Simulating Continuous Time Rating Transitions

Robert A. Jones

17 March 2003

This note describes how to simulate state changes in continuous time Markov chains. An important application to credit risk is the evolution of a borrower's credit rating. Default counts as a rating. Potential CFS contracts have payments contingent not only on time of default, but also on credit rating immediately prior (e.g., through effect on collateral support agreements). Thus a means of simulating the entire process is needed.

1. The continuous time transition generator

In discrete time Markov chains, the likelihood of change between n states is described by a $n \times n$ *transition matrix* $P \equiv [p_{ij}]$. Element p_{ij} is the probability of transition from state i at time t to state j at time $t+1$. Elements of P are non-negative and each row sums to 1.

In continuous time, the process is described by a *generator matrix* Q. Off-diagonal elements of Q are non-negative and each row sums to 0. The corresponding transition matrix over an interval of length t is e^{tQ}. The exponential of a matrix is defined in the same manner as the exponential function of a scalar: by series expansion. That is,

$$e^{tQ} = I + tQ + (tQ)^2/2! + (tQ)^3/3! + \dots \tag{1}$$

in which I denotes the identity matrix. As can be seen, setting $t = 1$ and supposing Q sufficiently small that the quadratic and higher terms can be ignored, Q is approximately $P - I$. The higher order terms capture the possibility of multiple transitions within the interval t. An absorbing state, such as default, would have all zeros in its row.

Note this series expansion permits ready computation of survival and default probabilities over intervals of arbitary length. If state n corresponds to default, then the n^{th} column of e^{tQ} gives the probability of default by time t when starting from each of the respective states at time 0.[1]

Off-diagonal elements of Q have the following interpretation. They are the intensities of independent Poisson processes of transition from state i to state j. Diagonal element q_{ii} is (the negative of) the intensity of arrival of transition to any state other than i. This suggests how to efficiently simulate the process:[2]

[1] For matrices typical of one year credit rating transitions, terms in the series have all matrix elements less than 10^{-9} by the fifteenth term.

[2] This follows a method suggested by Duffie and Singleton (1998).

1. Starting from state i at $t = 0$, draw a uniform $[0, 1]$ random variable u. Time to transition from state i is computed as $\Delta t = \ln u/q_{ii}$. This is an exponentially distributed random variable with mean $-1/q_{ii}$.[3]

2. Given that a move has occurred, the probability that the move is to state $j \neq i$ is $q_{ij}/\sum_{k\neq i} q_{ik}$. Partition the unit interval into subintervals of these lengths for $j = 1, \ldots, n,\ j \neq i$. To determine which state the transition is to, now draw another uniform random variable v. The subinterval in which it falls gives the next state j.

3. If new state is default (and that is absorbing), or if transition date exceeds maturity T relevant for the contract at hand, this path is done. Otherwise increment the current date, return to first step, and draw the next transition time.

Note that this can be exceptionally fast for low transition intensities since large time increments would tend to occur in step one.

2. Estimating the generator from rating transition data

For the discrete time Markov chain, elements of the transition matrix have maximum likelihood estimates

$$\hat{p_{ij}} = \frac{N_{ij}}{N_i} \tag{2}$$

where N_i is the number of firms in rating category i at the beginning of a period and N_{ij} is the number of that population that migrate to category j by the end of the period. A significant consequence of this is that if no transition from i to j occurs in a period, then the corresponding probability estimate is 0. This, using annual periods, is the way rating agency data is typically presented.

Lando and Skodeberg (2002) observe that this method essentially discards data on the exact timing of transitions within the period. This both reduces statistical efficiency and also misrepresents the probability of rare events such as indirect moves from high ratings to default over a year. Based on Kuchler and Sorensen (1997), they suggest the following direct estimator of the generator from observation over the interval from 0 to T:[4]

$$\hat{q}_{ij} = \frac{N_{ij}(T)}{\int_0^T Y_i(s)\, ds} \tag{3}$$

Here $N_{ij}(T)$ is the total number of transitions from i to $j \neq i$ over the interval, and $Y_(s)$ is the number of firms in rating class i at time s. The numerator counts the number of transitions

[3]This is the random time of the first occurrence of a Poisson event with intensity q_{ii}.

[4]Lando and Skodeberg (2002), p. 427.

observed over the interval; the denominator counts the number of firm-years spent in state i.[5]

If one's interest was solely in the transition matrix over a unit interval, their estimate of that would be $e^{\hat{Q}}$, which would generally differ from $\hat{P}$ obtained from discrete time data. Christensen and Lando (2002) explore how levels and confidence intervals for these probabilities differ under the two methods. Confidence intervals using their method are generally tighter than using the rating agency method. They obtained complete data on transition times from S&P, which, alas, we do not have. However the method might be applied in other settings for which we do have data.

3. Finding the generator for a given transition matrix

Some of the improved efficiency of the Lando/Skodeberg method comes from using aspects of the data discarded in the transition matrix estimator method, but some comes from imposing (rightly or wrongly) more restrictions on the stochastic process. They assume that the annual matrix is the result of a time-homogeneous Markov process from moment to moment throughout the year. The set of admissible P is larger than the set of e^Q for admissible Q.

But what if all we have is P, are willing to make this assumption, and need an estimate of Q to implement the continuous time simulation? This is the 'embedding problem'. Two solutions are available in the literature. Jarrow, Lando and Turnbull (1997), in modelling the term structure of credit spreads through rating transitions, approached it through an 'approximating assumption'. They discard firms that moved to non-rated during the year, and assume that each remaining had at most one transition during the year. This results in the following:[6]

$$\hat{q}_{ii} = \ln p_{ii} \text{ for } i = 1 \dots n \tag{4}$$

$$\hat{q}_{ij} = p_{ij}(\ln p_{ii})/(p_{ii} - 1) \text{ for } j \neq i \tag{5}$$

The exponential of the resulting $\hat{Q}$ does not give back the starting P, but they are fairly close.

A better approach is taken by Israel, Rosenthal and Wei (2001). When the embedding problem has a solution, they show that it can usually be expressed by the following series

[5] The estimator for $\hat{q}_{ii} = -\sum_{i \neq j} \hat{q}_{ij}$.

[6] Jarrow, Lando and Turnbull (1997), p.505.

expansion.[7]

$$\tilde{Q} = (P - I) - (P - I)^2/2 + (P - I)^3/3 - (P - I)^4/4 + \ldots \tag{6}$$

It is interesting to note that this corresponds to the series expansion for the logarithm function in the case of scalars, i.e., the inverse of the exponential function.

When this expansion results in negative off-diagonal elements, some fix is needed. They suggest setting the offending elements to 0 and adding their total value to all remaining entries of the same row in proportion to the absolute values of the latter (remember the diagonal element is negative). I.e., letting

$$G_i = |\tilde{q}_{ii}| + \sum_{j \neq i} \max(\tilde{q}_{ij}, 0) \qquad B_i = \sum_{j \neq i} \max(-\tilde{q}_{ij}, 0) \tag{7}$$

be the good and bad totals for each row, set

$$q_{ij} = \begin{cases} 0 & i \neq j \text{and } \tilde{q}_{ij} < 0 \\ \tilde{q}_{ij} - B_i |\tilde{q}_{ij}| / G_i & \text{otherwise if } G_i > 0 \\ \tilde{q}_{ij} & \text{otherwise if } G_i = 0 \end{cases} \tag{8}$$

This way of distributing the probability is arbitrary but appears to work fairly well.

IRW compare how close is the exponential of the resulting Q to P using the norm $\sum_{i,j} |p_{ij} - e^Q_{ij}|$. In their examples based on agency annual rating transition matrices, the resulting distance is about one tenth that of the JLT approximation. The IRW method is the one we should favor.

Note that the method can be applied if the input transition matrix is for an interval different from one year. For example, if the input matrix P in equation (6) is for two years, then the annual rate generator is obtained by dividing the result by two. Indeed, comparing the Q so obtained from transition matrices for different intervals gives some indication of the validity of the time-homogeneity assumed.

4. Risk-neutral valuation

Given either an estimated annual transition matrix, or data on the actual transition times, the preceeding sections equip us to simulate the objective process of rating and default transitions in continuous time. Monte Carlo simulation could then estimate the objective, or actuarial, expected discounted value of a contract with payoffs contingent on such events.[8]

---

[7] Convergence of the series is assured if all diagonal elements of P exceed 0.5 . See Israel, Rosenberg, Wei (2001), p.247.

[8] Let us assume for the moment a deterministic riskless interest rate process.

The assumption of an arbitrage-free environment (with usual qualifications) implies existence of an equivalent probability measure over transition time paths such that equilibrium security prices equal expected discounted values under this risk-neutral measure.[9] The notion of *equivalent* only implies that the objective and risk-neutral measures agree on which events have measure zero.

This isn't much to work with. In lieu of building a plausible general equilibrium model, the typical shortcut is to impose a simple parametrization of the difference between the two measures and calibrate the values of the parameters to a few observed market prices. The first simplification is to suppose that the process for rating and default transitions is also Markov under the risk-neutral measure. Note that this is purely an assumption, and not implied by absence of arbitrage. With it, the risk-neutral generator matrix can still be any admissible generator Q^* that agrees with the objective Q on which elements are 0.

Some simple parametrizations of the difference between the two measures (giving the risk premium for bearing such risks) are the following, in order of increasing complexity: (a) Scale up all elements of Q by a constant factor λ. (b) Scale up each row of Q by its own state-specific factor λ_i. (c) Scale up each row of Q by a time and state-dependent factor $\lambda(t)$. The latter is used by JLT in their calibration of a risk-neutral Q^* to the observed term structure of credit spreads. The scale factor is any strictly positive number. Values greater than one typically imply a positive risk premium for bearing default risk (i.e., expected return exceeding the risk-free rate); values less than one a negative premium. Note that these parametrizations all preserve which elements are zero and non-zero in the matrix.

To illustrate, suppose we have a single price observation on a single firm (e.g., the credit default swap spread for one maturity only), know its current credit rating, and know the default-free term structure. We could adopt parametrization (a) and choose λ so that the swap has exactly zero expected discounted value using the transition generator $Q^* = \lambda Q$. If more price data is available, we could either stick with (a) and choose λ to give a best least-squares fit to all the data, or fit the larger number of parameters of a more complex parametrization. The usual dangers associated with over-parametrization of course apply.

References

Christensen, Jens, and David Lando, "Confidence sets for continuous-time rating transition data," Department of Statistics and Operations Research, University of Copenhagen, September 2002.

Harrison, Michael J., and David Kreps, "Martingales and Arbitrage in Multi-period Se-

[9]Michael Harrison and David Kreps (1979).

curities Markets," *Journal of Economic Theory*, 20 (1979), 381-408.

Israel, Robert B., Jeffrey S. Rosenthal and Jason Z. Wei, "Finding Generators for Markov Chains via Empirical Transition Matrices, with Application to Credit Ratings," *Mathematical Finance*, 11 (April 2001), 245-265.

Jarrow, Robert A., David Lando and Stuart M. Turnbull, "A Markov Model for the Term Structure of Credit Risk Spreads," *Review of Financial Studies*, 10, no. 2 (Summer 1997), 481-523.

Kuchler, U., and M. Sorensen, *Exponential Families of Stochastic Processes*, New York: Springer-Verlag, 1997.

Lando, David, and T. M. Skodeberg, "Analyzing rating transitions and rating drift with continuous observations," *Journal of Banking and Finance*, 26 (2002), 423-444.