Example of Full Conditionals for Gibbs Sampler

This example is used in HW Problem Set 17 due Wednesday, November 18, 2009. It is taken from Example 3, p. 172, of "*Explaining the Gibbs Sampler*", by G. Casella and E. George, **American Statistician** (1992) 46, 167-174.

Define the joint mixed-type density (continuous in Y and discrete in (X, N)) of random variables (X, Y, N) by:

$$f(x, y, n) \;=\; C \binom{n}{x} y^{x+\alpha-1} (1-y)^{n-x+\beta-1} e^{-\lambda} \frac{\lambda^n}{n!} \tag{1}$$

where the constant $C = \Gamma(\alpha+\beta)/(\Gamma(\alpha)\,\Gamma(\beta))$ is defined by the relation

$$\sum_{n=0}^{\infty} \sum_{x=0}^{n} \int_0^1 f(x, y, n)\, dy \;=\; 1$$

While artificially constructed, this density could arise in the following way: $N = n$ is a *Poisson*(λ) distributed sample size, $Y \sim$ *Beta*(α, β) plays the role of a probability parameter with beta prior distribution, and given $(N, Y) = (n, y)$, the data $X = x$ represents the number of successes in a Bernoulli coin-toss experiment with n trials and success-probability y.

As indicated in the Casella and George paper (p.168), the conditional distribution of X given N obtained by integrating out y can be found explicitly, for $x = 0, 1, \ldots, n$, as:

$$f_{X|N}(x|n) \;=\; \int_0^1 f(x, y|n)\, dy \;=\; \binom{n}{x} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\,\Gamma(\beta)} \cdot \frac{\Gamma(x+\alpha)\,\Gamma(n-x+\beta)}{\Gamma(\alpha+\beta+n)}$$

which is called the ‘beta-binomial distribution.

Our exercise in HW 17 is to simulate an *iid* sequence $X_1, X_2, \ldots$ from the marginal distribution of X, which can be written explicitly but not in closed form:

$$P(X = x) \;=\; f_X(x) \;=\; \sum_{n=x}^{\infty} \binom{n}{x} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\,\Gamma(\beta)} \cdot \frac{\Gamma(x+\alpha)\,\Gamma(n-x+\beta)}{\Gamma(\alpha+\beta+n)} e^{-\lambda} \frac{\lambda^n}{n!}$$

Obviously (?!), for each choice of parameters (λ, α, β), one could simulate directly from this marginal probability mass function. How would you do it ? To simulate R *iid* X_r variates, you could first choose a number M so large that

$$P(\max_{1\le r\le R} X_r \ge M) \;\le\; .001$$

based on some analytic estimate, and then calculate the probability masses $q_x = f_X(x),\ x = 0,\,1,\,\ldots,\,M$ and put them in a vector along with $q_{M+1} \equiv 1 - \sum_{j=0}^{M} q_j$. Finally you could use the `R` function `sample` to generate *iid* random integers from the set $\{0,\,1,\,\ldots,\,M+1\}$ with probabilities from the probability vector $\mathbf{q}$. But all this is a little cumbersome, especially if you wanted a really large sample-size R and if you wanted the parameters $(\lambda,\,\alpha,\,\beta)$ also to be generated from some distribution in terms of other hyper-prior paramaters.

So it would be attractive to have some second method of generating variates X_r. Gibbs Sampling is a handy approach here, because, although the marginal of X is cumbersome, the conditional distributions of X given Y, N, of Y given (X, N), and of N given (X, Y) are all quite explicit and closed-form:

$$f_{X|Y,N}(x|y,n) \;=\; \texttt{dbinom}(x,n,y) \qquad\qquad \text{(by definition)}$$

$$f_{Y|X,N}(y|x,n) \;=\; \texttt{dbeta}(y,\; x+\alpha,\; n-x+\beta) \qquad\qquad \text{(by inspection)}$$

and for $n \ge x$,

$$f_{N|X,Y}(n|x,y) \;=\; \frac{f(x,y,n)}{\sum_{m\ge x} f(x,y,m)} \;=\; (1-y)^{n-x}\,\frac{\lambda^n}{(n-x)!}\Big/\sum_{m=x}^{\infty}(1-y)^{m-x}\,\frac{\lambda^m}{(m-x)!}$$

The result of this last calculation is:

$$f_{N|X,Y}(n|x,y) \;=\; I_{[n\ge x]}\,e^{-(1-y)\lambda}\;\cdot\;\frac{((1-y)\lambda)^{n-x}}{(n-x)!}$$

or equivalently,

$$f_{N-X\,|X,Y}(k|x,y) \;=\; \texttt{dpois}(k,(1-y)\lambda)$$

All of these conditional densities are really easy to simulate from in `R` !