A simulation estimator for testing the time homogeneity of credit rating transitions ☆

Nicholas M. Kiefer [a,b,*], C. Erik Larson [c]

[a] Departments of Economics and Statistical Sciences, Cornell University, United States
[b] Risk Analysis Division, Office of the Comptroller of the Currency, US Department of the Treasury, United States
[c] Fannie Mae, United States

Abstract

The measurement of credit quality is at the heart of the models designed to assess the reserves and capital needed to support the risks of both individual credits and portfolios of credit instruments. A popular specification for credit-rating transitions is the simple, time-homogeneous Markov model. While the Markov specification cannot really describe processes in the long run, it may be useful for adequately describing short-run changes in portfolio risk. In this 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 test of the null hypotheses of time-homogeneity that can be performed on the sorts of data often reported. We apply the tests to 4 data sets, on commercial paper, sovereign debt, municipal bonds and S&P-rated Corporates. The results indicate that commercial paper looks Markovian on a 30-day time scale for up to 6 months; sovereign debt also looks Markovian (perhaps due to a small sample size); municipals are well-modeled by the Markov specification for up to 5 years, but could probably benefit from frequent updating of the estimated transition matrix or from more sophisticated modeling, and S&P Corporate ratings are approximately Markov over 3 transitions but not 4.
© 2006 Elsevier B.V. All rights reserved.

* Corresponding author. Departments of Economics and Statistical Sciences, 490 Uris Hall, Cornell University, Ithaca, NY 14853-7601, US.
E-mail addresses: nmk1@cornell.edu (N.M. Kiefer), c_erik_larson@fanniemae.com (C.E. Larson).

1. Introduction

The measurement of levels and changes in credit quality is at the heart of the development of a whole spectrum of models designed to assess and support the dynamic management of credit risky assets. Several of these models are specifically designed to estimate the reserves and capital needed to support the risks of both individual credits and portfolios of credit instruments as functions of credit quality. Credit quality can be expressed in a variety of ways, but it is increasingly common to have credits assigned 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 (Basel Committee on Banking Supervision, 2004) requires institutions to rate assets by their 1-year probability of default and by their expected loss severity given default.

Banks, supervisors and other institutions are relying 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. As an example, we apply our methods to migration matrices for commercial paper over horizons of 30, 60, 90, 120, 180 and 270 days. 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 summary data that are commonly published? 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, 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 ijth element is the probability that a loan is in state j next period given 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 cause confusion with the 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 they are taken. Nevertheless, a test for time