Of Moody's and Merton: a structural model of bond rating transitions[1]

Michael Gordy
Board of Governors of the Federal Reserve System
Mail Stop 153, Washington, DC 20551 USA
202-452-3705
E-mail: Michael.Gordy@frb.gov

Erik Heitfield
Board of Governors of the Federal Reserve System
Mail Stop 153, Washington, DC 20551 USA
202-452-2613
E-mail: Erik.Heitfield@frb.gov

June 4, 2001

[1]The opinions expressed here are those of the authors and do not reflect the views of the Board of Governors or its staff.

1 Introduction

In their many variants on the basic framework of Merton (1974), structural models of credit risk rest on a more or less literal interpretation of the borrower's balance sheet. The firm's fixed liabilities constitute a barrier point for the value of its assets. If assets drop below that barrier, the firm is unable to support its debt and therefore defaults. Assuming the current asset value and fixed liabilities are observable, we need only specify the stochastic behavior of assets and debt issuance to determine the probability of firm default at any given horizon. In pricing applications of these models, one is interested in the stochastic processes under the risk-neutral measure. Risk management applications, which are the focus of this paper, require specification under the natural measure.

Once a model is specified, its unknown parameters may be calibrated to observed data by either direct or indirect methods. The direct approach requires that one collect detailed information on an obligor's balance sheet in order to estimate its fixed liabilities, which are generally assumed to be non-stochastic. The obligor's capacity to carry these liabilities depends on the market value of its assets, which cannot be directly observed. However, by treating equity as a put option on underlying assets, one can use observed equity prices and volatility to recover the current value and volatility of the of the obligor's assets. This procedure depends strongly on the assumed distribution for the asset return process. In practice, log-normality is nearly always imposed.[1]

Under the indirect approach, one starts with agency ratings of the type issued by S&P and Moody's. An obligor's current rating is taken to be a sufficient statistic for some structural measure of its credit quality. In credit risk management applications, it is generally assumed that obligors in the same rating grade share the same *distance to default*. We can think of the distance to default as a measure of an obligor's leverage relative to the volatility of its asset values. As the value an obligor's assets changes over time, its distance to default changes as well. If assets fall below the value of fixed liabilities the distance to default drops below zero, and the obligor becomes insolvent. Given assumptions about the asset return process, an obligor's

[1] The KMV model is a partial exception. Log-normality is assumed for the purpose of backing out the asset return process from the equity return process, but is replaced by an empirical distribution for the purpose of estimating a default probability. The resulting model is not entirely self-consistent, but offers much improved empirical fit over a strict log-normal model

distance to default is all that is needed to determine its default probability at a fixed horizon date. By examining historical patterns of default for each rating grade, one can estimate unknown parameters of the return distribution process as well as the distance to default associated with each grade.

A virtue of the indirect approach to modeling credit risk is that it leverages the comparative advantage of the ratings agencies in interpreting balance sheet information. Extracting a default barrier from accounting statements is not only time-intensive, but may require significant exercise of judgment in handling complex liability structures. The agencies have decades of experience at balance-sheet analysis, as well as access to private information that is not available in public filings. The potential drawback to the indirect method is that it relies on rather strong assumptions about the rating agencies' objective functions and methodologies. It is widely acknowledged that agency ratings can be slow to respond to new information. Less widely recognized is that ratings have traditionally had only a qualitative interpretation. Even now, the agency's judgment on a firm's one-year default probability is only one factor considered in rating assignment. Rating agencies may also consider the ability of the firm to withstand the trough of a business cycle (i.e., so-called "through the cycle" rating), as well as the loss a senior unsecured claimant is likely to experience in the event of default.

To explore the efficacy of the indirect approach, we begin with a benchmark rating assignment model. We assume that the agency observes the value of an obligor's distance to default at both the as-of date and at the horizon, and that the change from as-of date to horizon is normally distributed with unit variance.[2] By construction, firms within a given rating class then have a common distance to default at the as-of date and a common distribution over outcomes at the horizon. In the spirit of CreditMetrics and KMV Portfolio Manager, we introduce correlations in asset values across firms by expressing the change in distance to default as the sum of a systematic risk factor that is common to all obligors and an idiosyncratic risk factor that is unique to each obligor. The degree of correlation across obligors depends on the relative importance of the systematic risk factor.

This simple rating assignment model contains three types of parameters that must

[2] This assumption of unit variance is not restrictive. It follows directly for the definition of distance to default as the volatility normalized log inverse leverage ratio. However, as we shall see presently, the normality assumption is quite important.

be calibrated to observed data. A vector of *transition thresholds* partitions the range of all possible realized default distances into intervals. Each interval is associated with a particular final rating grade. A vector of *initial default distances* describes the mean of the distribution of final distances for each grade. Finally, *factor loadings* determine the correlation across firms in rating migrations and defaults. All three sets of parameters can be estimated using annual data on the number of firms in each rating grade and their final grade (or default status) at the end of one year.

Comparing the matrix of one-year rating transition probabilities implied by the benchmark model with a model-free transition matrix estimated directly from rating assignment data provides a test of the benchmark model's first-moment predictions. We find that the model does a good job of fitting the probability that a firm will retain its initial rating after one year, but it does a very poor job of fitting the probability that a firm will move two or more grades away from its initial rating.[3]

Our goal in this paper is to identify ways of relaxing the strong assumptions underlying both the structural default model and the benchmark rating assignment model to better capture observed features of rating transition data. We focus on generalizing the model in ways that are (1) parsimonious, and (2) consistent with an underlying structural model of firm and rating agency behavior.

We proceed along two lines. First, we relax the assumption that the change in an obligor's distance to default over the assessment horizon is normally distributed. An extensive body of research on stock price data suggests that equity returns, and by implication asset returns are most appropriately modeled as having been drawn from thick-tailed distributions (Mittnik and Rachev 1999). This is consistent with our finding that the benchmark model places unrealistically low weights on large rating changes. The normality assumption can be relaxed while preserving the mathematical structure of the Merton model by assuming that the distribution of the change in default distance is a member of the alpha-stable family. The alpha-stable distribution family includes the normal distribution as a limiting case, but it also includes a range of thicker-tailed distributions. Transition matrices estimated from an alpha-stable model do a better job of predicting large rating changes than those estimated from the normal model, however, they predict unrealistically high default probabilities for

[3]CreditMetrics uses a framework similar to our own but treats transitions from each initial rating grade as a distinct random process. This model generates very reasonable rating transition probabilities, but, like the KMV model, it cannot be fully reconciled with the Merton model from which it is derived.

high-grade obligors.

Our second approach to generalizing the benchmark model focuses on the information set available to the rating agencies. A key assumption of the benchmark model is that the rating agency observe an obligor's distance to default at the as-of date and the horizon date without error. If this assumption is relaxed, then at any point in time a firm may be much closer to the default boundary than is implied by its grade. The presence of incomplete information in the rating process may therefore help explain why the benchmark model does a poor job of predicting default probabilities for higher-rated firms. We find that a straightforward extension of the benchmark model along the lines of Duffie and Lando (1999) significantly improves predicted default probabilities without dramatically increasing the number of parameters that must be estimated.

2 Data

An important practical advantage of the indirect approach to modeling credit risk is that its data requirements are relatively modest. This study makes use of annual aggregate date on ratings by Moody's Investment Service. Moody's rates primarily large corporate and sovereign bond issuers. To ensure a relatively homogeneous population, we limit our sample to include only US, non-financial, non-sovereign obligors.[4]

Moody's uses seven rating grades ranging from 'Aaa' (high-quality investment grade) to 'C' (speculative grade). In 1983 Moody's refined its rating system by appending '+' or '-' to each major grade. Transitions rates in 1983 were far higher than those observed in any other year, suggesting that Moody's may have adjusted its rating criteria at the same time it added detail to its rating scale. For this reason we restrict our analysis to years after 1983. For simplicity and consistency with previous research data are aggregated to the major grade level.

For each June-to-June, twelve-month interval from June, 1984 to June, 2000 data are compiled on the number of obligors within each grade that transition to each other grade or default. If an obligor's rating is withdrawn over the course of a your, it is excluded from the sample for that year.[5]

[4] Nickell, Perraudin and Varotto (2000) find evidence that the properties of rating transitions differ by obligor domicile and industry.

[5] On average 4.9 percent of ratings are withdrawn in a given year. Ratings may be withdrawn if an obligor is a party to a merger, or if its public debt is retired.

Table 1 reports the number of obligors assigned to each grade at the beginning of a sample year. Most obligors are assigned to the middle grades, while the number of Aaa and C-rated obligors is exceptionally small. As we shall see presently, differences in the number of obligors assigned to each grade has implications for the way observed transition data are interpreted. The number of obligors assigned to the lower grades grew dramatically over the sample period, reflecting the advent of the market for junk bonds.

Table 2 reports the average one-year transition matrix taken over the 16 sample years. We will frequently refer to this table in the analysis that follows. It provides a target that a well specified structural model should be able to approximate. A few essential features of the empirical transition matrix are worth noting. First, the overwhelming majority of obligors remain in the same grade over the course of a year. By-and-large, rating transitions are relatively rare, but a significant fraction of those transitions that occur may span more than one grade. Second, average default frequencies rise dramatically as one moves from higher to lower grades. This indicates that, as expected, Moody's ratings convey important information about the one-year default probabilities of obligors. Third, obligors are more likely to be downgraded than to be upgraded. This fact is almost certainly due largely to the fact that firms' cost of debt depends on their rating. Firms that expect to receive a low rating are unlikely to issue debt. Thus, few speculative grade firms issue debt. On the other hand, those firms that have previously issued debt may decline to speculative grade. Thus, on average, firms will tend to have high ratings initially, and see these ratings decline over time. Fourth, no high-grade obligors defaulted, and no speculative-grade obligors transitioned to high investment grades over the course of a single year. Obviously, since no obligor is truly riskless, the fact that we observe no defaults for some grades must be taken to imply that the true probabilities of these defaults are small but non-zero. More generally, null transition frequencies are almost certainly a result of the relatively small number of obligors assigned to some grades, rather than an indication that some transitions are impossible.

Figure 1 provides information on observed upgrade and downgrade frequencies over time. A single bar in the figure corresponds to a particular initial rating grade and year. Each bar is partitioned into three intervals that show the proportion of obligors in a grade-year that were upgraded (top), remained in the same grade (middle), or were downgraded (bottom). Notice that the frequency of upgrades and downgrades varies substantially from grade to grade and across years. The high volatility of transitions frequencies for grades C/Caa and Aaa can be discounted, to some degree, because of the small numbers of obligors actually assigned those grades. However, significant temporal differences in observed upgrade and downgrade frequencies can also be observed in the middle grades which each contain large numbers of obligors. For example, 1998 saw higher than average upgrade frequencies for nearly every grade, while 1999 so higher than average downgrade frequencies. Rating transitions appear to be correlated across obligors, suggesting the presence of systematic shocks that affect the credit quality of many obligors at once.

3 A Baseline Structural Transition Model

3.1 The Default Model

The starting point for any structural model of rating transitions is a description of the stochastic process that drives obligor default. We use a simple discrete time version of Merton's (1974) now ubiquitous stochastic default model. According to this model, an obligor defaults when the value of its assets falls below the value of its liabilities, or equivalent when its inverse leverage ratio (the ratio of liabilities to assets) falls below one. In applications such as this one, it is convenient to work with the natural logarithm of the inverse leverage ratio. Let Y_0 denote an obligor's log inverse leverage ration at the current as-of date, and let Y_1 denote the value of this variable at a fixed future horizon date. The relationship between Y_0 and Y_1 is given by

$$Y_1 = Y_0 + \sigma u \tag{1}$$

u is a random variable with mean zero and unit volatility, The parameter σ captures the volatility of firm leverage.

By dividing both sides of (1) by σ we can express the model in units that are invariant to differences in leverage volatilities across obligors. Let $y_t = Y_t/\sigma$ so that (1) becomes

$$y_1 = y_0 + u.$$

y_0 and y_1 respectively measure the obligor's as-of and horizon distance to default. An obligor will default at the horizon date if $y_1 < 0$. If the marginal distribution of u is identical across obligors, then y_0 is sufficient to determine the marginal probability that a particular obligor will default at the horizon date. This probability is given by

$$p_0 = F(-y_0)$$

where $F(u)$ is the marginal cumulative density function for u.

u need not be independent across firms. Systematic shocks, such as unforeseen changes in macroeconomic conditions, may affect the asset values of many obligors simultaneously. We account for such shocks by expressing u as the weighted sum of two variables, a systematic risk factor x and an idiosyncratic risk factor ϵ;

$$u = \omega x + (1 - \omega^2)^{1/2} \epsilon.$$

ω, called the *factor loading*, is a model parameter that can range from zero to one. It describes the share of the change in distance to default explained by the systematic risk factor. x is common to all obligors, whereas ϵ is obligor-specific and is independent across obligors. For the time being, both variables are assumed to have standard normal distributed. x and ϵ are independent over time.[6]

The additive properties of normal random variables imply that u has a standard normal marginal distribution and is independent across time. The stochastic features of u have important implications for firm defaults. The normality assumptions, while convenient, is quite restrictive. It will be relaxed in Section 4. The assumption that u is independent over time implies that a firm's distance to default at the as-of date is an efficient predictor of its default probability at the horizon date. We will allow for the possibility that rating agencies do not make use of these efficient predictors in section 5.

3.2 The Rating Assignment Model

Assume that at the as-of date, the rating agency assigns grades to obligors on the basis of their probabilities of default at or before the horizon date. Given the structure

[6]More general implementations of the Merton framework allow for the possibility of more than one systematic risk factor, permitting a richer structure of correlations in asset values and defaults across obligors. In practice, the data requirements needed to estimate multiple risk factor models are substantially more restrictive than those needed to estimate single factor models.

of our default model, this is equivalent to assuming that an obligor's grade depends on y_0, its distance to default at the as-of date. To map default distances into ratings, the continuum of possible values of y_0 is partitioned into I intervals, each of which corresponds to a single rating grade (see Figure 2). Let γ_i denote the upper threshold for the ith grade so that an obligor is assigned grade i if $\gamma_{i-1} < y_0 \leq \gamma_i$. There is no upper threshold for the highest grade, so an obligor is assigned grade I if $y_0 > \gamma_{I-1}$. Furthermore, since $y_0 < 0$ implies than an obligor is currently insolvent, zero is the lower threshold for grade 1.

At the assessment horizon, the grades of those obligors that have not defaulted are reassessed. Grades are assigned at the horizon date based on y_1 using the same algorithm used for assigning grades at the as-of date. Obligors transition between grades as their distances to default change over time. The transition probabilities for an obligor between the as-of date and the horizon date depend on the obligor's initial distance to default, and the agency's chosen rating thresholds. For simplicity, we assume that all obligors in grade i at the as-of date share the same initial distance to default, y_i^0.[7]

3.3 Marginal Transition Probabilities

To understand the implications of our baseline modeling assumptions for estimated transitions probabilities it is helpful to contrast the structural model with a more general, but in some ways more limited, reduced form transition model. Marginal transition probabilities can be described by an $I \times I$ matrix of reduced form transition thresholds. Let g_{ij} be a scalar with the property that an obligor in grade i at the as-of date will arrive in grade j or lower at the horizon date if $u \leq g_{ij}$. The transition probabilities for obligors in grade i are then

$$\begin{aligned} p_{i0} &= F(g_{i0}) \\ p_{ij} &= F(g_{ij}) - F(g_{i,j-1}) \text{ for } 1 \leq j \leq I-1 \\ p_{iI} &= 1 - F(g_{i,I-1}). \end{aligned}$$

[7]If grades form a reasonably fine partition of the range of possible default distances, the assumption of homogeneity within grades is innocuous. This assumption dramatically simplifies exposition and model estimation from aggregate data. If richer obligor-specific data were used, this assumption could be relaxed.

where $j = 0$ denotes default.

One can estimate the reduced form transition thresholds directly from default frequency data of the sort presented in Section 2. Indeed, this is exactly what is done when the RiskMetrics Group's CreditMetrics model is calibrated to historical data (Gupton, Finger and Bhatia 1997)[Chapter 6]. Because there are exactly as many reduced form thresholds as transition probabilities, this approach produces estimated transition probabilities that exactly match the transition frequencies reported in Table 2. However, reduced form models provide only limited insights into the stochastic process driving firm defaults and the behavior of rating agencies. Observe in particular that the reduced form transition thresholds have no natural economic interpretation because they imply nothing about how transitions between pairs of grades are related to one another. A second drawback to highly parameterized models such as the reduced form specification is that they are likely to "over-fit" historical data, limiting their usefulness for forecasting applications.

The reduced form thresholds can be linked to the structural parameters of our baseline model by equating the rating transitions conditions implied by each. Under the reduced form model an obligor currently in grade i will transition to grade j or lower if $u \leq g_{ij}$, while under the baseline model the same obligor will transition to grade j or lower if $y_i^0 + u \leq \gamma_j$. Thus, the two models imply the same transition probability when $g_{ij} = \gamma_j - d_i^0$. Repeating this logic for all possible transitions yields the following mapping from structural to reduced form parameters:

$$g_{ij} = \begin{cases} y_i^0 & \text{if} \quad j = 0 \\ \gamma_j - y_i^0 & \text{if} \quad 1 \leq j \leq I - i \end{cases} . \tag{2}$$

The structural model is substantially more parsimonious than the reduced form model. There are I^2 reduced form transition thresholds, but only $2I - 1$ structural parameters. Transition probabilities predicted by the structural model cannot be expected to exactly match observed transition frequencies. However, if the assumptions underlying the structural model are reasonably accurate, the estimated and observed transition probabilities should be close to one-another.

3.4 Parameter Estimates

Parameters for the structural model are estimated using a two-stage, hybrid method-of-moments/maximum likelihood procedure. In the first stage, the vector of as-of default distances $(y_1^0 \ldots y_I^0)$ and the vector of agency rating assignment thresholds $(\gamma_1 \ldots \gamma_{I-1})$ are estimated by fitting the predicted structural transition probabilities to the observed transition frequencies reported in Table 2. In the second stage, the factor loading ω is estimated by simulated maximum likelihood, holding the first-stage parameter estimates fixed.

Column (a) of Table 3 presents the structural parameter estimates for the baseline model. These parameters exhibit expected relationships to one-another. The initial distances to default for obligors assigned to higher grades are greater than those of obligors assigned to lower grades. The initial distance to default for each grade lies between the upper and lower rating thresholds for that grade.

To determine how well the baseline model captures observed patterns in rating transitions, the structural parameters are used to predict a matrix of marginal transition probabilities (Table 4). Comparing this matrix with the matrix of observed transition frequencies in Table 2, we see that the model does an excellent job of fitting the probabilities that obligors will remain in the same grade. The predicted probabilities that obligors in a grade will transition to adjacent grades are also reasonable close to the observed frequencies. However, the baseline model dramatically understates the likelihood that an obligor will move two or more grades from its initial rating over the course of a year. For all but the lowest speculative grade, predicted default probabilities are far lower than observed default probabilities.

These results indicate that the structural model is too restrictive to accurately describe the process of rating assignment and default. In the next two sections we generalize the baseline model with an eye towards better capturing observed features of the data.

4 Thick-Tailed Asset Return Distributions

Assuming that the year-to-year change in an obligor's distance to default is normally distributed, while convenient, may not be particularly realistic. Normally distributed random variables have the property that the probability of relatively large or small realizations falls quickly toward zero. Thus, if most changes in distance to default are relatively small, a normal model will predict that the probability of large changes is negligible. This is precisely what our baseline model implies. Grade thresholds must be sufficiently far apart to ensure that the probability that an obligor remains in its own grade is high. But given these wide grades, the odds that a change in distance to default drawn from a normal distribution will lead to a transition across many grades is very nearly zero. Table 2 indicates that, in fact, such events are not terribly rare.

Research on equity return data suggests that the shocks affecting firms' asset values (and by implication their default distances) are best modeled as coming from a "thick-tailed" distribution. Unlike normal random variables, the probability of drawing an exceptionally large or small realization of a thick-tailed variable can be nontrivial, even if most draws of that variable tend to be relatively small. Thus, by allowing for the possibility that changes in default distances have a thick-tailed distribution, we permit a reasonable likelihood of multi-grade transitions in an environment where most obligors do not change grades.

4.1 The alpha-stable distribution family

Though there are many thick-tailed distributions to choose from, practical considerations impose some limitations on our modeling assumptions. An important characteristic of the baseline model is that the systematic risk factor and the idiosyncratic risk factor can be summed together to yield a change in default distance whose marginal distribution does not depend on the factor loading parameter ω. This means that all model parameters except the factor loading can be identified and estimated from data on average transition frequencies. It is this fact that allows us to separate our estimation of model parameters into two stages. More importantly, it affords a degree of robustness to possible errors in assumptions about systematic risk and transition correlations. For example, estimated default distances and grade thresholds can be consistently estimated, even if our one systematic risk factor assumption is incorrect. For these reasons we prefer to choose a thick-tailed distribution that preserves the separability between the factor loading and other model parameters.

The alpha-stable distribution family meets all of our modeling requirements. This family includes as special cases the normal distribution as well as the well-known, thick-tailed, Cauchy distribution (Fama and Role 1968). All members of this family are additively stable, so that if systematic risk factors and idiosyncratic risk factors are drawn from identical alpha-stable distributions, changes in default distances will also be drawn from an alpha-stable distribution. For simplicity, we restrict our analysis to symmetric alpha-stable distributions.[8]

The tail-thickness, or kurtosis, of alpha-stable distributions depends on a shape parameter α that can range from zero to two. $\alpha = 2$ yields a normal distribution, while $\alpha = 1$ yields a Cauchy distributions. More generally, lower values of α correspond to greater kurtosis. Figure 3 illustrates the effect of changes in α on the shape of an alpha-stable probability density function. Notice that a change in α does not simply affect the "spread" of the distribution, as would a change in variance. Rather, it moves probability mass between the shoulders and the tails of the distribution. Distributions with α parameters that lie between one and two have unbounded variance but bounded mean. Those with parameters that lie between zero and one have both unbounded variance and unbounded mean. Because we believe it is unrealistic to assume that expected changes in distance to default cannot be measured, we restrict α to the range $1 < \alpha \leq 2$.

4.2 Parameter Estimates

We incorporate the idea that the distribution of changes in distance to default has thick tails by assuming that the systematic risk factor x and the idiosyncratic risk factor ϵ are both drawn from appropriately normalized symmetric alpha-stable distribution with the same kurtosis parameter α,[9] and that

$$u = \omega X + (1 - \omega^{\alpha})^{1/\alpha} \epsilon. \tag{3}$$

This generalization adds one additional parameter to our model, the single shape parameter α.

Column (b) of Table 3 presents estimated structural parameters for the thick-tailed model. As before, the default distances and grade thresholds bare the expected relationships to one-another. These parameters are much larger in magnitude than the baseline parameters, owing to the different shape of the distribution of default

[8]The alpha-stable family includes a range of skewed distributions. Preliminary results were not sensitive to skewness, so we elected to focus only on symmetric distributions.

[9]In addition to the shape parameter, symmetric alpha-stable distributions depend on a location and a scale parameter. These are normalized to 1 and $1/\sqrt{2}$ respectively.

distances. The estimated value for α, 1.14, is far from 2, suggesting that the normality assumption of the baseline model is indeed overly restrictive.

Comparing Table 2 with Table 5 reveals that the more general model continues to effectively predict the probabilities that obligors will remain in their current grades. Moreover, higher, more realistic probabilities of multi-grade transitions are generated. However, so much mass is placed in the extreme tails of the default distance distribution that the model predicts unrealistically high default probabilities, particularly for investment grade obligors.

Even when a relatively general stochastic default process is assumed, a rating model in which an obligor's distance to default determines both its grade and its default probability does not appear able to explain observed features of the marginal transition matrix. In the section that follows, we generalize the rating assignment model to relax the strong link it imposes between the distance to default used for assigning an obligor to a grade and that obligor's probability of default.

5 Rating Assignment Errors

Duffie and Lando (1999) develop a term structure model in which bond investors observe an obligor's distance to default with some error. An implication of this model is that at any point in time, an obligor's forward looking default probability may actually be much higher than would be suggested by its measured distance to default. A similar approach can be taken to model bond rating transitions. If an obligor's true distance to default is imperfectly measured by the rating agency, then the strong relationship between distance to default, grade, and default probability implied by the baseline model can be relaxed.

5.1 A simple measurement error model

To capture the idea that a rating agency miss-measures distance to default, we must distinguish between an obligor "true" distance to default y_t, and the agency's measured distance to default $\hat{y}_t$. The relationship between the two is given by

$$\hat{y}_t = y_t + sv_t \tag{4}$$

where v_t is the measurement error at time t, and s is a parameter that describes the importance of the measurement error in determining the observed distance to default.. To allow for the possibility that measurement errors are correlated over time, we assume that v_t follows the AR1 process,

$$v_t = \psi v_{t-1} + (1 - \psi^\alpha)^{1/\alpha} e_t \tag{5}$$

where e_ts are *iid* draws from an alpha-stable distributions. When $\psi = 1$ measurement errors are constant over time. When $\psi = 0$, new and independent measurement errors are made each period.

We assume that the agency uses only its measured default distances to assign ratings. It should be noted that this assumption implies a degree of naivete‘ on the part of the rating agency. It rules out the possibility that the rating agency updates its assessment of an obligor's distance to default as it observes that an obligor has not defaulted in the past. Incorporating such dynamics would imply that rating transitions do not follow a first order Markov process. While relaxing the Markov assumption would almost certainly prove fruitful, this generalization must await future research.

Let $\hat{\gamma}_j$ denote the upper assignment threshold associated with grade j, and let $\hat{y}_i^0$ denote the measured distance to default for obligors initially assigned to grade i. As before, an obligor is assigned to grade j when $\hat{\gamma}_{j-1} < \hat{y}_t \leq \hat{\gamma}_t$. The probability that an obligor in grade i will end the assessment period in grade j is

$$\begin{aligned}
p_{ij} &= \mathrm{P}\left[\hat{\gamma}_{j-1} < \hat{y}_1 \leq \hat{\gamma}_j \mid \hat{y}_i^0\right] \\
&= \mathrm{P}\left[\hat{\gamma}_{j-1} - \hat{y}_i^0 < u + sv_1 - sv_0 \leq \hat{\gamma}_j - \hat{y}_i^0 \mid \hat{y}_i^0\right] \\
&= F\left(\frac{\hat{\gamma}_j - \hat{y}_i^0}{\left(1 + ((1-\psi^\alpha) + (1-\psi)^\alpha)s\right)^{1/\alpha}}\right) - F\left(\frac{\hat{\gamma}_{j-1} - \hat{y}_i^0}{\left(1 + ((1-\psi^\alpha) + (1-\psi)^\alpha)s^\alpha\right)^{1/\alpha}}\right).
\end{aligned}$$

$F(\cdot)$ is the cumulative density function for the normalized symmetric alpha-stable distribution with shape parameter α. The last equality follows from the additive properties of alpha-stable variables.

Importantly, an obligor defaults only when its true distance to default falls below zero. Thus, the default probability for an obligor in grade i is

$$\begin{aligned} p_{i0} &= \mathrm{P}\left[y_1 \le 0 \mid \hat{y}_i^0\right] \\ &= \mathrm{P}\left[u - sv_0 \le -\hat{y}_i^0 \mid \hat{y}_i^0\right] \\ &= F\left(\frac{-\hat{y}_i^0}{(1+s^\alpha)^{1/\alpha}}\right). \end{aligned}$$

To more easily compare the measurement error model with our earlier models, it is helpful to rewrite the above equations in terms of parameters that more closely match those of the baseline model. We will use the following reparameterization:

$$\begin{aligned} \gamma_j &= \frac{\hat{\gamma}_j}{(1+((1-\psi^\alpha)+(1-\psi)^\alpha)s)^{1/\alpha}} \\ y_i^0 &= \frac{\hat{y}_i^0}{(1+((1-\psi^\alpha)+(1-\psi)^\alpha)s)^{1/\alpha}} \\ q &= \frac{(1+((1-\psi^\alpha)+(1-\psi)^\alpha)s)^{1/\alpha}}{(1+s^\alpha)^{1/\alpha}}. \end{aligned}$$

Substitution then yields a familiar mapping from structural parameters to reduced form transition thresholds:

$$g_{ij} = \begin{cases} qy_i^0 & \text{if} \quad j = 0 \\ \gamma_j - y_i^0 & \text{if} \quad 1 \le j \le I-1 \end{cases}. \tag{6}$$

The only difference between equation (6) and our baseline mapping is the new structural parameter q. q drives a wedge between transition probabilities and default probabilities. All else equal, default probabilities may be either higher or lower than those implied by the baseline model depending on the magnitude of q, which is a function of the overall importance of measurement errors (s) and the extent to which these errors are correlated over time (ψ).

5.2 Parameter Estimates

Columns (c) and (d) of Table 3 show two sets of estimates of the structural parameters in equation (6). In the first specification, the measurement error adjustment factor q is assumed to be constant across grades. In the second specification, we allow for the possibility that this adjustment factor might itself depend on an obligor's initial grade. This situation could arise, for example, if obligors with lower credit quality have poorer accounting practices, or alternatively, if rating agencies monitor lower quality obligors more rigorously. To more easily nest these two models, q is expressed as

$$q = \exp(\beta_0 + \beta_1 y_i^0),$$

and β_0 and β_1 are estimated. In the constant measurement error specification $\beta_1 = 0$.

A cursory examination of the matrix of marginal transition probabilities implied by the constant measurement error specification (Table 6) reveals that this model is not a great improvement over the thick-tailed model without measurement errors. Predicted default probabilities for high grade bonds are far greater than those actually observed.

In contrast, the matrix of transition probabilities implied by the grade-varying measurement error specification shown in Table 7 exhibits many of the salient features of Table 2. The probability that obligors remain in their current grades or transition to adjacent grades match closely. Implied default probabilities fall dramatically as one moves from lower to higher grades, and the default probabilities associated with the highest grades appear quite reasonable. The largest discrepancies between the predicted probability matrix and the empirical frequency matrix occur for transitions of greater than one grade. The probabilities assigned by the structural model are substantially higher than those observed. Nonetheless, the effectiveness of the grade-varying measurement error specification as compared with the baseline specification is impressive, particularly in light of the fact that it contains only three additional parameters (α, β_0, and β_1).

6 Conclusions

The primary goal of this analysis was to develop a structural model of bond rating transitions. Structural empirical models provide two distinct benefits over more highly-parameterized reduced form models. First, economic theory can generate powerful restrictions, substantially reducing the number of parameters that must be estimated. In applications such as this one where data are limited, parsimony is important. Models with fewer parameters are generally more easily identified by available data, and their parameters can be estimated more efficiently. Highly parameterized models have a tendency to "over-fit" observed data. reducing the effectiveness of out-of-sample forecasts. The second advantage of structural models is that they provides a rigorous, internally consistent framework from which we can draw economically meaningful inferences. Because the parameters that characterize structural models have economic interpretations, they lend themselves to scrutiny on theoretical as well as empirical grounds.

With regard to the first benefit of structural modeling, our empirical exercise was largely successful, The grade-varying measurement error model is able to described most of the important features of an observed transition matrix containing 49 cells using only 16 parameters. The second advantage of structural modeling was more elusive. The parameter estimates for the grade-varying measurement error model seem to imply unrealistically large differences in the importance of measurement errors across grades. Thus, although this model does a good job describing observed features of the data, it is less convincing as a description of the bond rating process. These findings suggest that the process that drives rating transitions cannot be closely tied to the process that drives default through a single solvency measure such as the distance to default.

One possible explanation for our findings is that rating agencies apply different standards for assigning grades at different points in the business cycle. Such "through-the-cycle" rating would imply that transition probabilities change in a predictable way over time. For example, if an agency applied stricter rating criteria during expansions than during recessions, then the matrix of observed transition probabilities would differ depending on whether an assessment period begins during an expansion or a recession. Bahar and Negpal (1999) and Bangia, Diebold and Shuermann (2000) find empirical evidence for this view. The idea of through-the-cycle rating can be incorporated into our structural framework by allowing the upper grade transition thresholds to depend on macroeconomic variables such as an indicator for recessions. This extension will be the subject of future research.

Year	Caa/C	B	Ba	Baa	A	Aa	Aaa	Total
1984	2	117	225	224	347	165	40	1120
1985	4	140	260	232	359	192	36	1223
1986	9	172	326	233	356	168	45	1309
1987	12	253	410	231	337	168	44	1455
1988	7	282	372	249	338	161	51	1460
1989	7	309	391	241	337	141	47	1473
1990	12	283	348	245	346	136	46	1416
1991	15	227	273	256	337	135	39	1282
1992	16	183	260	259	333	119	35	1205
1993	16	215	264	279	333	108	28	1245
1994	36	311	295	303	351	101	27	1424
1995	50	364	290	310	363	90	25	1492
1996	64	391	317	332	380	96	25	1605
1997	62	429	324	380	390	90	27	1702
1998	105	581	393	425	400	87	25	2016
1999	136	612	408	487	409	90	19	2161
Average	34.6	304.3	322.3	292.9	357.3	127.9	34.9	1474.1

Table 1: Number of obligors in each rating grade, by year.

Initial	Year-End Grade							
Grade	Default	Caa/C	B	Ba	Baa	A	Aa	Aaa
Caa/C	0.1638	0.7420	0.0646	0.0202	0.0094	0.0000	0.0000	0.0000
B	0.0779	0.0263	0.8309	0.0556	0.0067	0.0017	0.0007	0.0002
Ba	0.0172	0.0033	0.0932	0.8311	0.0490	0.0056	0.0003	0.0002
Baa	0.0008	0.0007	0.0087	0.0473	0.8930	0.0485	0.0009	0.0000
A	0.0000	0.0000	0.0021	0.0079	0.0549	0.9164	0.0184	0.0004
Aa	0.0000	0.0000	0.0004	0.0008	0.0056	0.0986	0.8878	0.0069
Aaa	0.0000	0.0000	0.0000	0.0000	0.0000	0.0037	0.0710	0.9252

Table 2: Average one-your transition frequencies.

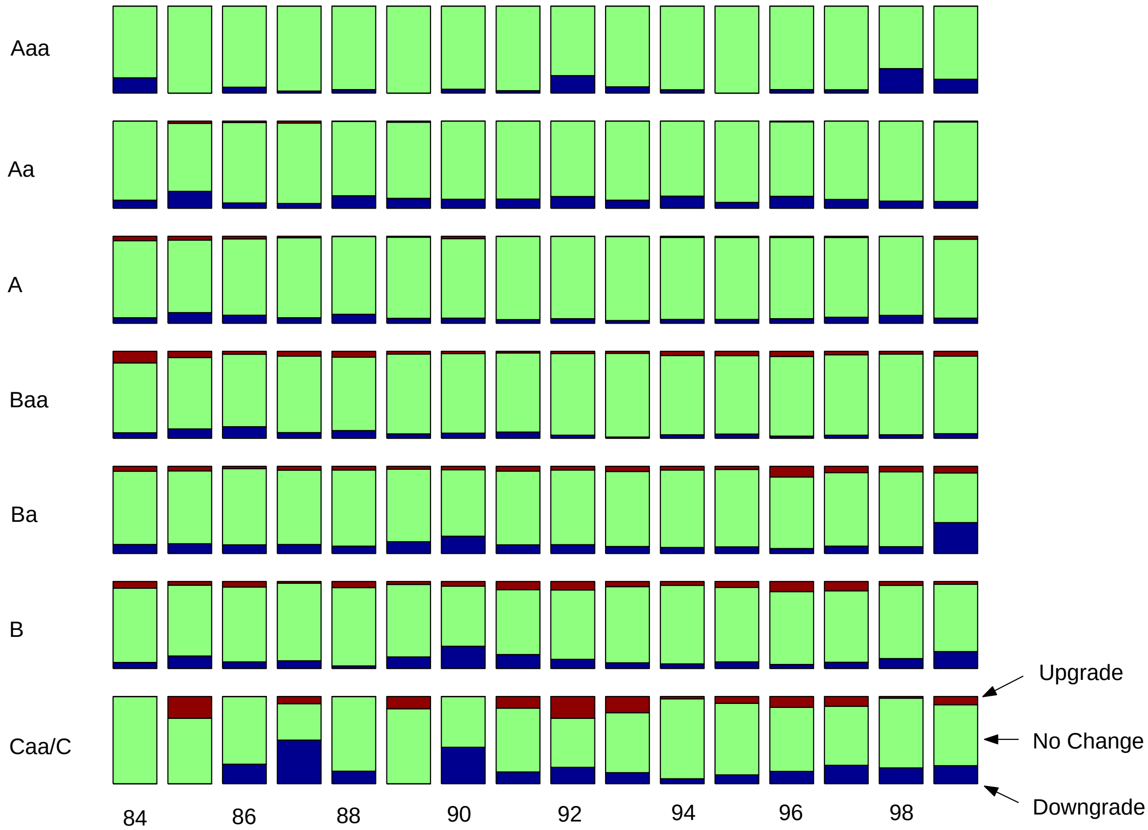

Figure 1: Shares of obligors upgrades and downgraded over one year, by initial grade and year.

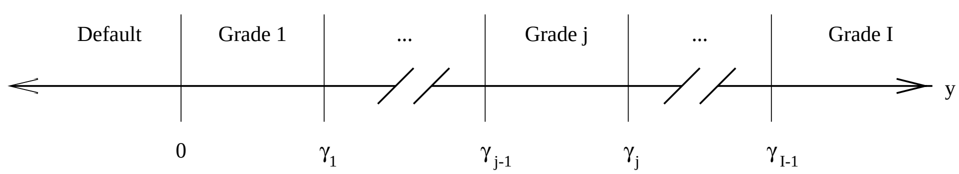

Figure 2: Partition of the range of possible default distances into rating grades.

	(a) Baseline Model	(b) Thick-Tailed Asset Returns	(c) Thick-Tailed Returns with Constant Rating Errors	(d) Thick-Tailed Returns with Variable Rating Errors
Upper Grade Thresholds (γ_j)				
Caa/C	2.2912	2.9792	4.8741	337.0140
B	5.0617	7.4792	8.0733	340.5062
Ba	7.8656	12.5855	11.4024	344.1889
Baa	11.0913	19.4199	15.4614	348.8219
A	17.9753	34.5541	20.9271	355.7309
Aa	22.8465	57.7742	28.6274	368.4784
Initial Distance to Default (y_i)				
Caa/C	0.9781	1.1361	3.4107	335.4426
B	3.5475	4.6893	6.2550	338.4850
Ba	6.2682	9.1440	9.4034	341.9108
Baa	9.4412	15.7076	13.3549	346.3965
A	12.6060	22.1702	17.2854	350.8508
Aa	19.2270	36.3217	22.3180	357.2070
Aaa	23.2878	60.1735	31.3163	370.3277
Tail Shape Parameter (α)	**2**	1.14	1.54	1.40
Measurement Error Parameters				
Intercept (β_0)	**0**	**0**	-1.2112	-71.5605
Slope (β_1)	**0**	**0**	**0**	0.1961

Restricted parameter values are shown in bold.

Table 3: Structural parameter estimates.

Initial Grade	Year-End Grade: Default	Caa/C	B	Ba	Baa	A	Aa	Aaa
Caa/C	0.1640	0.7414	0.0946	0.0000	0.0000	0.0000	0.0000	0.0000
B	0.0002	0.1043	0.8305	0.0650	0.0000	0.0000	0.0000	0.0000
Ba	0.0000	0.0000	0.1138	0.8311	0.0551	0.0000	0.0000	0.0000
Baa	0.0000	0.0000	0.0000	0.0575	0.8930	0.0495	0.0000	0.0000
A	0.0000	0.0000	0.0000	0.0000	0.0649	0.9351	0.0000	0.0000
Aa	0.0000	0.0000	0.0000	0.0000	0.0000	0.1053	0.8945	0.0001
Aaa	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0747	0.9253

Table 4: Predicted transition probabilities for baseline model.

Initial	Year-End Grade							
Grade	Default	Caa/C	B	Ba	Baa	A	Aa	Aaa
Caa/C	0.1664	0.7330	0.0763	0.0120	0.0051	0.0036	0.0016	0.0020
B	0.0345	0.0747	0.8279	0.0440	0.0097	0.0051	0.0020	0.0021
Ba	0.0159	0.0092	0.0873	0.8382	0.0354	0.0090	0.0026	0.0023
Baa	0.0085	0.0023	0.0070	0.0359	0.9000	0.0394	0.0042	0.0028
A	0.0058	0.0010	0.0024	0.0059	0.0488	0.9248	0.0079	0.0033
Aa	0.0033	0.0003	0.0007	0.0011	0.0025	0.0974	0.8887	0.0060
Aaa	0.0018	0.0001	0.0002	0.0003	0.0005	0.0020	0.0699	0.9252

Table 5: Predicted transition probabilities for thick-tailed asset returns model.

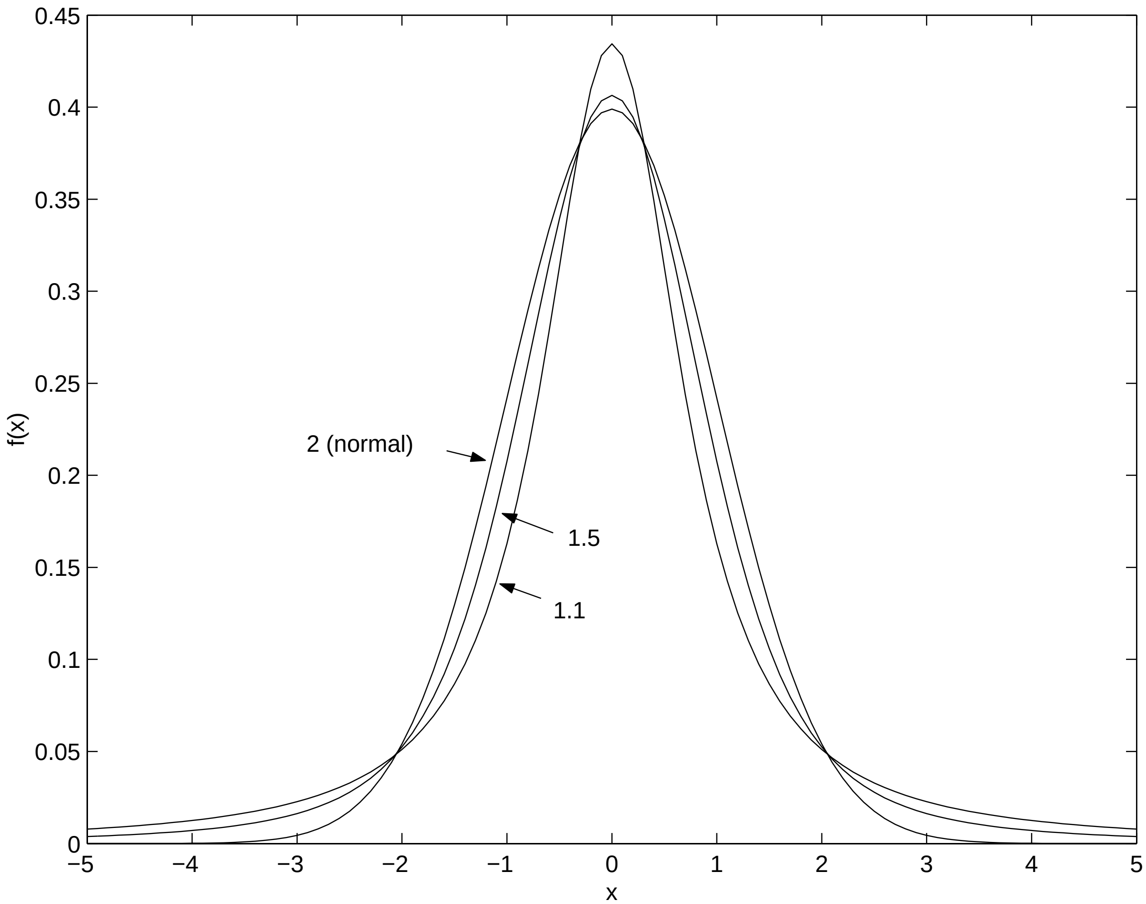

Figure 3: Three alpha-stable probability destribution functions with α parameters 2, 1.5, and 1.1.

Initial	Year-End Grade							
Grade	Default	Caa/C	B	Ba	Baa	A	Aa	Aaa
Caa/C	0.1669	0.7364	0.0853	0.0067	0.0022	0.0011	0.0006	0.0007
B	0.0621	0.0444	0.8283	0.0555	0.0059	0.0019	0.0009	0.0009
Ba	0.0283	0.0000	0.1014	0.8324	0.0469	0.0047	0.0015	0.0011
Baa	0.0150	0.0000	0.0050	0.0476	0.8942	0.0439	0.0036	0.0015
A	0.0096	0.0000	0.0014	0.0040	0.0570	0.9178	0.0151	0.0023
Aa	0.0063	0.0000	0.0005	0.0010	0.0031	0.0993	0.8893	0.0054
Aaa	0.0036	0.0000	0.0002	0.0002	0.0005	0.0015	0.0717	0.9253

Table 6: Predicted transition probabilities for constant measurement error model.

Initial	Year-End Grade							
Grade	Default	Caa/C	B	Ba	Baa	A	Aa	Aaa
Caa/C	0.1663	0.7377	0.0804	0.0086	0.0032	0.0017	0.0010	0.0011
B	0.0706	0.0352	0.8292	0.0519	0.0076	0.0028	0.0014	0.0012
Ba	0.0241	0.0000	0.0968	0.8331	0.0439	0.0062	0.0022	0.0014
Baa	0.0064	0.0000	0.0061	0.0440	0.8949	0.0423	0.0045	0.0019
A	0.0018	0.0018	0.0019	0.0049	0.0542	0.9190	0.0139	0.0026
Aa	0.0003	0.0018	0.0007	0.0012	0.0035	0.0979	0.8898	0.0048
Aaa	0.0000	0.0010	0.0002	0.0002	0.0005	0.0014	0.0714	0.9252

Table 7: Predicted transition probabilities for grade-varying measurement error model.

References

Bahar, R. and Negpal, K. M.: 1999, Dynamics of rating transition, *Technical report*, Standard and Poor's. Working paper.

Bangia, A., Diebold, F. X. and Shuermann, T.: 2000, Ratings magration and the business cycle, with applications to credit portfolio stress testing. Working paper.

Duffie, D. and Lando, D.: 1999, Term structures of credit spreads with incomplete accounting information. Working paper.

Fama, E. F. and Role, R.: 1968, Some properties of symmetric stable distributions, *American Statistical Assiciation Journal* .

Gupton, G., Finger, C. and Bhatia, M.: 1997, Creditmetrics – technical document, *Technical report*, The RiskMatrics Group.

Merton, R. C.: 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* **29**, 449–470.

Mittnik, S. and Rachev, S.: 1999, *Financial Modeling and Option Pricing with Alternative Stable Models*, Wiley.

Nickell, P., Perraudin, W. and Varotto, S.: 2000, Stability of rating transitions, *Journald of Banking and Finance* **24**, 203–227.