

Testing Simple Markov Structures for Credit Rating Transitions

Nicholas M. Kiefer and C. Erik Larson

OCC Economics Working Paper 2004-3

Abstract

Models abound that analyze changes in credit quality. These models are designed to determine the reserves and capital needed to support the risks of individual credits as well as portfolios of credit instruments. Historical information on the transition of credit exposures from one quality level, or rating, to another is often used to estimate models that describe the probabilistic evolution of credit quality. A popular specification is the simple, time-homogeneous Markov model. While the Markov specification cannot describe credit processes in the long run, it may be useful for describing short-run changes in portfolio risk. In this convenient specification, the entire stochastic process can be characterized in terms of estimated transition probabilities. However, the simple homogeneous Markovian transition framework is restrictive. We propose a simple test of the null hypotheses of time-homogeneity that can be performed on the sorts of data often reported. The test is applied to data sets on municipal bonds, commercial paper, and sovereign debt. We find that municipal bond ratings transitions are adequately described by the Markov model for up to five years, that commercial paper on a 30-day transition scale seems Markovian up to six months (the extent of the available data), and that the transitions of sovereign debt ratings are adequately described by the Markov model (a result that may derive from the limited data of small sample sizes).

The views expressed in this paper are those of the authors alone, and are not necessarily those of the Office of the Comptroller of the Currency or the Department of the Treasury. Nicholas M. Kiefer is a Professor in the Departments of Economics and Statistical Sciences, Cornell University, Uris Hall, Ithaca, NY 14853. C. Erik Larson is a Senior Financial Economist in the Office of the Comptroller of the Currency's Risk Analysis Division, Washington, DC 20219.

Testing Simple Markov Structures for Credit Rating Transitions

Nicholas M. Kiefer

C. Erik Larson

Introduction

Models abound that analyze changes in credit risk. The ultimate aim of these models is to determine the reserves and capital needed to support the risks of not only individual credits but also portfolios of credit instruments. Credit quality is increasingly denoted by one or more ratings that summarize, over some specified horizon, a credit's probability of default, rate of loss given default, or both. For example, the proposed Basel II agreement (2004) requires institutions to rate assets by their one-year probability of default and by their expected loss severity given default.

Banks, supervisors, and other institutions rely upon these systems to produce accurate, stable representations of the risks of credit loss for current and future populations of similar credit exposures. Financial institutions use this information in portfolio selection. Historical information on the transition of credit exposures from one quality level or rating to another is used to estimate various models that describe the probabilistic evolution of credit quality.

A popular specification is the simple, time-homogeneous Markov model. With this specification, the stochastic processes can be specified completely in terms of transition probabilities. These correspond nicely to summary data that are often available and reported (though not without problems, as we note below). In particular, details on the history of individual assets are not required under this specification. Suppose we accept the simple Markov chain specification for describing credit rating transitions. Is there anything we can say about how to check the simple Markov structure using the sort of summary data that are commonly published (i.e., the transition tables for one, two, and five years for rated corporate bonds)? Important restrictions on credit

ratings transitions are implied by the simple Markov specification, and these restrictions suggest methods by which specification tests of the model can be made.

Given the relative ease with which the simple Markov model can be estimated or manipulated, it is not surprising that several practitioners have embedded the Markov framework into their credit quality tracking and risk assessment/management models. However, the adoption of the simple Markovian transition framework is not without cost; restrictive assumptions on the nature of credit transitions are required for the validity of this simple model. In fact, many of these assumptions are unrealistic, and are likely to be violated by the types of credit transitions considered by practitioners. Therefore, it would appear that a diagnostic test for the validity of the simple Markov chain model specification would be a valuable tool to add to the credit risk modeler's toolkit. In the next section, we illustrate how such a test can be constructed, and we discuss its statistical properties.

We consider testing time-homogeneity. This is not required by the general Markov specification, but seems to be a featured assumption in practice (and is an assumption that makes much empirical work possible!). Time-homogeneity means that the transition matrix P whose ij th element is the probability that a loan is in state j next period given that it is in I this period, is constant over time. This is a strong assumption. For example it might be thought that these transition probabilities would depend on macroeconomic conditions, or conditions specific to the industrial sector of the loan. Time-homogeneity is sometimes referred to as the property of having “stationary transition probabilities.” This is probably bad terminology, as it may be confused with stationarity of the stochastic process determined by these transition probabilities. Markov chains are not usually stationary, in the sense that the joint distribution of N successive observations may be different depending on where the N are taken. Nevertheless, a test for time homogeneity of ratings transitions is one diagnostic test that can be used to evaluate the adequacy of the simple Markov specification.

This paper proceeds as follows: section two introduces the simple Markov chain model, and describes how its parameters might be estimated with the sorts of data frequently available. Section three then develops a test of the null hypotheses of time-homogeneity that is implied by the simple

Markov specification. Before proceeding, it is useful to put the goals in perspective. First, credit rating transitions cannot “really” be Markovian. To see this, note that a key probability is the transition into default. Default is an absorbing state in any sensible specification. If changes in credit ratings could be strictly described as an absorbing Markovian system, all assets would eventually be in default. This is ridiculous. However, this is not the issue. The issue instead is whether the Markovian specification adequately describes credit rating transitions over rather short periods. Our test will allow researchers to determine which assets display Markovian characteristics and over what periods.

The Markov Chain Model

One relatively simple probability model for ratings transitions that is increasingly being used by financial practitioners is the Markov chain model. In the simple, discrete Markov chain model, the states that a stochastic process X_t may occupy at (discrete) time t form a countable or finite set. It is convenient to label the states by the positive integers. We will denote the number of unique states in a finite-dimensional Markov chain by the integer K . The probability of $X_{t+1} = j$, given that $X_t = i$, called a one-step transition probability, is denoted $p_{ij}(t)$. This representation emphasizes that transition probabilities in general will be functions of both the initial and the final states, and the time of transition. When the transition probabilities are independent of the time variable (the usual case in financial applications), the Markov process is said to have stationary transition probabilities. In this case we may write $p_{ij}(t) = p_{ij}$. It is common to organize the transition probabilities among all states into a transition matrix of the form

$$P = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1K} \\ p_{21} & p_{22} & \cdots & p_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ p_{K1} & p_{K2} & \cdots & p_{KK} \end{bmatrix}$$

Note that for each row, the sum of transition probabilities will be required to equal one, $\sum_{j=1}^K p_{ij} = 1$. This is because the state must either remain unchanged (with probability p_{ii}), or move to one of the alternative states.

When all possible states for the Markov process are similarly organized in a vector X , the equation describing the probabilistic evolution of the Markov process across states can be represented as:

$$X_{t+1} = X_t P$$

The simplicity of this last equation explains why the simple Markov chain model has proven so attractive to those seeking to describe the ratings transition behavior of credit. Note the interpretations possible: If X_t identifies the risk category of a single asset, for example $X_t = (0,0,1,0 \dots,0)$, say, indicating the asset is in the third risk group, then $X_{t+1} = (p_{31}, p_{32}, \dots, p_{3K})$ is the probability distribution across risk states in time $t+1$ for an asset in risk state 3 at time t .

Alternatively, if X_t is itself the distribution of assets across risk categories at time t , then X_{t+1} is the corresponding distribution in time $t+1$. The same probabilistic framework applies to single assets and the portfolio as a whole.

In the non-homogeneous case the transition probability matrix should be indexed by t , P_t and the transition equation is $X_{t+1} = X_t P_t$. In this case the model remains Markovian in that the distribution across states in the next period depends only on the distribution in the current period (and not on the distributions in previous periods). That is, the transitions are memoryless. History does not matter given the current state. In the credit ratings context, this result is equivalent to saying that given a credit's current rating, the likelihood that the credit will move to any other rating level, or that it will keep its current one, is independent of its past ratings history. While an interesting characteristic of Markovian transitions, this property seems to run against the widely held view that actual credit ratings changes exhibit "momentum." Conditional on the current rating level, a future credit downgrade is perhaps more likely if the credit had experienced a downgrade in the previous period than if it had experienced an upgrade or no change in rating; see, for example Carty and Fons (1993) or Bangia, Diebold, Kronimus, Schagen, and Schuermann (2002).

Consider obtaining the distribution across states in period $t+2$, given the distribution in period t . By substitution,

$$X_{t+2} = X_{t+1}P = X_tP^2$$

in the time-homogeneous case. In general, $X_{t+m} = X_tP^m$. In the non-homogeneous case we have

$$X_{t+2} = X_{t+1}P_{t+1} = X_tP_tP_{t+1}$$

and $X_{t+m} = X_tP_tP_{t+1}\dots P_{t+m-1}$.

Estimation of Markov Chain Transition Parameters

Estimation of the transition probabilities in a simple Markov chain can be carried out in a straightforward manner by counting the number of changes from one state to another that occur during a specified sample period. The estimation and inference problem here is classical; leading early papers are Anderson and Goodman (1957), Billingsley (1961), and the application-focused paper Chatfield (1973). We can for demonstrative purposes consider a simple two-state chain, which in our applied context might have credit quality states labeled low-risk and high-risk. For simplicity in this discussion, we do not include an absorbing state.

We will specify the single-period transition matrix as

$$P = \begin{bmatrix} p_{11} & 1 - p_{11} \\ 1 - p_{22} & p_{22} \end{bmatrix}$$

Let n_{ij} be the number of times, in a sample of size N that there is a move from state i to state j . The log-likelihood function for the collected data under the assumed Markov representation is then given by

$$\ln L(P | 1\text{-period data}) = n_{11} \ln p_{11} + n_{12} \ln(1 - p_{11}) + n_{21} \ln(1 - p_{22}) + n_{22} \ln p_{22}$$

Maximization of this function with respect to the unknown parameters p_{11} and p_{22} in this case yields closed form solutions for the maximum-likelihood estimators. These estimates are given by

$$\hat{p}_{11} = n_{11} / (n_{11} + n_{12}), \quad \hat{p}_{22} = n_{22} / (n_{22} + n_{21})$$

These estimators and their obvious generalizations to the K-state case are appropriate when we observe one-step transitions.

Suppose we have data on two-period transitions, and we wish to estimate p_{11} and p_{22} , the one-period transition probabilities in a time-homogeneous Markov chain. First, note that the two-period transition probabilities, denoted $p_{11}(2)$ and $p_{22}(2)$, are each functions of the underlying one-period transition probabilities through the relation $P(2) = P^2$. Specifically,

$$P(2) = \begin{bmatrix} p_{11}(2) & 1 - p_{11}(2) \\ 1 - p_{22}(2) & p_{22}(2) \end{bmatrix} = \begin{bmatrix} p_{11}^2 + (1 - p_{11})(1 - p_{22}) & 1 - \dots \\ 1 - \dots & p_{22}^2 + (1 - p_{11})(1 - p_{22}) \end{bmatrix}$$

Letting $n_{11}(2)$ be the number of observations remaining in state 1 for two periods, $n_{12}(2)$ the number moving from 1 to 2, etc., the loglikelihood for estimating p_{11} and p_{22} from two-period transition data is

$$\ln L(P | 2 - period\ data) = n_{11}(2) \ln p_{11}(2) + n_{12}(2) \ln(1 - p_{11}(2)) + n_{21}(2) \ln(1 - p_{22}(2)) + n_{22}(2) \ln p_{22}(2)$$

where of course the $p_{ij}(2)$ are defined above in terms of the p_{ij} . At this point we note that maximizing this loglikelihood with respect to the $p_{ij}(2)$ is easy, leading to the obvious estimators in terms of the sample fractions. This will be useful shortly. When one-period and two-period data are available, the total likelihood is the sum of the loglikelihoods

$$\begin{aligned} \ln L(P | 1 \& 2 - period\ data) &= n_{11} \ln p_{11} + n_{12} \ln(1 - p_{11}) + n_{21} \ln(1 - p_{22}) + n_{22} \ln p_{22} \\ &+ n_{11}(2) \ln p_{11}(2) + n_{12}(2) \ln(1 - p_{11}(2)) + n_{21}(2) \ln(1 - p_{22}(2)) + n_{22}(2) \ln p_{22}(2) \end{aligned}$$

which should be maximized with respect to p_{ij} . Extension to more than two periods and more than two states is immediate. While the lack of data at the intervening period might at first seem strange, it is actually a very common feature of ratings data. For instance, published ratings might be produced only annually, while the credit risk modeler might believe the Markov transition process to have a monthly or quarterly frequency. Alternately, the modeler might think that the Markov transition process has an annual frequency, but only cumulative transition data over a multi-year period is available for analysis.

The basic parameters are the separate one-period transition matrices P_1, P_2 , up to P_T , where T is the longest time length for which transitions are observed. Note that this does not correspond directly to the way in which the data are observed. Observed, typically, are one-period, two-period, up to perhaps 15-period transitions. These are quite different from one-period transitions observed at time 1, time 2, up to time 15. When multiple single-period transitions are observed, the above methods can be used to estimate period-specific parameters, and a chi-squared statistic can be formed to test the equality of parameters across periods. This is explored by Anderson and Goodman, and treated recently in the credit transition application by Thomas, Edelman and Crook (2002).

In our setting, when the summary statistics consist of multiple-period transition information, the problem is simplified by a reparametrization to the multiple-period transition probabilities. Of course these depend on the same underlying parameters P_1 through P_T . This dependence is complicated, since the t -step transition data $n_{ij}(t)$ depend on the t -step transition matrix, which is $P_1 P_2 \dots P_{t-1}$. For our purposes, however, separate estimates of the P_t are not really needed. What is needed is the maximized value of the loglikelihood function. This is easily obtained by reparametrizing to $P(1)=P_1, P(2) = P_1 P_2, P(3) = P_1 P_2 P_3$, etc. The loglikelihood is easily maximized with respect to these new parameters. Indeed the MLE's are simply

$$\hat{P}_{ij}(t) = n_{ij}(t) / \sum_{k=1}^K n_{ik}(t).$$

Solving for the underlying time-dependent one-step transition probabilities is difficult but unnecessary. The maximized value of the loglikelihood is clearly identical, whether parametrized in terms of the $P(t)$ or the P_t .

In the application to ratings transitions, the worst state is the absorbing state, default. With the convention that state 1 is the best (AAA in S&P and Fitch, Aaa in Moody's, etc.) and state K the worst, we have that the Kth row of P is $(0,0,\dots,1)$; there are no transitions out of default. Thus, a $K \times K$ transition matrix has $(K-1)^2$ parameters. One row is known, and one column is restricted since the row sums are one. There are other restrictions as well ($p_{ij} \in [0,1]$) which make a reparametrization convenient before maximization. The point is that the number of parameters is not unduly large in practical applications and it does not increase with the number of periods available in the time-homogeneous case.

The Likelihood Ratio Test

We are now in a position to test for the hypothesis that the transition data are generated by a time-homogeneous Markov chain. To test, we will use a likelihood ratio statistic that compares the restricted likelihood computed under the null that the transitions are generated by a time-homogeneous Markov chain, with an alternative likelihood for the observed transitions computed without imposing the homogeneity restrictions. The statistic is computed as $-2\ln LR$, where

$$\ln LR = \max_P \ln L(P | \text{all data}) - \max_{P(1), \dots, P(T)} \ln L(P(1), \dots, P(T) | \text{all data})$$

Using standard arguments, this statistic is asymptotically $\chi^2(q)$, where the degrees of freedom q is the number of restrictions imposed. Suppose there are K states and transitions 1, 2, through T periods. Then the restricted model has $(K-1)^2$ parameters, as we have seen, and the unrestricted model has $T(K-1)^2$ parameters. Hence $q = (T-1)(K-1)^2$.

Computation of the unrestricted maximized loglikelihood is trivial, as closed form solutions exist. Computation of the restricted MLE is a little more difficult when multiperiod transitions are available. We have found it useful to reparametrize with a logit transform to new parameters

$$\alpha_{ij} = \exp(p_{ij}) / \sum_{j=1}^K \exp(p_{ij})$$

and to maximize numerically with respect to the $(K-1)^2$ “free” parameters α_{ij} . This is a fairly easy maximization using any of many available software packages. Convenient starting values are provided by the logit transforms of the sample fractions making up the 1-period transition data. Finally, we rule out zeros among the transition probabilities, taking the position that no transition, however unlikely, is really impossible. To this end, we bound the estimated transition probabilities from below by 10^{-10} .

As the likelihood ratio test against the general class of non-time homogeneous models, the statistic can be expected to have maximum local power against this broad class of alternatives. As usual, more power can be achieved if the class of alternatives is further restricted, but this seems of little interest. It is useful to consider whether the test has power against other departures from the time-homogeneous Markov model. Departures of interest include alternatives with more dependence than the Markov specification allows. For example, does next period’s distribution of assets across risk categories depend not only on the current distribution but on previous distributions as well? If so, the process is not Markovian. Our test will have power against these alternatives if these processes can be better approximated by a time-inhomogeneous process than by a time-homogeneous process. While no proof is provided, this property certainly seems likely, so it can be expected that our test has power against a fairly broad class of alternatives to the time-homogeneous Markov specification.

Data and Applications

While several institutions have their own historical data and ratings assignments with which to estimate Markov transition models, several alternative sources exist from which credit rating transition matrices might be obtained. Transition matrices are regularly published by the major rating agencies, including Moody's and Standard and Poor's; they are also available from risk management consulting companies and vendors of credit risk models including RiskMetrics, KMV, and Kamakura, Inc. Unfortunately, reporting standards in this area are fairly lax. Often, transition

fractions are reported with no indication of the size of the underlying data set, or the initial distribution of assets across risk categories. Since many of these transitions are low-probability events, large samples are necessary for precise estimation. This information should be routinely provided as a matter of sound statistical practice.

We have assembled several data sets with which to illustrate our technique and demonstrate its usefulness. These are:

1. Municipal Bonds

These data are from a well-known Standard and Poor's study published on the web (Standard & Poor's, 2001). This data set contains information on rating transitions of municipal bonds through eight rating categories and a nonrated category from 1986 through 2000. Average one-year, two-year, and through 15-year transition matrices are reported. This is a rare study that gives sufficient detail to be realistically useful. We have run tests based on the one-year and two-year data, the one-, two-, and three-year data, and so on up to the full data set including one-, two-, ...up to 15-year transitions. Results are shown in Table 1. We see here that the Markov specification is adequate for describing annual ratings transitions for periods up to about five years. The model is fairly reliable for one-, two-, three-, and even four-year transitions, but begins to fail when looking more than 5 years out. The implication for practice is that municipal bond rating transitions can over reasonable periods be described by a simple time-homogeneous model, but the transition probabilities should probably be updated, and the new ones used, every few years.

Table 1: Municipal Bonds 1986-2000

Transitions	LR Statistic	DF	Pr($\chi^2 > LR$)
1,2	42.96	49	0.715
1,2,3	92.33	98	0.642
1 to 4	154.2	147	0.327
1 to 5	237.6	196	0.022
1 to 6	373.9	245	2.07E-07
1 to 7	526.8	294	2.00E-15
1 to 8	714.7	343	0
1 to 9	972.6	392	0
1 to 10	1107	441	0
1 to 11	1567	490	0
1 to 12	2056	539	0
1 to 13	2490	588	0
1 to 14	2866	637	0
1 to 15	3177	686	0

Calculations based on Standard & Poor's (2001)

2. Commercial Paper

These data are from a study by Moody's (2000). This study is also commendable in the detail provided. Commercial paper defaults are extremely rare, although there is some migration among rating categories (here P-1, P-2, P-3, and NP; P is for prime). Moody's goal is that commercial paper with any prime rating should never default. Since commercial paper, in contrast to the municipals studied above, are short-term assets, the transitions examined are 30-day transitions. The Moody's study reports transitions over 30, 60, 90, 120, 180, 270 and 365 days. Note that the "missing" powers of P at 150, 210, 240, 300, and 330 days present no problem for our methods. The data up to the 180-day transitions exhibit expected patterns— for example, the transitions into default (and generally transitions out of the initial state) increase with the length of the time span.

However, for 270 and 365 days, the transitions from NP into default are zero. This appears problematic and is perhaps due to definitions Moody's has used and the fact that commercial paper is rarely extended for long periods. We have chosen simply not to use the data for 270 and 365 days. Another question arises as to the treatment of Moody's

category WR – withdrawn. These withdrawals occur because commercial paper is not rolled over, so the asset size becomes negligible, or the market has otherwise lost interest in the offering. It is apparently not a synonym for default or a decline in creditworthiness. Consequently, we treat these as censored. Again, to illustrate the use of the test, we calculate the statistic based on increasing numbers of periods. Results are given in Table 2.

Table 2: Commercial Paper 1972-1999

Transitions	Chi-Square	DF	P-value
1,2	14.95	16	0.529
1,2,3	18.83	32	0.969
1.2.3.4	23.78	48	0.999
1 to 4, 6			

Calculations based on Moody's (2000)

There is no serious evidence against the Markov specification. Recall that the time scale here is 30 days; commercial paper is generally short lived. The Markov model “works” over the period available (six months). Six months is probably also the relevant period for applications. Note that the sample size here is enormous, so these results are fairly firm. For a different interpretation of the failure to reject, see our next discussion on sovereign debt.

3. Sovereign Debt

These data are from Standard & Poor’s (2003). Again, suitable detail is provided. Here, it appears that overlapping cohorts were used to calculate the transition matrices. For example, in a sample of 15 years, there is 1 15 year cohort, 2 14, 3 15, etc. Transitions among eight rating categories including default are recorded. In the case of sovereign debt, default must be defined particularly carefully, because some of the debt to preferred creditors is often serviced after default. S&P discusses this complication in appendix 2 of its update on defaults. One-, three-, five-, and seven-year transition rates are reported. Our results are reported in Table 3.

Table 3: Sovereign Debt 1975-2002

Transitions	Chi-Square	DF	P-value
1,3	17.33	49	1.0000
1,3,5	44.32	98	1.0000
1,3,5,7	85.14	147	1.0000

Calculations based on Standard & Poor's (2003)

There is simply no rejecting the time-homogeneous Markovian specification. It should be noted that the sample sizes are relatively small here (all well under 100), and this is perhaps driving the result. Nevertheless, it is doubtful that a more complicated model would be supported by the data.

There are many other studies reporting and discussing transition rates among risk ratings. We make no claim to have uncovered even a small fraction of the competent studies. Unfortunately, many studies do not report sufficient information for the reader to calculate sample sizes. This is poor statistical reporting, and readers should naturally treat such reports with some skepticism. Ratings changes do not occur often and some are very rare. Thus, the transition rates reported are essentially estimators of very small probabilities. The sampling standard error of an estimator \hat{p} of a binomial probability p is approximately $(p(1-p)/n)^{1/2}$. With $p = 0.0002$, say, the standard error of the estimator (with expected value 0.0002) at a sample size $n=1000$ is 0.00045, more than twice the expected value of the estimator. It is not until $n=5000$ that the standard error is even equal to the expected value of the estimator. Thus, large sample sizes are required for precise estimation of small (or large) probabilities. This fact probably explains the failure to reject time homogeneity in the case of sovereign debt. For such debt, sample sizes are less than 100 and many transitions are low-probability events. No test can be expected to have much power. This is however, essentially the correct inference. The lack of data because of small sample sizes suggests that richer models cannot be supported.

Conclusions

The time-homogeneous Markov model for transitions among risk categories is widely used in areas from portfolio management to bank supervision and risk management. It is well-known that these models can be overly simple as descriptions of the stochastic processes for riskiness of assets. Nevertheless, the model's simplicity is extremely appealing. We propose a likelihood ratio test for the hypothesis of time-homogeneity. Because reparametrization is convenient, the test is simple to compute, requiring numerical estimation of only the restricted model, a $(K-1)^2$ -parameter problem where K is the number of risk categories. The test can be based on summary data often reported by rating agencies or collected within banks. We recommend that the test be interpreted as determining whether or not transitions over particular periods can be adequately modeled as Markov chains. (We do not recommend using the test to determine whether the underlying process is Markovian. It is not: the prediction would be that everything eventually defaults.)

We believe that some transitions can be usefully modeled as Markovian for several years. For example, municipal bond ratings can be adequately described as time-homogeneous Markovian on an annual scale over as much as five years. For transitions over a longer period, a different model should be used. As a practical matter, we conclude simply that the estimated transition matrix be updated regularly and then used to forecast for up to five years. For a sample of commercial paper obligations on a 30-day time scale, we find that the time-homogeneous Markov model adequately describes transitions for at least six periods (the length of our available data). This result is fairly firm, as the data set involved is large, though of course rating transitions over 30-day periods are rare. In an application to sovereign debt with transitions on an annual scale, we cannot reject the Markov specification, but qualify this result heavily because the sample sizes are small. That is, the data do not indicate that the Markov model can be improved, but they do not indicate that it works very well either; there is simply very little information on these rare transitions.

References

Anderson, T.W. and L. A. Goodman. "Statistical Inference About Markov Chains," *Ann. Math. Statist.*, 28 (1957) , 89-110.

Bangia, A., F. Diebold, A. Kronimus, C. Schagen, and T. Schuermann (2002). "Ratings Migration and the Business Cycle, with Applications to Credit Portfolio Stress Testing," *Journal of Banking and Finance*, 26(2-3), 445-474.

Basel Committee on Banking Supervision (June 2004). "International Convergence of Capital Measurement and Capital Standards: A Revised Framework," Bank for International Settlements.

Billingsley, P. "Statistical Methods in Markov Chains," *Ann. Math. Statist.*, 32 (1961), 12-40.

Carty, Lea V. and Jerome S. Fons. "Measuring Changes in Corporate Credit Quality," *Moody's Special Report* (November 1993), New York: Moody's Investors Service.

Chatfield, C. "Statistical Inference Regarding Markov Chain Models," *Applied Statistics* 22:1 (1973), 7-20.

Karlin, S. and Taylor, H. M. *A First Course in Stochastic Processes. Second edition.* Academic Press, New York, 1975.

Berthault, D. Hamilton, and L. Carty. "Commercial Paper Defaults and Rating Transitions, 1972-2000," *Global Credit Research, Special Comment* (October 2000), New York: Moody's Investors Service. <http://riskcalc.moodysrms.com/us/research/trans/60917.asp>.

Standard & Poor's (2001a), "Default Research: Municipal Rating Transitions and Defaults." <http://www2.standardandpoors.com/NASApp/cs/ContentServer?pagename=sp/Page/HomePg>

Standard & Poor's (2001b), "Default Research: Municipal Rating Transitions and Defaults: Related Tables 4-5." <http://www2.standardandpoors.com/NASApp/cs/ContentServer?pagename=sp/Page/HomePg>

Standard & Poor's (2003), "2002 Defaults and Rating Transition Data Update for Rated Sovereigns."

Thomas, L. C., D. B. Edelman, and J. N. Crook (2002). *Credit Scoring and its Applications*, SIAM, Philadelphia.