The Dynkin Martingale for General Markov Processes

by Lones Smith

Given is a Markov process $\{x_n\}$, where $x_n \mapsto x_{n+1}$.

Step 1: Constructing *Dynkin's martingale* for any markov process.

Let u be an arbitrary real function defined on the same range as the Markov process. Define $Lu(x) = E[u(x_n)|x_n] - u(x_n)$. Thus, the L functional operates on a function to tell us how much it is expected to rise under the operation of the Markov process. For general Markov processes (see Dynkin's texts, vol. I-II), the following is a martingale:

$$y_n = u(x_n) - \sum_{k=0}^{n-1} Lu(x_k)$$

In other words, the value of the chosen function in period n less the sum of its (progressively) expected increments is a martingale. To see this fact:

$$\begin{aligned} E[y_{n+1}|y_n] &= E\left[u(x_{n+1}) - \sum_{k=0}^{n} Lu(x_k) \middle| u(x_n) - \sum_{k=0}^{n-1} Lu(x_k)\right] \\ &= y_n + E[u(x_{n+1}) - u(x_n) - Lu(x_n)|y_n] \\ &= y_n + \Big(E[u(x_{n+1})|y_n] - u(x_n) - Lu(x_n)\Big) \end{aligned}$$

which is y_n, by definition of L.

Step 2: Application of Dynkin's Martingale to Hitting Times.

Iterating the martingale identity for any time fixed time horizon T yields $y_0 = E[y_T|y_0]$. Hence,

$$E[u(x_T) - \sum_{k=0}^{T-1} Lu(x_k)|x_0] = u(x_0)$$

In fact, we can assume by a standard trick that T is in fact a random (not necessarily uniformly bounded) stopping time – for instance, the first time a process hits a boundary. [Otherwise, simply put $\tau = min\langle t, T\rangle$, and then let $t \to \infty$, and apply Lebesgue's Dominated Convergence Theorem.] The name of the game is to choose special functions u satisfying inequalities that are either always true or at least true in our range (i.e. until x_n hits a barrier, say, at the stopping time T).

Assume that we have the inequality $Lu \leq c$, some constant $c < 0$. Then

$$\begin{aligned} E[u(x_T)|x_0] - cE[T|x_0] &\leq u(x_0) \\ \Rightarrow E[T|x_0] \leq \{u(x_0) - E[u(x_T)|x_0]\}/(-c) &< \infty \end{aligned}$$

if $E[u(x_T)|x_0]$ is uniformly bounded as t explodes.

In other words, if we can find a "test function" $u(x)$ such that $E[u_{n+1}|u_n] \leq u_n - c$, then we can conclude a finite hitting time. This should be hardly surprising, at an intuitive level. Simply transform the original Markov process $\{x_n\}$ into one $\{u_n\}$ with boundedly negative drift.

Now let's apply it in a nontrivial example. Assume our process on the reals is

$$x_{n+1} = \begin{cases} x_n + U & \text{chance } p + b/x_n \\ x_n - D & \text{chance } 1 - p - b/x_n \end{cases}$$

where $pU = (1-p)D$, b is a real constant, and $2b(U+D)+pU^2+(1-p)D^2 < 0$. In other words, for large x, the process has a zero drift, but the vanishing inverse terms in x somehow negatively contribute to the drift.

Consider an arbitrary test function $v(x)$. Then

$$Lv(x) = (p + b/x)v(x + U) + (1 - p - b/x)v(x - D) - v(x)$$

Then one can check that for the identity function $v_1(x) = x$, we have $Lv_1(x) = (b/x)[U + D]$. Hmmm ... not much use. Next, try $v_2(x) = x^2$. One then gets $Lv_2(x) = 2b(U + D) + (b/x)(U^2 - D^2) + pU^2 + (1 - p)D^2$. So letting $v = v_2 + (D - U)v_1$ yields $Lv(x) = 2b(U + D) + pU^2 + (1 - p)D^2 \equiv c$.

Regardless of whether we get entirely rid of the b/x term, our transformed process has a negative drift for large x iff

$$c = 2b(U + D) + pU^2 + (1 - p)D^2 < 0$$

With a constant negative drift, we would expect a finite hitting time for any lower boundary.