Finite-State Approximation Of VAR Processes: A Simulation Approach

Stephanie Schmitt-Grohé* Martín Uribe†

July 2009
Revised November 2014

This paper presents a simulation approach to discretizing autorregressive models. The methods delivers a transition probability matrix that captures well the autocovariance function of the continuous-valued process. The main advantage of the approach is that it does not impose normality on the innovation of the VAR process. The method makes it easy to handle multivariate processes, as well as autorregressive processes of order higher than one. Matlab code is provided.

Consider the multivariate autorregressive process

$$x_t = Ax_{t-1} + \Omega\epsilon_t, \tag{1}$$

where x_t is a column vector of length m, A is an $m \times m$ matrix of constants, Ω is an $m \times r$ matrix of constants, and ϵ_t is an $r \times 1$ i.i.d. disturbance with mean $\emptyset$ and standard deviation I. All eigenvalues of A are assumed to lie inside the unit circle. The AR(1) process (1) could be the AR(1) represenation of an autorregressive process of order higher than one.

Discretize the support of element i of the state vector x_t using n_i points from $-\bar{x}^i$ to $+\bar{x}^i$. Simulation tests suggest that a reasonable value for $\bar{x}^i$ is $\sqrt{10}\sigma_i$, where σ_i is the unconditional standard deviation of element i of x_t.[1] Using grids with equally spaced points yields a satisfactory balance between simplicity and accuracy. Let $n = \prod_{i=1}^{m} n_i$ be the number of possible values of the discretized state. The discretized state, denoted S, is a $n \times m$ matrix. Element (i, j) of S indicates the value of the jth element of x_t in state i. Let Π be the $n \times n$ transition probability matrix associated with the state S. The purpose of this note is to compute Π. To this end, simulate the continuous-valued process (1) to create an artificial time series X of length T (after possibly burning a number of initial draws). Thus, the matrix X is of size $T \times m$. Let X^d be the discretized version of X. Specifically, for $t = 1, \ldots, T$, row t of X^d is given by the row of S with the minimum Euclidian distance to row t of X. Let Π_1 be an $n \times n$ matrix with all elements equal to zero. For $t = 1, \cdots, T$, define the matrix Π_t as follows: $\Pi_t(i, j) = \Pi_{t-1}(i, j) + 1$ if $S_i = X^d_{t-1}$ and $S_j = X^d_t$, and 0 otherwise, where S_i denotes the ith row of S. Finally, define Π as $\Pi(i, j) = \Pi_T(i, j)/\sum_{h=1}^{n} \Pi_T(i, h)$. The matlab script tpm.m produces the matrices Π and S given A, Ω, n_i, and T.

*Columbia University, CEPR, and NBER. E-mail: stephanie.schmittgrohe@columbia.edu.
†Columbia University and NBER. E-mail: martin.uribe@columbia.edu.

[1]Formally, σ_i is the square root of the ith diagonal element of the covariance matrix Σ, implicitly given by $\Sigma = A\Sigma A' + \Omega\Omega'$.