity on a corporate bond in that particular month and the yield to maturity on a government bond of the closest maturity in the same month. For the benchmark treasuries, we use the CRSP fixed term indices, which provide monthly yield data on notes and bonds of 1, 2, 5, 6, 10, 20 and 30 target years to maturity. We assume that each quoted price in the Warga data is at the end of the month when the CRSP indices are published, but this should have little impact on the measured spreads. As can be seen from Panel A, the average spread is about 108 basis points over this sample period (1980-1997), similar in magnitude to spreads reported in the other studies. We find that the magnitude of π_{KMV} and π_{Naive} are smaller than the values reported in Table 1, suggesting that firms that have issuances in the bond market are better credit risks. The average maturity outstanding for the bonds in our sample is about 10 years and the Coupon rate is around 8.3%.

In Panel B of Table 7 we report the results of regressing bond yield spreads on a number of

¹¹As in the summary statistics in Table 1, we do not winsorize the expected default frequency measures from the Merton Model and the naive alternative, since these are naturally bounded between 0 and 1.

explanatory variables. Looking at the results, it appears that both π_{KMV} and π_{Naive} are significantly related to bond yield spreads. However, looking again at the results it quickly becomes clear that while spreads are correlated with both π_{KMV} and π_{Naive} , the coefficients on these default probabilities are too low. For example, given the coefficient of 0.5 for π_{Naive} in Model 1, if π_{Naive} increased from zero to five percent, the expected bond yield would increase by just 2.5 basis points. The magnitude of the coefficients can be explained by the fact that bond rating dummies are included in these regressions, and bond ratings capture a large fraction of the variation in spreads. The regression coefficients must be interpreted as capturing the explanatory power of our probability measures conditional on being in a particular ratings class.

In univariate regressions, the magnitude and statistical significance of π_{KMV} is much smaller than that of π_{Naive} . Combining both π_{KMV} and π_{Naive} in one spread regression makes the coefficient of π_{KMV} become insignificant. When other explanatory variables are included in the regression, the coefficient on π_{Naive} loses some of its significance but remains statistically distinguishable from zero. π_{KMV} is less significant, both statistically and economically.

Overall, the regressions indicate that π_{KMV} is not strongly related to bond yield spreads after conditioning on bond ratings. This is consistent with the hazard model and out of sample results discussed previously.

IV. Conclusion

We examine the accuracy and the contribution of the KMV-Merton default forecasting model. Looking at hazard models that forecast default, the KMV-Merton model does not appear to be a sufficient statistic for default. It appears to be possible to construct an accurate default forecasting model without considering the iterated KMV-Merton default probability. The naive probability that we propose, which captures both the functional form and the same basic inputs of the KMV-Merton probability, performs surprisingly well. Looking at out of sample forecasting ability, it is fairly simple to construct a model that outperforms the KMV-Merton model without using the KMV-Merton default probability as an explanatory variable. However, hazard models that use the KMV-Merton probability with other covariates have slightly better out of sample performance than models which omit the KMV-Merton probability. Looking at CDS implied default probability regressions and bond yield spread regressions, the KMV-Merton probability does not appear to

be a significant predictor of either quantity when our naive probability, agency ratings and other explanatory variables are accounted for. We conclude that the KMV-Merton probability is a marginally useful default forecaster, but it is not a sufficient statistic for default.

References

- Basel Committee on Banking Supervision, 1999, *Credit Risk Modelling: Current Practices and Applications*.
- Berndt, Antje, Rohan Douglas, Darrell Duffie, Mark Ferguson, and David Schranz, 2004, Measuring default risk premia from default swap rates and EDFs, working paper, Stanford University.
- British Bankers' Association, 2002, *BBA Credit Derivatives Report 2001/2002*.
- Campbell, John Y., Jens Hilscher, and Jan Szilagyi, 2004, In search of distress risk, working paper, Harvard University.
- Campbell, John Y., and Glen B. Taksler, 2003, Equity volatility and corporate bond yields, *Journal of Finance*, 58, 2321-2349.
- Chava, Sudheer, and Robert Jarrow, 2004, Bankruptcy prediction with industry effects, *Review of Finance*, forthcoming.
- Cox, David R. and D. Oakes, 1984, *Analysis of Survival Data*, Chapman and Hall, New York.
- Crosbie, Peter J. and Jeffrey R. Bohn, 2001, *Modeling Default Risk* (KMV LLC).
- Du, Yu, and Wulin Suo, 2004, Assessing credit quality from equity markets: Is a structural approach a better approach? working paper, Queen's University.
- Duffie, Darrell, and Kenneth J. Singleton, 2003, *Credit Risk: Pricing, Measurement, and Management*, Princeton University Press, Princeton, NJ.
- Duffie, Darrell, and Ke Wang, 2004, Multi-Period corporate failure prediction with stochastic covariates, working paper, Stanford University.
- Elton, Edwin J., Martin J. Gruber, Deepak Agrawal, and Christopher Mann, 2001, Explaining the rate spread on corporate bonds, *Journal of Finance* 56, 247-277.
- Falkenstein, Eric, and Andrew Boral, 2001, Some empirical results on the Merton model, *Risk Professional*, April 2001.
- Hillegeist, S. A., E. K. Keating, D. P. Cram, and K. G. Lundstedt, 2004, Assessing the probability of bankruptcy *Review of Accounting Studies* 5-34.
- Hull, J., M. Predescu, and A. White, 2004, The relationship between credit default swap spreads, bond yields and credit rating announcements, *Journal of Banking and Finance*, forthcoming.

- Jones, E., S. Mason and E. Rosenfeld, 1984, Contingent claims analysis of corporate capital structures: an empirical investigation, *Journal of Finance* 39, 611-625.
- Kealhofer, Stephen, and Matthew Kurbat, 2001, *The Default Prediction Power of the Merton Approach, Relative to Debt Ratings and Accounting Variables* (KMV LLC).
- Longstaff, Francis A., Sanjay Mithal, and Eric Neis, 2004, Corporate yield spreads: Default risk or liquidity? New evidence from the credit-default swap market, forthcoming, *Journal of Finance*.
- Merton, Robert C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449-470.
- Saunders, Anthony and Linda Allen, 2002, *Credit Risk Measurement: New Approaches to Value at Risk and Other Paradigms*, John Wiley and Sons, New York.
- Shumway, Tyler, 2001, Forecasting bankruptcy more accurately: a simple hazard model, *Journal of Business* 74, 101-124.
- Sobehart, Jorge R., and Sean C. Keenan, 1999, *An Introduction to Market-Based Credit Analysis* (Moody's Investors Services).
- Sobehart, Jorge R., and Sean C. Keenan, 2002, Hybrid contingent claims models: A practical approach to modelling default risk, in *Credit Rating: Methodologies, Rationale, and Default Risk*, edited by Michael Ong and published by Risk Books, 125-145.
- Sobehart, Jorge R., and Sean C. Keenan, 2002, The need for hybrid models, *Risk*, February 2002, 73-77.
- Sobehart, Jorge R., and Roger M. Stein, 2000, *Moody's Public firm Risk Model: a Hybrid Approach to Modeling Short Term Default Risk* (Moody's Investors Services).
- Stein, Roger M., 2000, *Evidence on the Incompleteness of Merton-type Structural Models for Default Prediction* (Moody's Investors Services).
- Vassalou, Maria and Yuhang Xing, 2003, Default risk in equity returns, *Journal of Finance* forthcoming.
- Warga, A., 1998, Fixed income data base, working paper, University of Houston.

Table 1: Summary Statistics

Table 1 reports summary statistics for all the variables used in the KMV-Merton Model and the hazard models. E is the market value of equity in millions of dollars and is taken from CRSP as the product of share price at the end of the month and the number of shares outstanding. F is the face value of debt in millions of dollars (computed as Compustat item 45 + 0.5 * Compustat item 51). r is the risk-free rate measured as the 3 month Treasury-bill rate. The past returns variable, $r_{it-1} - r_{mt-1}$, is the difference between the prior year return of the firm and the return on the CRSP value weighted index during the same period, and NI/TA is the firm's ratio of net income to total assets. V is the market value of firm assets in millions of dollars, σ_V is the asset volatility measured in percentage per annum, and μ is the expected return on the firm's assets. All three of these variables are generated as the result of solving the KMV-Merton model for each firm-month in the sample using the iterative procedure described in the text. π_{KMV} is the expected default frequency in percent and is given by equation (7). Naive σ_V is calculated by equation (10), and the firm's equity return from the previous year, r_{it-1} , is used as a proxy for the firm's expected asset return to calculate the naive probability of default, π_{Naive} . Our naive probability, π_{Naive} , is calculated according to equation (13). All variables except the default probabilities are winsorized at the first and ninety-ninth percentiles. Our sample spans 1980 through 2003, containing 1,016,552 firm-months with complete data.

Panel A: Means, Standard Deviations, and Quartiles

Variable	Mean	Std.Dev.	Min	Quantiles			
				0.25	Mdn	0.75	Max
E	808.80	2453.15	1.21	18.56	76.52	394.83	17534.72
F	229.92	729.66	0.02	2.67	15.56	96.65	5175.50
r (%)	6.46	2.82	1.01	4.85	5.85	7.89	16.72
$r_{it-1} - r_{mt-1}$ (%)	-8.69	63.02	-99.89	-46.79	-14.21	16.94	272.00
NI/TA	-1.08	6.99	-41.13	-0.94	0.73	1.85	7.83
V	1072.33	3228.60	1.52	26.43	105.24	530.12	22949.32
σ_V (%)	56.00	36.83	10.03	30.41	46.32	70.61	230.19
μ (%)	3.25	57.17	-253.58	-21.72	4.36	29.34	210.37
π_{KMV} (%)	10.95	23.32	0.00	0.00	0.01	6.41	100.00
Naive σ_V (%)	50.67	30.97	10.48	28.17	42.29	64.70	162.70
r_{it-1} (%)	13.75	82.07	-85.45	-27.01	2.27	34.13	294.94
π_{Naive} (%)	8.95	20.57	0.00	0.00	0.00	3.46	100.00

Panel B: Correlations

Corr (σ_V , Naive σ_V)	=	0.8748
Corr (π_{KMV} , π_{Naive})	=	0.8642

Table 2: Comparison with Moody’s KMV EDF

Table 2 compares the expected default frequency computed by Moody’s KMV corporation and the methods used in this paper. We obtain the data for Moody’s KMV EDF and asset volatility for 80 firms for August 2000 from the article ‘Ranking America’s top debt issuers by Moody’s KMV Expected Default Frequency’, by Ronald Fink, in *CFO Magazine*, November 2003. The second column of the table provides the rank correlations between the various measures listed in the first column. All correlations are significant at the 0.1% level or lower.

Correlation	Estimate
Rank Corr(Moody’s π_{KMV} , Our π_{KMV})	0.788
Rank Corr(Moody’s π_{KMV} , Our π_{Naive})	0.786
Rank Corr(Moody’s σ_V , Our σ_V)	0.574
Rank Corr(Moody’s σ_V , Our Naive σ_V)	0.853

Table 3: Hazard Model Estimates

Table 3 reports the estimates of several Cox proportional hazard models with time-varying covariates. There are 15,018 firms and 1,449 defaults in the sample. π_{KMV} is the KMV-Merton probability of default, π_{Naive} is our naive alternative, $\ln(E)$ and $\ln(F)$ are the natural logarithms of market equity and face value of debt respectively. $1/\sigma_E$ is the inverse of equity volatility, measured with daily data over the previous year, $r_{it-1} - r_{mt-1}$ is the stock's return over the previous year minus the market's return over the same period, and NI/TA is the firm's ratio of net income to total assets. A positive coefficient on a particular variable implies that the hazard rate is increasing in that variable, or that the expected time to default is decreasing in that variable. This table shows that the KMV-Merton default probability is not a sufficient statistic for forecasting default and that the naive default probability measure is at least as important as the KMV-Merton measure. Standard errors are in parentheses (***) Significant at one percent level, ** Significant at five percent level, * Significant at 10 percent level).

Dependent Variable: Time to Default							
Variable	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7
π_{KMV}	3.635*** (0.068)			1.697*** (0.142)	3.272*** (0.077)	0.230 (0.164)	
π_{Naive}		4.011*** (0.067)		2.472*** (0.147)		1.366*** (0.178)	1.526*** (0.138)
$\ln(E)$			-0.472*** (0.014)		-0.164*** (0.015)	-0.247*** (0.024)	-0.255*** (0.023)
$\ln(F)$						0.263*** (0.020)	0.269*** (0.020)
$1/\sigma_E$						-0.506*** (0.047)	-0.518*** (0.046)
$r_{it-1} - r_{mt-1}$						-0.819*** (0.081)	-0.834*** (0.080)
NI/TA						-0.044*** (0.002)	-0.044*** (0.002)

Table 4: Out of Sample Forecasts

Table 4 reports on the success of various forecasting quantities by sorting firms each quarter by forecast and counting the fraction of defaults that correspond with each decile of the forecast variable. Panel A examines accuracy over the entire period for which we have data, from 1980 to 2003. There are 350,662 firm-quarters in our sample, with 1449 defaults. π_{KMV} is the KMV-Merton probability of default, π_{Naive} is our naive alternative, E is market equity, $r_{it-1} - r_{mt-1}$ is the stock's return over the previous year minus the market's return over the same period, and NI/TA is the firm's ratio of net income to total assets. Panel B only considers defaults from 1991 to 2003, and it includes the fitted values of two hazard models (models 6 and 7 from Table 3) as predictors. These models are estimated each quarter using all available data in each quarter, and the resulting coefficients are used to form the predictors assessed in columns 4 and 5 of Panel B. Again, this table shows that the naive variable works slightly better than the KMV-Merton quantity. It also shows that a reduced form model that neglects to consider the KMV-Merton probability can perform better than π_{KMV} out of sample.

Panel A: 1980 - 2003
350,662 firm-quarters, 1449 defaults

Decile	π_{KMV}	π_{Naive}	E	$r_{it-1} - r_{mt-1}$	NI/TA
1	64.9	65.8	35.7	44.4	46.8
2	15.1	14.3	17.5	25.1	23.8
3	6.0	6.7	14.3	9.2	10.6
4	4.6	4.1	9.1	5.4	5.9
5	2.9	2.4	6.1	2.9	4.2
6 - 10	6.5	6.7	17.3	13.0	8.7

Panel B: 1991 - 2003
226,604 firm-quarters, 842 defaults

Decile	π_{KMV}	π_{Naive}	Model 6	Model 7
1	68.8	70.3	77.1	76.8
2	15.3	12.6	10.5	10.5
3	5.1	6.2	4.9	4.9
4	3.0	3.4	1.8	2.1
5	1.9	1.8	1.2	1.0
6 - 10	5.9	5.7	4.5	4.7

Table 5: Alternative Predictors

Table 5 reports on the success of three alternative predictors, or predictors that calculate KMV-Merton default probabilities in alternative ways. One predictor, $\pi_{\text{KMV}}^{\mu=r}$, is calculated in exactly the same manner as π_{KMV} , except that the expected return on assets used for π_{KMV} is replaced by the risk-free rate, r . A second alternative predictor, $\pi_{\text{KMV}}^{\text{simul}}$, is calculated by simultaneously solving equations (2) and (5). This predictor avoids the iterative procedure in the text, estimating equity volatility with one year of historical returns data and using r as the expected return on assets. The third alternative predictor, $\pi_{\text{KMV}}^{\text{imp}\sigma}$, uses the option-implied volatility of firm equity (implied σ_E) to simultaneously solve equations (2) and (5). Our sample for $\pi_{\text{KMV}}^{\mu=r}$ and $\pi_{\text{KMV}}^{\text{simul}}$ is the same as the samples described in Table 1, including 1,016,552 firm-months from 1980 through 2003. Our sample for $\pi_{\text{KMV}}^{\text{imp}\sigma}$ spans 1996 through 2003, containing 101,201 firm-months with complete data. Panel A reports summary statistics on our alternative predictors, Panel B reports correlations between all of our predictors, and Panel C reports on the out of sample predictive success of our alternatives. As in Table 4, Panel C sorts all firm-quarters by each predictor and then counts the number of defaults that occur among firms in each decile of the predictor. In Panel C, the results for $\pi_{\text{KMV}}^{\mu=r}$ and $\pi_{\text{KMV}}^{\text{simul}}$ in the second and third columns are directly comparable to the results in Panel A of Table 4. However, because of the different sample size, the result for $\pi_{\text{KMV}}^{\text{imp}\sigma}$ are not comparable to any results in Table 4. To provide a performance benchmark for the $\pi_{\text{KMV}}^{\text{imp}\sigma}$ results, the fifth and sixth columns of Table 5 report on the success of π_{KMV} and π_{Naive} using the subset of firms for which implied volatility is available.

Panel A: Summary Statistics (in percent)

Variable	Mean	Std.Dev.	Min	Quantiles			
				0.25	Mdn	0.75	Max
$\pi_{\text{KMV}}^{\mu=r}$	7.71	17.97	0.00	0.00	0.01	3.77	100.00
$\pi_{\text{KMV}}^{\text{simul}}$	8.13	20.76	0.00	0.00	0.00	1.48	100.00
Implied σ_E	58.47	26.54	4.01	39.25	52.48	72.27	500.00
$\pi_{\text{KMV}}^{\text{imp}\sigma}$	4.11	14.70	0.00	0.00	0.00	0.10	100.00

Panel B: Correlation Matrix

	π_{KMV}	π_{Naive}	$\pi_{\text{KMV}}^{\mu=r}$	$\pi_{\text{KMV}}^{\text{simul}}$
π_{Naive}	0.8642			
$\pi_{\text{KMV}}^{\mu=r}$	0.8575	0.7486		
$\pi_{\text{KMV}}^{\text{simul}}$	0.8338	0.9755	0.7220	
$\pi_{\text{KMV}}^{\text{imp}\sigma}$	0.6858	0.9624	0.4102	0.6259

Panel C: Out of Sample Forecasts

Decile	$\pi_{\text{KMV}}^{\mu=r}$	$\pi_{\text{KMV}}^{\text{simul}}$	$\pi_{\text{KMV}}^{\text{imp}\sigma}$	π_{KMV}	π_{Naive}
1	60.0	65.1	84.1	80.7	83.0
2	17.7	15.0	8.0	9.1	9.1
3	8.0	7.7	4.6	3.4	5.7
4	4.1	3.4	0.0	5.7	1.1
5	3.4	3.2	1.1	0.0	0.0
6 - 10	6.8	5.6	2.2	1.1	1.1
Defaults	1,449	1,449	88	88	88
Firm-Quarters	350,662	350,662	36,274	36,274	36,274

Table 6: CDS Spread Regressions

Table 6 reports on a comparison of the probability of default implied by credit default swap (CDS) spreads with π_{KMV} and π_{Naive} . We obtain the data on CDS spreads from www.creditrade.com for the period December 1998 to July 2003. CDS spread is the credit default swap spread in basis points. π_{CDS} is the probability of default backed out from the CDS spread. All other measures are described in table 1 of the paper. The total number of firm-month observations is 3,833. Panel A reports summary statistics, Panel B reports correlations, and Panel C reports the results of regressing the default probability implied by the CDS spread on π_{KMV} , π_{Naive} and several other predictive variables. In Panel C, standard errors are shown in parentheses (***) Significant at one percent level, ** Significant at five percent level, * Significant at 10 percent level).

Panel A: Summary Statistics (in percent)

Variable	Mean	Std.Dev.	Min	Quantiles			
				0.25	Mdn	0.75	Max
CDS Spread	165.89	170.49	9.50	60.83	100.00	204.30	1650.00
π_{CDS}	3.36	3.33	0.20	1.26	2.08	4.19	29.59
π_{KMV}	8.39	20.98	0.01	0.01	0.02	2.75	100.00
π_{Naive}	4.55	13.54	0.01	0.01	0.01	0.73	99.79

Panel B: Correlations

Corr(π_{CDS} , $\pi_{\text{KMV}}^{\mu=r}$)	0.1594
Corr(π_{CDS} , π_{KMV})	0.3150
Corr(π_{CDS} , π_{Naive})	0.5090

Panel C: Regressions

Dependent Variable: π_{CDS}

Variable	Model 1	Model 2	Model 3	Model 4
Const.	.03*** (.0005)	.03*** (.0004)	.03*** (.0004)	.16*** (.007)
π_{KMV}	.05*** (.004)		-0.001 (.003)	.009** (.004)
π_{naive}		.13*** (.007)	.13*** (.008)	.06*** (.008)
$\ln(E)$				-.007*** (.0005)
$\ln(F)$.003*** (.0006)
$1/\sigma_E$				-0.014*** (.0007)
r_{it-1}				-0.0012 (.001)
Obs.	3833	3833	3833	3833
R^2	0.10	0.26	0.26	0.40

Table 7: Bond Yield Spread Regressions

Table 7 reports the results of bond yield spread regressions. Spread is the difference between the yield to maturity on the bond and the yield of the closest maturity treasury in basis points. σ_E is the standard deviation of equity returns. Maturity is the remaining time to maturity in years of the bonds. Coupon is the coupon rate on the bond issue. r is the risk free rate measured as the 3 month t-bill rate. Amount is the dollar amount of the bond issue. π_{KMV} is the expected default frequency in percent, given by equation (7), and π_{Naive} is the corresponding naive default probability measure given by equation (13). All observations except the default frequency measures are winsorized at the first and ninety-ninth percentiles. The data span 1980 through 1997 and there are 61,776 bond-months with complete data in our sample. Panel A reports summary statistics for the sample used in the regressions and Panel B reports the regression results. In addition to the variables described in Panel A, all regressions have year, rating and 1 digit sic code dummies. Heteroscedasticity consistent standard errors are shown in parentheses (***) Significant at one percent level, ** Significant at five percent level, * Significant at 10 percent level). This table shows that the naive default probability can replace the KMV-Merton probability in explaining bond yields.

Panel A: Summary Statistics

Variable	Mean	Std.Dev.	Min	Quantiles			
				0.25	Mdn	0.75	Max
Spread (bp)	108.09	69.43	28.11	67.53	90.31	121.75	605.91
σ_E (%)	27.67	8.73	6.60	22.34	26.36	31.06	250.38
Maturity	10.32	8.20	1.00	4.59	7.79	12.63	39.25
Amount	190,000	120,000	15,700	100,000	150,000	250,000	750,000
r (%)	5.50	1.39	3.18	4.94	5.54	5.87	16.72
Coupon (%)	8.36	1.46	4.50	7.25	8.38	9.38	14.25
π_{KMV} (%)	9.42	27.49	0.00	0.00	0.00	0.00	100.00
π_{Naive} (%)	2.02	10.71	0.00	0.00	0.00	0.00	99.49

Table 7: Bond Yield Spread Regressions (continued)

Panel B: Regressions						
Dependent Variable: Bond Yield Spread						
Variable	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Const.	125.77*** (16.55)	118.27*** (16.59)	126.2*** (16.59)	153.11*** (17.43)	147.28*** (17.36)	153.8*** (17.46)
σ_E	.86*** (.05)	.95*** (.05)	.86*** (.05)	.79*** (.06)	.87*** (.06)	.75*** (.06)
Maturity	1.49*** (.02)	1.5*** (.02)	1.49*** (.02)	1.5*** (.02)	1.51*** (.02)	1.5*** (.02)
Ln(Amount)	-5.78*** (.38)	-5.92*** (.39)	-5.81*** (.39)	-8.13*** (.49)	-8.27*** (.49)	-8.01*** (.49)
r	-.46 (.34)	.23 (.34)	-.46 (.34)	-.61* (.37)	.07 (.37)	-.60 (.37)
Coupon	3.31*** (.2)	3.31*** (.21)	3.31*** (.2)	3.49*** (.22)	3.53*** (.22)	3.51*** (.22)
Coverage < 5				-1.4* (.85)	-3.1*** (.91)	-2.12** (.89)
5 <= Coverage < 10				-6.37*** (.71)	-7.37*** (.73)	-6.65*** (.73)
10 <= Coverage < 20				-1.79*** (.68)	-2.13*** (.69)	-1.87*** (.68)
Operating Income to Sales				-36.11*** (4.46)	-40.93*** (4.87)	-38.78*** (4.85)
Long Term Debt to Assets				12.11*** (2.57)	20.08*** (2.83)	16.29*** (2.78)
Total Debt to Capitalization				8.88*** (1.62)	5.82*** (1.59)	4.26*** (1.57)
π_{Naive}	.5*** (.03)		.49*** (.03)	.63*** (.04)		.6*** (.04)
π_{KMV}		.08*** (.008)	.007 (.007)		.19*** (.02)	.1*** (.01)
Rating Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Industry Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	61776	61776	61776	51831	51831	51831
R^2	0.70	0.70	0.70	0.72	0.71	0.72