Application: Extended Markov Chains

Monday, September 22, 2008
11:49 AM

Let's continue our discussion of the quality control example, turning to the question of the fraction of shipped products which are defective in the long run. Natural to try and attack this with the law of large numbers just as we did for the question of the fraction of products which are inspected. But there'll be a twist…

$$\beta = \frac{\cancel{\#\text{ defective shipped products}}}{\#\text{ shipped products}}$$

D_n = # defective shipped products from nth inspection until just before $n+1$ inspection

S_n = # shipped products from nth inspection until just before $n+1$ inspection

$$\beta = \lim_{n\to\infty} \left(\frac{\sum_{n'=1}^{n} D_{n'}}{\sum_{n'=1}^{n} S_{n'}} \right)$$

$$= \lim_{n\to\infty} \left(\frac{\frac{1}{n}\sum_{n'=1}^{n} D_{n'}}{\frac{1}{n}\sum_{n'=1}^{n} S_{n'}} \right)$$

$$S_n = g(X_n) = \begin{cases} 0 & \text{for } X_n = 0 \\ 1 & 1 \leq X_n \leq M-1 \\ r & \text{for } X_n = M \end{cases}$$

So the law of large numbers for Markov chains can be applied (as we've already checked irreducibility and aperiodicity) to show that:

$$\lim_{n\to\infty} \frac{1}{n} \sum_{n'=1}^{n} S_{n'} = \sum_{j=0}^{M} \pi_j \, g(j)$$

$$= \pi_0 (0) + \sum^{M-1} 1 \cdot \pi_j + r\pi_M = (1 - \pi_0 - \pi_M) + r\pi_M$$

$$= \pi_0 (0) + \sum_{j=1}^{M-1} 1 \cdot \pi_j + r \pi_M = (1 - \pi_0 - \pi_M) + r \pi_M$$

$$= 1 - p + (r-1)(1-p)^M$$

Now let's turn to the long-time average in the numerator. Let's try to express D_n as a function of the state of the Markov chain:

$$D_n = h(X_n) = \begin{cases} 0 & \text{for} \quad 0 \leq X_n \leq M-1 \\ \text{random!} & X_n = M \end{cases}$$

$\downarrow$

binomial distribution for $r-1$ Bernoulli trials with "success" probability p

Strictly speaking, this means we can't use the law of large numbers for Markov chains because that corresponded to taking time averages of deterministic functions of the state of the Markov chain. Intuitively, one can guess correctly how to fix the law of large numbers to handle this by just replacing the random variable by its average (conditioned on the state of the Markov chain).

But we will take a more rigorous look at this just for the purpose of illustrating the concept of extending Markov chains. So we will extend our Markov chain by adjoining a random variable keeping track of the number of defective products between inspection n and $n+1$ and then this will make the above calculation into a time average of a deterministic function of the state of the extended Markov chain and so Law of Large Numbers can be applied to this extended Markov chain in the straightforward way.

Y_n = # defective products between inspection n and $n+1$

We will now take $(X_n, Y_n) = \vec{U}_n$

as our (extended) Markov chain. But wait, is it a Markov chain? It's fairly easy to write down the stochastic update rule.

$$X \; = \min(X \; + 1, M) \; Z$$

where $Z_n = \begin{cases} 1 & \text{w/prob } 1-p \\ 0 & \text{w/prob } p \end{cases}$

$$Y_n = \begin{cases} 0 & \text{for } 0 \le X_n \le M-1 \\ \xi_n & \text{for } X_n = M \end{cases}$$

$\xi_n \sim B(r-1,p)$: binomial distributed variable: $r-1$ trials, p = success prob

$$m \sim B(n,p): \quad P(m=j) = \underset{\frac{n!}{j!(n-j)!}}{\binom{n}{j}} p^j (1-p)^{n-j} \quad \text{for } j = 0,\dots,n$$

$$(X_{n+1}, Y_{n+1}) = w(X_n, Y_n, \underset{\uparrow}{Z_n}, \underset{\uparrow}{\xi_{n+1}})$$

indep r.v.'s

So we have established a stochastic update rule with w implicitly defined by the above relations.

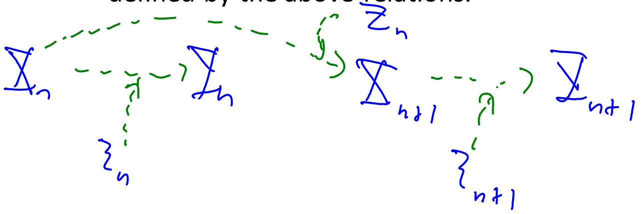

Now the number of defective shipped products from the nth inspection up to just before n+1 inspection is a deterministic function of the Markov chain:

$$D_n = h(X_n, Y_n) = \begin{cases} 0 & \text{if } 0 \le X_n \le M-1 \\ Y_n & \text{if } X_n = M \end{cases}$$

So this motivates to invoke the law of large numbers for Markov chains…but we first need check irreducibility and aperiodicity so that the LLN applies. We would run into trouble if we took the state space of the extended Markov chain as $\{0,1,\dots,M\} \times \{0,\dots,r-1\}$

the LLN applies. We would run into trouble if we took the state space of the extended Markov chain as $\{0,1,\ldots,M\} \times \{0,\ldots,r-1\}$

because there is no possible way to visit states (x,y) with $x < M,\ y > 0$

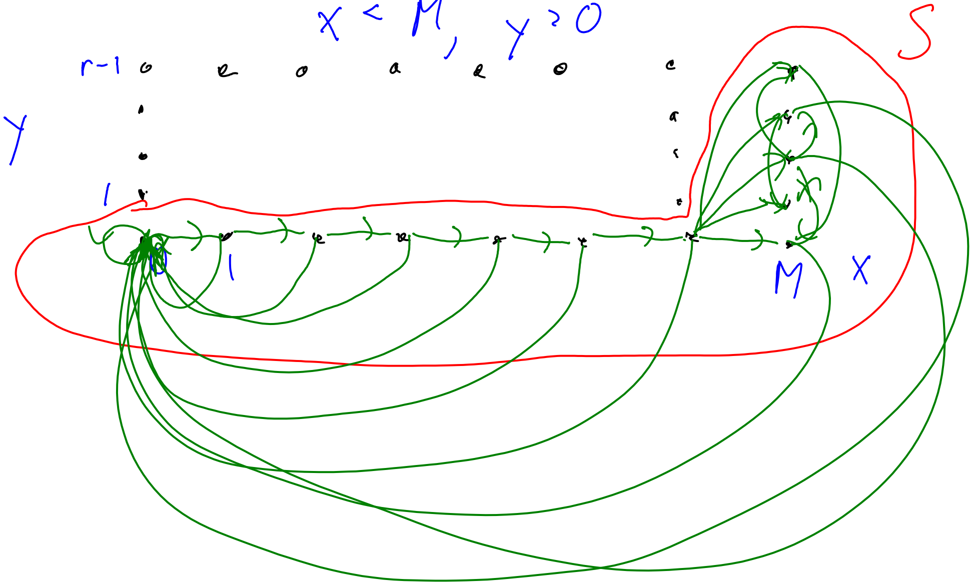

Define the state space S to only be the accessible subset of this product space.

$$S = \{(x,y) : (x \in \{0,1,2,\ldots,M-1\} \text{ and } y = 0) \text{ or } (x = M \text{ and } y \in \{0,1,2,\ldots,r-1\})\}$$

With this definition, Markov chain is:

- irreducible: Every state can reach (0,0), and from (0,0) you can reach any state (x,0) in x steps, and from (M,0) you can reach any state (M,y) in one step.
- aperiodic: (0,0) has period 1 since it can visit itself in one epoch.

Now we can use the Law of Large Numbers for this extended Markov chain:

$$\lim_{n\to\infty} \frac{1}{n} \sum_{n'=1}^{n} D_n = \lim_{n\to\infty} \frac{1}{n} \sum_{n'=1}^{n} h(X_n, Y_n)$$

$$= \sum_{(x,y)\in S} h(x,y)\, \tilde{\pi}_{xy}$$

How to compute the stationary distribution for the extended Markov chain? Could start from scratch, solve the equations directly like we did for the original Markov chain. Here the projection $(X_n, Y_n) \to X_n$ yields our original Markov chain, for from scratch, solve the equations directly like we did for the original Markov chain.

Here the projection $(X_n, Y_n) \to X_n$ yields our original Markov chain, for which we computed a stationary distribution π_x

$$\pi_x = \sum_{y:(x,y)\in S} \hat{\pi}_{x,y}$$

Reason: $\{X_n = x\} = \bigcup_{y:(x,y)\in S} \{X_n = x, Y_n = y\}$

$$\hat{\pi}_{x,0} = p(1-p)^x \quad \text{for } 0 \le x \le M-1$$

$$\sum_{y=0}^{r-1} \hat{\pi}_{M,y} = (1-p)^M$$

Notice from the definition of the Markov chain that for any (including the stationary) probability distribution:

$$P(Y_n = y \mid X_n = M) = P(\xi_n = y \mid X_n = M)$$
$$= \binom{r-1}{y} p^y (1-p)^{r-1-y}$$

$$P(X_n = M, Y_n = y) = P(Y_n = y \mid X_n = M)\, P(X_n = M)$$

(law of cond. prob.)

$$= \binom{r-1}{y} p^y (1-p)^{r-1-y}\, \underset{\substack{\| \\ \pi_M = (1-p)^M}}{P(X_n = M)}$$

in stat dist.

$$\hat{\pi}_{M,y} = \binom{r-1}{y} p^y (1-p)^{r-1-y} (1-p)^M \quad \text{for } y = 0,1,2,\dots,r-1$$

So returning to our law of large numbers formula:

$$\lim_{n\to\infty} \frac{1}{n} \sum_{n'=1}^{n} D_n = \sum_{(x,y)\in S} h(x,y)\, \tilde{\pi}_{x,y}$$
$$= \sum_{x=0}^{M-1} 0\, \hat{\pi}_{0,y} + \sum_{y=0}^{r-1} y\, \tilde{\pi}_{M,y}$$
$$= \sum_{y=0}^{r-1} y \binom{r-1}{y} p^y (1-p)^{r-1-y} (1-p)^M$$

$$= \sum_{y=0}^{r-1} y \binom{r-1}{y} p^y (1-p)^{r-1-y} (1-p)^M$$

$$= (1-p)^M \left(\underbrace{\sum_{y=0}^{r-1} y \binom{r-1}{y} p^y (1-p)^{r-1-y}}_{\langle \xi \rangle \, = \, (r-1)p} \right)$$

$$= (1-p)^M (r-1) p$$

By the way, repeating this argument formally would show in general:

$$\lim_{n\to\infty} \frac{1}{n} \sum_{n'=1}^{n} Y_{n'} = \sum_{j \in S} \pi_j \, \mathbb{E}_{\tilde{\pi}} [Y_n | X_n = j]$$

↑ conditional expectation w.r.t. stationary distribution in extended state space

provided (X_n, Y_n) can be formulated as an irreducible, aperiodic Markov chain.

So putting together the above calculations we have the fraction of defective products which are shipped (in the long run) is:

$$\beta = \frac{\lim_{n\to\infty} \frac{1}{n} \sum_{n'=1}^{n} D_{n'}}{\lim_{n\to\infty} \frac{1}{n} \sum_{n'=1}^{n} S_{n'}} = \frac{(1-p)^M (r-1) p}{(1-p) + (r-1)(1-p)^M}$$

$$= \frac{(1-p)^{M-1} (r-1) p}{1 + (r-1)(1-p)^{M-1}}$$

Remarks:
How is this fraction of shipped products which are defective related to the effort of inspection (fraction of products inspected)?

$$I = \frac{1}{1 + (r-1)(1-p)^M}$$

Note: $(r-1)(1-p)^M = \dfrac{1-I}{I}$

So $\beta = \dfrac{p\left(\frac{1-I}{I}\right)}{1-p+\frac{1-I}{I}} = \dfrac{p(1-I)}{(1-p)I + 1 - I}$

$$\beta = \frac{p(1-I)}{\quad}$$

$$1 - pI$$

Interesting point about this formula: There is no dependence in this relationship on the governing parameters M,r. The only thing that matters here is how much effort you're willing to devote to inspection (and that completely determines the quality of the shipped products) or vice versa. What's even more disconcerting is that you can show that the relationship between the quality of shipped products and the inspection effort is exactly the same as if products were sampled completely at random with the same inspection effort. This is just an artefact of our assumption that the defects occur independently in each product -- clearly then nothing clever helps. If the defects occur with some type of pattern, then sampling strategies will have effectiveness that depends on their design (which should reflect statistical properties of the defects.