BAYESIAN PRACTICE

1.

(1) Your clothes are in the dryer. You are trying to find a pair of socks — any pair. The number N of pairs of socks in the dryer is unknown to you — it could be arbitrarily many. Each pair is distinct. You start drawing socks from the dryer, one at a time, looking for a pair and putting the unpaired socks in the laundry basket. The nth sock you draw matches one of the previous draws, so you stop.
 (a) Show that a flat prior over $N = 1, \dots, \infty$ leads to an improper (non-integrable) posterior.
 (b) Suppose your prior was uniform over $N = 1, \dots, 100$. Plot the posterior pdf and find the posterior median for the cases of $n = 3$ and $n = 10$.
 (c) Describe a computer algorithm to generate a sample from the posterior distribution of N given an observation $n = 10$, using the flat prior on N between 1 and 100. [Since this is not an in-class exam, you might actually run the algorithm, which would be more fun. There is no unique best way to do this. A version of MCMC could work, but a more direct approach, producing an i.i.d. sample, is also possible.]

(2) The notation in this example is meant to suggest analogies with the estimation of an autoregressive parameter. Suppose that $\rho \mid \hat{\rho}$ is distributed uniformly on an interval of length $1 - \hat{\rho}$ centered at $\hat{\rho}$. The marginal distribution of $\hat{\rho}$ is uniform on $(0, 1)$.
 (a) Find the conditional distribution of $\hat{\rho} \mid \rho$ for $\rho \geq .5$. (Covering the cases where $\rho < .5$ as well is quite feasible, but involves some messier math.)
 (b) Show that if we observe $\hat{\rho}$ and treat $\hat{\rho}$ as an estimator of an unknown value for ρ, it is biased downwards (in the frequentist sense) when $\rho > .5$.
 (c) Plot the pdf of $\hat{\rho} \mid \rho$ for $\rho = .5$ and $\rho = .99$.
 (d) Is the ratio of the bias to the standard deviation larger for values of ρ close to one? (In the case of an autoregressive parameter, this ratio does increase as ρ approaches 1, but the rate of increase goes to zero as ρ approaches 1.)

(3) If $\hat{\beta}$ is the Bayesian posterior mean for an unknown parameter β, constructed with a proper prior, it is impossible that $\hat{\beta}$ is biased in the same direction for all values of β. That is, if $E[\hat{\beta} \mid \beta] > \beta$ for some values of β, we must have $E[\hat{\beta} \mid \beta] < \beta$ for other values of β. Prove this. [It can be proved in two lines, if you think of the right argument.] How can this result be reconciled with the fact that the OLS estimator $\hat{\rho}$ of the autoregressive parameter in a univariate autoregression is biased downward for all ρ in $(0, 1)$ and yet is also the posterior mean for β under a flat prior?

Date: January 13, 2009.