AN EMPIRICAL INVESTIGATION OF CREDIT TRANSITION MATRICES

Adrian M. Cowan
The George Washington University
School of Business and Public Management
Department of Finance
2023 G Street, N.W.
Washington, D.C. 20052
Phone: 703/757-9060
Email: cowana@gwu.edu

May 2001

AN EMPIRICAL INVESTIGATION OF CREDIT TRANSITION MATRICES

Abstract

Recent literature suggests using a discrete, state space Markov chain in credit ratings as a proxy for the default process in the valuation of credit risky securities. Several unresolved empirical issues arise in the application of Markov models. The purpose of this paper is to examine the time, age and industry characteristics of credit transition matrices over the period from January 1976 to January 1998. Using the Fixed Income Database from the University of Houston, I provide evidence contrary to results presented in previous studies. I generally find that conditional probabilities are age and industry invariant but not time invariant.

AN EMPIRICAL INVESTIGATION OF CREDIT TRANSITION MATRICES

A relatively recent development in the valuation of credit risky securities involves using an exogenous default process based on a first-order Markov chain in which credit rating transitions serve as a proxy for default. Examples of such models include the studies of Jarrow et al. (1997) and Kijima et al. (1998). However, several unresolved empirical issues surround the use of these models in the valuation process. This paper investigates whether a first-order Markov chain model in credit ratings accurately depicts the default process and is an accurate component in the pricing of risky debt.

This present paper is important for two reasons. First, it provides empirical research in an area that is not easily explored given the paucity of public data. Most of this work is developed by the industry with largely proprietary databases and remains unpublished. Second, this paper provides support for the integration of the bankruptcy and contingent claims pricing literature. The modeling of default as a finite state, Markov chain was first advanced by the bankruptcy literature. In fact, Altman (1989) suggests that his early work presents a Markov chain approach to the modeling of bankruptcy. At the same time, option pricing is dependent upon an understanding of bankruptcy. In the seminal work of Black and Scholes (1973), the firm's equity is viewed as an option that is valuable at the time a firm's debt matures only if the debt can be paid in full. It is the combination of bankruptcy modeling with valuation models, such as presented by Jarrow and Turnbull (1995), which establish the critical link between the bankruptcy literature and the contingent claims pricing literature.

The purpose of this paper is to investigate the underlying assumptions regarding time, business sector and aging in the development of the historical credit transition matrix. Each of these assumptions holds important implications for valuation techniques using Markov processes to model the dynamics of the default process. These assumptions are both individually and collectively important in linking theoretical and empirical models of credit risk. I analyze the empirical credit migration process as evidenced by the major credit rating transitions from January 1976 through January 1998 as assigned by Moody's Investor Services.

For purposes of this paper, it is important to distinguish between the transition matrix using the actual probability measure and that using pseudo-probabilities. As noted in Jarrow et al. (1997), the risk neutral credit rating process is not necessarily Markovian and may not be time-homogeneous. Tests conducted for my paper reflect the empirical transition matrix as measured using historical changes in credit ratings and not the empirical generator matrix that incorporates the risk premia necessary to convert the matrix into the pseudo-probabilities under a martingale process.

Jarrow et al. (1997) assume that transitional probabilities of credit rating migration are time homogeneous. However, there is no empirical evidence to support this assumption. Each transitional probability in such models represents the conditional probability of moving to a given credit rating given the current credit rating. In a time-homogeneous Markov chain estimates of transitional probabilities are independent of time; i.e., there is no change in the n-step transition matrix regardless of the time period in which the probabilities are measured. Other authors, such as Altman and Kao (1992), suggest that rating migration may be time sensitive. They analyze credit ratings drift using a sample of more than 7,000 new bond issues from 1970 through 1988. Kijima et al. (1998) incorporate non-homogeneity in their theoretical model, but fail to test this assumption empirically. The present paper extends this literature as it provides formal tests for time homogeneity. The test results provide evidence that there is little support for such an assumption.

Altman and Kao (1992) briefly address the importance of the economic sector of the firm in the credit migration process. Although they do not specifically measure the significance of this factor in the transition matrix, they analyze the propensity for bonds to experience multiple rating changes and the direction of those changes by industry sector. By performing nonparametric tests on their published results, it is possible to show that the differences by industry are statistically significant. This result provides some indication that differences in credit migration may exist by economic sector. Just as business cycles would be expected to affect time homogeneity, industry life cycles might be expected to influence business sector effects on credit transition matrices. I test the null hypothesis that the credit transition matrix is constant across business sectors. I use the same three business sectors of industrial, financial and utilities as Duffee (1998) uses to test bond yield spread behavior. I find that the results are highly sensitive to the weighting methodology I employ in the estimation of the transition matrix.

Another factor assumed to influence conditional probabilities of credit migration is the age of a bond. Altman (1989) attempts to identify the estimated probability of default and loss from default over a specific time horizon given a bond issue's initial bond rating. Asquith et al. (1989) provide a unique aging analysis of high yield bonds revealing that two important sources of bias in default studies are the exclusion of distressed exchanges and the failure to consider bond age. The authors provide tables of default rates from time of issue to support a finding that default rates are lower immediately after issue and rise over time. However, this study is limited in scope as it focuses on original issue, non-investment grade bonds. In addition, the authors propose that results of other studies may differ based on alternate economic conditions. Altman and Kao (1992) analyze bond credit rating transitions from 1970 to 1988 and provide preliminary evidence that the transition matrix is a function of a bond's age. A weakness of their approach is that it does not compare the n-step transition between older groups of bonds with newly issued bonds. In addition, a weakness of all of the previous studies is that they do not provide formal tests of the influence of bond age on default probabilities. I conduct nonparametric tests that show little evidence of aging effects.

In summary, the purpose of this paper is to investigate empirically the validity of the assumptions of time, business sector and age invariance in the development of the credit transition matrix. I find no evidence to support an assumption of time homogeneity or age variation in transitional probabilities. I find only limited support for a business sector impact, with the results quite sensitive to the weighting methodology. Whereas the lack of age and industry effects simplifies modeling efforts, the lack of time homogeneity complicates the empirical estimation of such models.

The remainder of this paper is organized into 4 sections. Section 1 describes the theoretical foundation. Section 2 details the methodology, including the data sampling and estimation procedures. Section 3 documents the empirical findings and limitations of the paper. Finally, Section 4 provides concluding remarks and suggestions for potential future research.

1. Integration of Bankruptcy and Contingent Claims Models

The significance of empirical transition matrices arises through their use in valuation models of credit-risky securities. Both bankruptcy theory and contingent claims pricing play key roles in the development of such models. Bankruptcy provides the foundation for the estimation of the probability of default and the recovery rate that provide salient information in pricing credit risk. Consequently, current advances in valuation incorporate such default modeling in the valuation of risky debt, credit risky derivatives and the related estimation of the term structure of credit spreads.

The bankruptcy literature, which is developed to a large extent by Altman, lays the foundation for a class of security valuation models that incorporate credit risk. Altman and Saunders (1997) present an insightful historical development of the bankruptcy literature, including mortality rate models. Such models, as exemplified by the mortality rate models of Altman (1989) and Asquith et al. (1989), are based on credit rating migration. These models represent the theoretical foundation for the development of empirical transition matrices.

Mortality rate models are grounded in the theory of capital markets, which tells us that investors require a premium to compensate for risk. Therefore, the probability of default of various credit classes of bonds should be reflected in the net spread over the risk free rate. These models attempt to obtain probabilities of default from past data on bond defaults by credit grade and years to maturity. And it should be noted that these models reflect the approach utilized by the credit rating agencies, such as Moody's or Standard & Poor's.

The integration of credit risk analysis with the pricing of securities exposed to credit risk is a natural extension of the above research. It can easily be shown that the cost of risky debt equals the risk free rate of return plus a risk premium. Many theoretical models of the term structure of risk premiums have been developed within the framework of a contingent claims pricing of credit risk. Jarrow et al. (1997) cite three major contingent claims models in the literature to date, each of which is arbitrage free. The models can be differentiated from one another based on the modeling of the bankruptcy process, which again highlights the importance of the previous bankruptcy literature. The newest class of contingent claims models is the one introduced by Jarrow and Turnbull (1995) and Jarrow et al. (1997). Specifically, their paper presents a contingent claims valuation model that explicitly incorporates credit rating information into the valuation methodology. Within this model, risky debt pays off a fraction of each promised dollar in the event of bankruptcy, but bankruptcy itself is now given as an exogenous process. The bankruptcy process is modeled as a finite state Markov process in the firm's credit ratings. Model tractability is enhanced by the fact that debt seniority is managed through different recovery rate assumptions and the model can be employed with any risk-free term-structure model. Jarrow et al. (1997) suggest the use of a historical transition matrix for model implementation which gives rise to the need to empirically evaluate the characteristics of such a matrix.

2. Methodology

The credit transitional matrices are tested across time, industry, and bond age. These tests are accomplished following a three-step process. I draw the sample, estimate the transition matrices and develop the test statistics. Each of these steps is detailed in the following sections.

2.1 Data Sample

This paper uses the Fixed Income Database from the University of Houston to draw all sample data. The database consists of month-end data for bonds that comprise the Lehman Brothers Bond Indexes. Separate data samples are drawn to test each specific hypothesis. The methodology for drawing samples is similar across hypotheses, and the detail of each is provided below. Changes in credit ratings are measured as a major credit rating change as reported by Moody's Investor Services. The Fixed Income Database is extended in 1992 to include non-investment grade firms. Transition matrices are extended to include more rating categories as data allows.

The testing of time homogeneity requires the time series of annual credit ratings per firm and per bond. Tests include the annual time series of observations from the Fixed Income Database from January 1976 to January 1998 for the following Standard Industrial Classifications (SIC): industrial, utilities, and finance. Each row of transitional probabilities in the Markov chain is estimated separately based on the initial Moody's credit rating of the sample. I use two separate weighting methodologies when drawing samples. The first sample is weighted by firm. Each firm is assigned the rating given by Moody's Investor Services to its senior, regular, outstanding debt. The alternate sample is weighted by outstanding debt amount. The firm-weighted approach is consistent with industry practice. The latter methodology is consistent with Altman et al.'s (1992) weighting by the bond's weighting in the portfolio. Each matrix is also estimated over multiple investment horizons. All sample sets are limited to the regular, senior bonds of each firm. The sample excludes bonds with call options, put options, convertible options and sinking funds. Bonds that mature within the measurement period are excluded to avoid distorting the results by elimination through maturity. In addition, I include all bonds issued up to one month prior to the measurement period. No attempt is made to identify distressed exchanges, although this presents a potential for future research.

The sample data sets described above are further divided into industry sectors to test for industry invariance of the credit transition matrices. Industry is identified by the two-digit SIC code, as follows:

03: Industrial,

04/05/06: Telephone/Electric/Other Utility,

07: Finance.

Finally, sample data sets as described for use in the tests for time homogeneity are expanded to include issue date to capture the influence of aging on credit migration. Due to the lack of information on firm age, I use a single weighting methodology of debt amount outstanding in testing for the significance of age in transition matrices. Age is measured as months since issuance and is used to develop portfolios of similar age bonds of the same credit quality. The same exclusion criteria are applied as set forth for the previous hypotheses.

2.2 Estimation of Transition Matrices

The results of this paper are sensitive to the methodology employed in the estimation of the transition matrix. The crucial decision factor is the selection of the "appropriate" weighting in the estimation of the transitional probabilities. There are many weighting schemes from which to select, such as weighting by dollar amount outstanding, by firm, or by number of bonds. And each of these weights can be further categorized by age, industry, etc. The weighting methodologies I use in this paper are consistent with the industry and with the literature. I note where the weighting impacts the results.

The procedure used to estimate credit transition matrices is similar across hypotheses. The time-homogeneous Markov chain is defined on a finite state space with S = (0, 1, . . ., K) and is measured over various time steps. The states are possible credit classes with 1 being the highest and K-1 being the lowest class prior to default. State K is assumed to be an absorbing state of default, although this assumption may be relaxed. Thus, the finite-state, homogeneous, Markov chain is specified by a K x K transition matrix, as follows:

$$P = \begin{pmatrix} p_{00} & p_{01} & \cdots & p_{0K} \\ p_{10} & p_{11} & \cdots & p_{1K} \\ \vdots & \vdots & \vdots & \vdots \\ p_{K-1,0} & p_{K-1,1} & \cdots & p_{K-1,K} \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (1)$$

Given that the credit process must transition into some state and probabilities are non-negative, we have the following:

$$p_{ij} \geq 0, \qquad i, j \geq 0 \quad (2)$$

and

$$\sum_{j=0}^{K} \mathrm{p}_{\mathrm{ij}} = 1, \qquad i = 0,1,\ldots,K \quad (3)$$

The transition probability matrix is not known, but rather is estimated. Based on the observed outcomes and following Guttorp (1995), I estimate the maximum likelihood estimator $\hat{p}$ that maximizes the likelihood of the parameter p, L(p). The log likelihood function, l(p), that is maximized given a multinomial distribution with n observations and K credit categories is as follows:

$$l(p) = \log L(p) = \log \frac{n_i!}{n_{i0}! \cdots n_{iK}!} + \sum_{j=0}^{K} n_i \log p_{ij} \qquad i,j = (0,1,\ldots,K) \qquad (4)$$

$$\text{where : } n_i = \sum_{j=0}^{K} n_{ij}, \quad \text{and}$$

n_{ij} is the observed count from credit rating i transitioning into credit rating j

By maximizing l(p) over all p that sum to 1, the following maximum likelihood estimator (mle) is obtained:

$$\hat{p}_{ij} = \frac{n_{ij}}{n_i}, \qquad i,j = (0,1,\ldots,K) \qquad (5)$$

Thus, I obtain the following, finite-state, K x K transition matrix:

$$\hat{P} = \begin{pmatrix} \hat{p}_{00} & \hat{p}_{01} & \cdots & \hat{p}_{0K} \\ \hat{p}_{10} & \hat{p}_{11} & \cdots & \hat{p}_{1K} \\ \vdots & \vdots & \vdots & \vdots \\ \hat{p}_{K-1,0} & \hat{p}_{K-1,1} & \cdots & \hat{p}_{K-1,K} \\ 0 & 0 & 0 & 1 \end{pmatrix} \qquad (6)$$

2.3 Log Likelihood Ratio Statistic

I evaluate the three hypotheses being tested using a standard goodness of fit, Chi-square test. Numerous sets of tests statistics are calculated for each null hypothesis given multiple years of matrices and industries.

The first set of tests is designed to analyze whether the historical probabilities are time invariant. The null hypothesis is as follows:

$$H_0\!:p_{ij}(t=0)=p_{ij}(t=1)=\cdots=p_{ij}(t=T)\,, \qquad i,j \;=\; (0,1,\ldots,K)$$

The second set of test analyzes the significance of using an industry-specific transition matrix using the following null hypothesis across constant timeframes:

$$H_0\!:p_{ij}(SIC_t=I)=p_{ij}(SIC_t=F)=p_{ij}(SIC_t=U)\,, \qquad i,j \;=(0,1,\ldots,K)$$

where : I = Industrial, F = Financial, and U = Utilities.

Finally, if there is no affect from aging, then there should be no statistical difference between the transition matrix estimated using pooled, older bond data as compared with the transition matrix estimated using newly issued bond data. The pooled bonds naturally exclude the "new issues" to maintain the independence of the samples. Thus, the third set of tests are designed to capture any age-specific influence based on the following null hypothesis:

$$H_0\!:p_{ij}(new_t)=p_{ij}(old_t)\,, \qquad i,j \;=\; (0,1,\ldots K)$$

Some variation of the following likelihood ratio test statistic is developed to test each null hypothesis. The variation of the test simply depends on the number of matrices being analyzed. Under the independence assumption, we have a multinomial distribution with the following observations from each category with probability θ_j.

$$n_{\bullet j}=\sum_i n_{ij}\;, \qquad \text{i, j} = (0, 1, \ldots, \text{K}) \qquad (7)$$

The likelihood is as follows:

$$l(\theta) = \sum_{j=0}^{K-1} n_{\bullet j}\theta_j + n_{\bullet K}\left(1 - \sum_{j=0}^{K-1}\theta_j\right) \tag{8}$$

This likelihood is maximized by the following estimate:

$$\hat{\theta}_j = {n_{\bullet j}}/{n} \tag{9}$$

Finally, the log likelihood ratio statistic for testing the null hypothesis is provided below. This statistic asymptotically has a χ^2 distribution with K(K+1)–K = K^2 degrees of freedom.

$$2\left(l(\hat{P}) - l(\hat{\theta})\right) = 2\sum_{i=0}^{K}\ \sum_{j=0}^{K} n_{ij} \log \frac{n_{ij}/n_i}{n_{\bullet j}/n} \tag{10}$$

3. Empirical Results

My results indicate that the use of historical transitional probabilities in the development of the generator matrix, as suggested by Jarrow et al. (1997), is complicated by several factors. The results of the tests and the analysis of these results are presented in the following sections.

3.1 Time Homogeneity

Time homogeneity is not a requirement for the use of an empirical transition matrix. Formally, if $P\{X_{t+s} = j | X_s = i\}$ is independent of *s*, then the Markov chain is said to have stationary or homogeneous transition probabilities. Although it is always possible to make the transition probabilities a function of time, the homogeneity assumption greatly simplifies the estimation procedure. Unfortunately, I do not find any compelling evidence to support the assumption of time homogeneity as shown in Table 1. The overall tests of all years clearly lead to a rejection of the null hypothesis of time invariance, regardless of weighting methodology or number of credit categories.

[INSERT TABLE 1]

The 23-year sample period of 1976 through 1998 provides for the estimation of 22 historical, one-step transition matrices. For the purposes of testing for time homogeneity, matrices are tested both in consecutive and non-consecutive time periods. When analyzing stability across consecutive years, all of the tests on firm-weighted matrices lead to a failure to reject the null hypothesis of time invariance at the 5% level of significance. This result, which holds regardless of the number of rating categories, suggests that the assumption of time homogeneity may be appropriate for very short-term investment horizons. However, there is a danger in the selection of only successive years in the development of the Chi-square tests as the results must hold independent of the time period selected. By comparing matrices across non-consecutive time intervals, I find that almost every time period leads to a rejection of the null hypothesis at the 5% significance level as shown in Table 2.

[INSERT TABLE 2]

The test results are sensitive to the weighting methodology. The dollar-weighted matrices exhibit less stability than the firm-weighted matrices. Although the firm-weighted probabilities are homogeneous across all consecutive, annual time periods, the dollar-weighted matrices lead to a rejection of the null hypothesis at the 5% level of significance for 17 out of 21 consecutive years. The results are even more definitive when comparing non-consecutive time periods. As shown in Table 2, tests of all non-consecutive time period matrices result in a rejection of the null hypothesis of time invariance at the 5% significance level.

It could be argued that the results depicted for the entire sample period are distorted by the change in the sample from the 1992 period forward when noninvestment grade bonds are added to the sample. This is an example of incomplete information censoring. However, whether or not the matrix is measured using additional credit categories for the high-yield debt categories does not materially affect the results.

Jarrow et al. (1997) suggest that time homogeneity appears to be a more reasonable assumption for investment grade bonds than for speculative grade bonds based on evidence contained in Moody's Special Report (1992). If so, then the contribution to the Chi-square test statistics from these categories should be minimal. However, this is not the case in the current data sample. The detailed results reveal that overall, the categories of credit ratings superior to Baa contribute 52% to the test statistic over the entire sample period. Alternatively, in comparing the transitional probabilities estimated for the final seven years of the sample period, the higher credit rating categories contribute approximately 32% to the Chi-square test statistic. A Chi-square test statistic calculated based on a reduced matrix limited to higher credit ratings does not exhibit stationary transition probabilities regardless of estimation period.

Finally, it is possible that a change in the industry composition of the dataset affects the time homogeneity results. There is no attempt in the development of the samples in each year to maintain a constant mix across industries. Given the influence of business sector measured in the same time period for dollar-weighted matrices, this may influence the time-homogeneity results on a dollar-weighted basis across time periods. The approximate industry composition in each year of the sample period is provided in Figure 1. Unfortunately, the maintenance of a constant mix severely limits the sample primarily due to the small number of firms in the utility and finance industries. Nevertheless, as this difference might at least in part explain the results, additional insight might be gained in the future by constraining the sample to isolate time effects relative to industry effects.

[INSERT FIGURE 1]

3.2 Industry Effects

If the default process does not differ by industry, then the transitional probabilities estimated from alternate industries should not differ. When measured on a firm basis, the results suggest that the default process for firms is relatively stable across industries. This is consistent with the findings of Duffee (1998) that corporate yield spreads relative to Treasuries do not differ by business sector. However, these results do not hold for dollar-weighted matrices.

These results illustrate most clearly the importance of the weighting methodology employed in the development of the transition matrix. As presented in Table 3 below, I fail to reject the null hypothesis of industry invariance for the overall test across the industrial, utility and finance industries based on firm-weighted matrices. The result holds for all annual transition periods subsequent to 1991/92. However, I obtain the opposite result when considering the dollar-weighted matrices measured in the same time frames. I reject the null hypothesis of industry invariance at the 5% significance level in each period. This result suggests that there may be an industry life cycle that influences the probability of default and leads to little consistency across industries.

[INSERT TABLE 3]

Several factors contribute to the mixed results. Firms typically have multiple senior-level bonds outstanding. When one of these bonds experiences a credit rating change, the remainder of the bonds will also. Therefore, estimating the transitional probability matrix without weighting by issuer will bias the transition probabilities toward large issuers.

Industry classifications may indeed be important in the determination of credit migration. However, time series data may be insufficient to obtain statistically significant results. In fact, this lack of data motivates the J.P. Morgan CreditMetrics' indirect approach. CreditMetrics' estimates credit migration based on the underlying firm asset value and employs an individual firm transitional approach.

In recognition of this limited data problem, this paper is limited to the use of the 2-digit SIC Code for the industry categorization of firms, despite the fact that the 3-digit

SIC Code would provide more detailed industry breakdowns. In addition, the data on the Lehman Bros. Fixed Income Database is not available at the 3-digit SIC Code level and thus mandates the grosser partition of industry classifications. A potential for further research is the partitioning of this dataset by matching the file against Compustat data.

Although not directly related, Duffee's (1998) study of mean yield spreads on noncallable bonds relative to Treasury yields supports the results obtained with firm-weighted matrices. He regresses the one-month change in the mean yield spread on noncallable bonds within industry, rating and maturity cohorts. Using joint estimation across the same three business sectors as used in the present study, he finds no evidence that different business sectors react differently to Treasury yields. To the extent that such yield spreads represent the market price of risk in default, this is consistent with a default process that is consistent across business sectors.

3.3 Aging Effects

In a continuous-time Markov chain, the amount of time spent in each state *i*, before proceeding to an alternate state, is exponentially distributed. It is important to recognize that the amount of time that a process spends in state *i* and the next state visited must be independent random variables. Otherwise there would be a violation of the Markovian assumption since knowledge of how long the process has been in state *i* would be relevant to the prediction of the subsequent state.

Altman and Kao (1992) suggest that there is an aging affect present in transition matrices. The authors did not conduct formal tests, but rather relied on the observed changes in the transition matrices over time as bonds mature. Based on this preliminary evidence, I segregate the transition matrices in select time periods between those bonds 1, 3 and 5 years since issue from other bonds in the sample. Again, if there are no differences in the two matrices developed from each segregated sample, this would indicate no age affects in the matrices. The results of these tests are shown in Table 4.

[INSERT TABLE 4]

Contrary to the findings of Altman and Kao (1992), in general I find that there is no age effect in the transition matrices. Analyzing 1-step, 3-step and 5-step transition probabilities as provided in Table 5, there is little evidence to support an assumption of age variance. I fail to reject the null hypothesis of age invariance at the 5% level of significance in all 5-step and 3-step transition matrices. And although the results of the 1-step transition matrices are not conclusive, the results in the 1-step matrices reflect a tendency towards age invariance. The results of the Chi-Square tests support a rejection of the null in four out of six years of 1-step matrices. I conduct these tests based on 7 rating categories due to insufficient data on new issue non-investment grade bonds. Clearly, it is likely that the expansion of the matrices to the 9 categories would lead to a failure to reject the null over the complete sample period. Such a result supports the use of the Markov chain approach to the default process.

Despite these results, there is still a potential for age effects in credit migration. There is no test of the aging hypothesis on a firm-weighted basis given lack of information regarding the age of the firm. Anecdotal evidence suggests that the greatest contribution is provided by credit rating agencies at the beginning of the firm's "public" life cycle when the information asymmetries are greatest. This asymmetry is related to the age of the firm and not the age of the bond. A study of the impact of firm age awaits future research.

3.4 Limitations

Some limitations exist that should be considered in the analysis of the results of this paper. Although I do not believe that any materially affect the results, they represent potential complications in the research.

The use of the Lehman Bros. Fixed Income Database introduces the obvious limitation of selection bias. The database consists of a set of bonds that are held within Lehman Bros. index portfolios. Thus, the sample is not a random sample of the total population of bonds. This limitation cannot be controlled for within my paper and suggests potential limitations in the ability to generalize my paper to the investment universe. However, the relatively large size of the database that includes over 70,000 instruments mitigates this problem.

Credit ratings serve as a proxy for the unobserved creditworthiness. This proxy represents subjective ratings as determined by individuals at the various rating agencies. The use of this measure as a proxy is further complicated by the fact that rating agencies don't always agree on ratings and therefore split ratings exist. For example, Altman (1982) finds a long median lag of 6 to 7 months between a rating change by one of the major rating agencies and the subsequent change by another. Therefore, the estimation of the transition matrix necessarily is impacted by the selection of the agency used for ratings. As noted previously, the ratings of Moody's Investor Services are used in this paper. The use of average credit ratings in future research presents a potential for mitigating the sensitivity of the results to the selected rating agency.

There are bonds that are not rated based on such factors as the outstanding balance of the debt (outstanding debt < $25 million) or for insufficient information. Therefore, although it's possible to go from a valid credit rating category to not rated (NR), the subsequent information on the disposition of the issue is not available. Given the lack of relevant information, all bonds that transition into the NR category during the sample period are eliminated from the paper. It should be noted that the simple elimination of this category may bias the transitional probabilities downward. If firms whose bonds are categorized as NR tend to default more frequently, this would tend to underestimate default probabilities in general. In addition, given that the lower credit rating categories tend to have higher NR probabilities, this is likely the case.

Firms exchange debt in lieu of bankruptcy on certain occasions. Such distressed exchanges tend to underestimate true defaults as noted within Asquith et al. (1989). However, no attempt is made within this paper to control for this issue given that the fixed income database does not provide data on distressed exchanges.

4. Conclusions

The focus of this paper is the analysis of the characteristics of the default process as proxied by the credit migration process. Given a discrete state space, continuous time Markov model of default, I study the degree to which time, business sector and bond age impact transition matrices. Drawing upon well-developed, statistical literature on the properties of such matrices, I use standard procedures in estimation and testing with a dataset covering an extensive 23-year period.

The valuation of credit risky securities using an exogenous Markov process in the firm's ratings requires a generalization to time dependent conditional probabilities to be consistent with empirical results. Although such a generalization complicates the estimation, it doesn't make such models obsolete. The results of this paper also suggest that the overall methodology conforms with the empirical evidence. I do not find strong evidence in support of age affects as suggested by Altman and Kao (1992). In addition, whether a generalization is required for the industry affects may require more refinement in the test for industry affects; i.e., the ability to isolate industry at more refined levels of categorization. Therefore, there is much to support the further development of Markov chain models of credit risky securities as well as further studies of the properties of transition matrices.

In light of the depth of this research topic, there is a broad range of potential avenues for future research, including possibilities for theoretical development that are particularly interesting. The first alternative is the delineation of the transition matrix that reflects the "true" default process as compared with the credit rating process. In the spirit of the Tversky et al. (1986), the default process model that agents use to value credit risky securities may differ significantly from the credit rating process. Economic agents may simplify the model to approximate the true but more complex model. Such an approach would likely lead to a richer integration of the bankruptcy and valuation literature. Another alternative involves the investigation of the impact of the asymmetry of information on the credit migration process. Consistent with market microstructure literature, the asymmetry of information between the rating agencies and the bond issuer is anticipated to be greatest when the firm is young.

We are only just beginning to explore the integration of the bankruptcy literature and the asset pricing literature. Because the true model of default cannot be observed, there are many aspects of credit migration that lend themselves to both empirical and theoretical research.

References

Altman, E., 1982, "Computerized Bond Rating Replication: Worthwhile or Futile?" *Journal of Bank Research*, 12, 250--268.

Altman, E., 1989, "Measuring Corporate Bond Mortality and Performance," *The Journal of Finance,* 44, 909--922.

Altman, E., 1998, "The Importance and Subtlety of Credit Rating Migration," *Journal of Banking and Finance,* 22, 1231--1247.

Altman, E., et al., 1992, "The Implications of Corporate Bond Ratings Drift," *Financial Analysts Journal,* 48, 64--75.

Altman, E., et al., 1997, "Credit risk measurement: Developments over the last 20 years," *Journal of Banking and Finance*, 21, 1721--1742.

Asquith, P., et al., 1989, "Original Issue High Yield Bonds: Aging Analyses of Defaults, Exchanges, and Calls," *The Journal of Finance*, 44, 923--952.

Black, F. and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *The Journal of Political Economy*, 81, 637--654.

Duffee, G., 1996, "On Measuring Credit Risks of Derivative Instruments," *Journal of Banking & Finance,* 805--833.

Duffee, G., 1998, "The Relation Between Treasury Yields and Corporate Bond Yield Spreads," *The Journal of Finance*, 53, 2225--2241.

Guttorp, P., 1995, *Stochastic Modeling of Scientific Data*, Chapman and Hall, London, U.K.

Jarrow, R., et al., 1995, "Pricing Derivatives on Financial Securities Subject to Credit Risk," *The Journal of Finance*, 50, 53--85.

Jarrow, R., et al., 1997, "A Markov Model for the Term Structure of Credit Risk Spreads," *The Review of Financial Studies*, 10, 481--523.

Kijima, M., 1997, *Markov Processes for Stochastic Modeling*, Chapman & Hall, London.

Kijima, M., et al., 1998, "A Markov Chain Model for Valuing Credit Risk Derivatives," *Journal of Derivatives*, 6, 97--108.

Mattson, A., et al., 1989, "A Simple Check of the Time Homogeneity of Markov Chains," *Journal of Forecasting,* 8, 65--72.

Tversky, A., et al., 1986, “Rational Choice and the Framing of Decisions,” *The Journal of Business,* 59, 251--278.

Table 1. Overall χ^2 Test Null Hypothesis: All Transition Probabilities are Time Invariant		
	FIRM-WGTD PROBABILITIES	**DOLLAR-WGTD PROBABILITIES**
ENTIRE SAMPLE ('76-'98) **7 Credit Categories**		
χ^2 Test Statistic	3570.63*	19066.85*
χ^2 Critical Value @ 5% = 1104.74		
FINAL 7 YEARS ('92-'98) **9 Credit Categories**		
χ^2 Test Statistic	1656.10*	5583.27*
χ^2 Critical Value @ 5% = 538.39		

*Reject the null hypothesis of time invariance at the 5% level of significance

Table 2. Testing of Time Homogeneity Over Various Investment Horizons
Results of χ^2 Test Based on 7 Credit Rating Categories

Matrix #1 Year End 1/31	Matrix #2 Year End 1/31	Firm-Weighted Matrices	Dollar-Weighted Matrices
1977	1980	14.01	**70.25***
1980	1983	40.43	**148.89***
1983	1986	**76.83***	**76.83***
1986	1989	**68.55***	**159.18***
1989	1992	57.33	**384.45***
1992	1995	**105.19***	**347.14***
1995	1998	**68.58***	**161.68***
1980	1985	60.83	**213.31***
1985	1990	**89.13***	**192.90***
1990	1995	**93.53***	**237.01***
1977	1987	**78.58***	**266.45***
1985	1995	**178.27***	**389.43***
1987	1997	**184.09***	**390.95***
Chi-Square Critical Values @ 5%		**66.34**	**103.01**

*Reject the Null Hypotheses of Time Invariance at the 5% Level of Significance

Table 3. χ^2 Test of Industrial vs. Utility vs. Financial Firms

Null Hypothesis: All Transition Probabilities are the Same Across Industries

FIRM-WEIGHTED PROBABILITIES

	χ^2 Test Statistic	
	7 Credit Categories	**9 Credit Categories**
1992/93	94.84	102.15
1993/94	110.70	110.72
1994/95	96.38	99.00
1995/96	108.04	117.45
1996/97	116.46	129.91
1997/98	120.28	125.55
χ^2 Critical Value @ 5%	***122.11***	***192.70***

DOLLAR-WEIGHTED PROBABILITIES

	χ^2 Test Statistic	
	7 Credit Categories	**9 Credit Categories**
1992/93	216.16*	221.22*
1993/94	237.62*	237.62*
1994/95	299.63*	306.97*
1995/96	343.41*	368.66*
1996/97	362.14*	381.10*
1997/98	326.08*	354.42*
χ^2 Critical Value @ 5%	***122.11***	***192.70***

**Reject the Null Hypothesis of Industry Invariance at the 5% Level of Significance*

Table 4. Chi-Square Tests of Aging Affects

Based on Dollar-Weighted Matrices and 7 Credit Rating Categories

One Step Transition Probabilities
1 Yr Since Issue vs. >1 Yr Since Issue

	Test Statistic	P-Value
1992/93	26.13	99.70%
1993/94	111.15	0.00%
1994/95	184.28	0.00%
1995/96	52.33	36.61%
1996/97	56.56	21.36%
1997/98	59.20	15.09%

Three Step Transition Probabilities
3 Yrs Since Issue vs. >3 Yrs Since Issue

	Test Statistic	P-Value
1992/95	52.31	34.68%
1993/96	65.66	5.60%
1994/97	65.94	5.35%
1995/98	27.55	99.43%

Five Step Transition Probabilities
5 Yrs Since Issue vs. >5 Yrs Since Issue

	Test Statistic	P-Value
1992/97	48.52	49.24%
1993/98	49.83	44.01%

Highlighted numbers with p-values ≤ 5% indicate a rejection of the null hypothesis at the 5% level

Figure 1
Industry Mix of Sample Firms in Select Years