Chris Edmond ⟨cpedmond@unimelb.edu.au⟩

Notes on computing and simulating Markov chains

Consider a two-state Markov chain (x, P, π_0) with transition matrix

$$P = \begin{pmatrix} 1-p & p \\ q & 1-q \end{pmatrix}$$

for $1 > p, q > 0$. Then this Markov chain has a unique invariant distribution $\bar{\pi}$ which we can solve for as follows

$$0 = (I - P')\bar{\pi}$$

or

$$\begin{aligned} \begin{pmatrix} 0 \\ 0 \end{pmatrix} &= \left[\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} - \begin{pmatrix} 1-p & q \\ p & 1-q \end{pmatrix} \right] \begin{pmatrix} \bar{\pi}_1 \\ \bar{\pi}_2 \end{pmatrix} \\ &= \begin{pmatrix} p & -q \\ -p & q \end{pmatrix} \begin{pmatrix} \bar{\pi}_1 \\ \bar{\pi}_2 \end{pmatrix} \end{aligned}$$

Carrying out the calculations, we see that

$$\begin{aligned} 0 &= p\bar{\pi}_1 - q\bar{\pi}_2 \\ 0 &= -p\bar{\pi}_1 + q\bar{\pi}_2 \end{aligned}$$

These two equations only tell us one piece of information, namely

$$\bar{\pi}_1 = \frac{q}{p}\bar{\pi}_2$$

But we also know that these elements must satisfy

$$\bar{\pi}_1 + \bar{\pi}_2 = 1$$

So we can solve these two equations in two unknowns to get

$$\bar{\pi}_1 = \frac{q}{p+q}, \qquad \bar{\pi}_2 = \frac{p}{p+q}$$

Notice the following properties:

- As $q \to 0$, the state x_1 becomes a **transient state** and the state x_2 becomes an **absorbing state**: once the chain leaves x_1, it never returns. Since this happens with probability 1 if we run the chain long enough, the stationary distribution will be degenerate with $\bar{\pi} \to (0, 1)$.
- Similarly, as $p \to 0$, the state x_2 becomes a transient state and the state x_1 becomes an absorbing state and $\bar{\pi} \to (1, 0)$.
- If $p = q$, the chain is **symmetric** and the stationary distribution is just $\bar{\pi} = (0.5, 0.5)$. More generally, for any n state symmetric Markov chain, the uniform distribution with $\bar{\pi}_i = \frac{1}{n}$ is a stationary distribution.

Because any dependence on transient states washes out in the long run, it is often easy to simplify the computation of a stationary distribution. For example, if $n = 3$ and

$$P = \begin{pmatrix} 0.7 & 0.2 & 0.1 \\ 0 & 0.5 & 0.5 \\ 0 & 0.9 & 0.1 \end{pmatrix}$$

The state x_1 is transient (once you leave it, you never return), so the invariant distribution can be found by considering the sub-matrix in the lower right corner, namely

$$P_{\text{sub}} = \begin{pmatrix} 0.5 & 0.5 \\ 0.9 & 0.1 \end{pmatrix}$$

Letting $p = 0.5$ and $q = 0.9$ we see that P has stationary distribution with $\bar{\pi}_1 = 0$ and non-zero elements

$$\begin{aligned} \bar{\pi}_2 &= \frac{q}{p+q} = \frac{0.9}{0.5+0.9} = 0.6429 \\ \bar{\pi}_3 &= \frac{p}{p+q} = \frac{0.5}{0.5+0.9} = 0.3571 \end{aligned}$$

Now test your intuition. Is it obvious that a Markov chain with transition matrix

$$P = \begin{pmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & 0 & 1 \end{pmatrix}$$

has stationary distribution $\bar{\pi} = (0, 0, 1)$?

A. Computing stationary distributions in Matlab

If the transition matrix P is regular, with $1 > p_{ij} > 0$ for each i, j, then the Markov chain has a unique stationary distribution and it can be computed with brute force. Just take P^T for T a large number. Then you will get back a matrix with identical rows that are equal to the chain's stationary distribution. For example, if

$$P = \begin{pmatrix} 0.7 & 0.2 & 0.1 \\ 0 & 0.5 & 0.5 \\ 0 & 0.9 & 0.1 \end{pmatrix}$$

Then

$$P^{1000} = \begin{pmatrix} 0 & 0.6429 & 0.3571 \\ 0 & 0.6429 & 0.3571 \\ 0 & 0.6429 & 0.3571 \end{pmatrix}$$

and the stationary distribution is $\bar{\pi} = (0, 0.6429, 0.3571)$, which is what we would get from the analytic approach taken above. This is not very elegant. A neater approach is to use Matlab to compute the eigenvalues and vectors of P.

Step 1. Compute matrix of eigenvalues and eigenvectors of P' (remember the transpose!). In Matlab,

$$[\texttt{V},\texttt{D}] = \texttt{eig}(\texttt{P}')$$

gives a matrix of eigenvectors V and a diagonal matrix D whose entries are the eigenvalues of P'.

Step 2. Now since P is a transition matrix, one of the eigenvalues is 1. Pick the column of V associated with the eigenvalue 1. With the matrix P given above, this will be the second column and we will have

$$\texttt{v} = \texttt{V}(:,2) = \begin{pmatrix} 0 \\ -0.8742 \\ -0.4856 \end{pmatrix}$$

Step 3. Finally, normalize the eigenvector to sum to one, that is

$$\texttt{pibar} = \texttt{v}/\texttt{sum}(\texttt{v}) = \begin{pmatrix} 0 \\ 0.6429 \\ 0.3571 \end{pmatrix}$$

B. Simulating Markov chains in Matlab

In the problem set, you will have to simulate an n state Markov chain $(\mathbf{x}, P, \pi_0)$ for $t = 0, 1, 2, ..., T$ time periods. I use a bold $\mathbf{x}$ to distinguish the vector of possible state values from sample realizations from the chain. Iterating on the Markov chain will produce a sample path $\{x_t\}_{t=0}^{T}$ where for each t, $x_t \in \mathbf{x}$. In the exposition below, I suppose that $n = 2$ for simplicity so that the transition matrix can be written

$$P = \begin{pmatrix} p_1 & 1-p_1 \\ p_2 & 1-p_2 \end{pmatrix}$$

Step 1. Set values for each of $(\mathbf{x}, P, \pi_0)$.

Step 2. Determine the initial state, x_0. To do this, draw a random variable from a **uniform distribution** on [0,1]. Call that realization ε_0. In Matlab, this can be done with the `rand()` command. If the number $\varepsilon_0 \leq \pi_{0,1}$, set $x_0 = \mathbf{x}(1)$. Otherwise, if $\varepsilon_0 > \pi_{0,1}$ set $x_0 = \mathbf{x}(2)$.

Step 3. Draw a vector of length T of independent random variables from a uniform distribution on [0,1]. Call a typical realization ε_t. Again, in Matlab this can be done with the command `rand(T,1)`. Now the current state is $x_t = \mathbf{x}(i)$, check if $\varepsilon_t \leq p_{i,1}$. If so, the state transits to $x_{t+1} = \mathbf{x}(j)$ with $j \neq i$. Otherwise, if $\varepsilon_t > p_{i,1}$, the state remains at i and $x_{t+1} = \mathbf{x}(i)$. Iterating in this manner builds up an entire simulation.

The attached Matlab file `markov_example.m` is a function file that implements this procedure for an arbitrary chain $(\mathbf{x}, P, \pi_0)$ and specified simulation length T.

C. Example

Let the growth rate of log GDP be

$$x_{t+1} \equiv \log(y_{t+1}) - \log(y_t)$$

and suppose that $\{X_t\}$ follows a 3 state Markov chain with

$$x = (\mu - \sigma, \mu, \mu + \sigma) = (-0.02, 0.02, 0.04)$$

and transition probabilities

$$P = \begin{pmatrix} 0.5 & 0.5 & 0 \\ 0.03 & 0.90 & 0.07 \\ 0 & 0.20 & 0.80 \end{pmatrix}$$

The idea here is that the economy can either be shrinking, $x = \mu - \sigma < 0$, growing at its usual pace, $x = \mu > 0$, or growing even faster. If the economy is in recession, there is about a 50/50 chance of reverting back to the average growth rate. If the economy is growing at its average pace, there is a slight probability of it falling into recession and a slightly bigger probability of it growing even faster, etc. We simulate a Markov chain on x and then recover the level of output via the sum

$$\log(y_t) = \log(y_0) + \sum_{k=1}^{t} x_k, \qquad t \geq 1$$

Attached is Matlab code which produces a simulation of this stochastic growth process assuming $\pi_0 = (0, 1, 0)$, simulation length $T = 100$ and initial condition $\log(y_0) = 0$.

Chris Edmond

22 August 2004