Stat 544 – Spring 2005
Homework assignment 4

Due on Wednesday, March 9 in TA's office by 5:00 pm

Problem 1

We wish to obtain draws y from a truncated (below) normal distribution with mean μ and variance σ^2. Suppose that $\mu = 0$ and $\sigma^2 = 1$, so that

$$p(y|\mu, \sigma^2) \propto \exp\left(-\frac{y^2}{2}\right), \; y \geq a$$

The naive method to sample from this distribution, is simply to draw values from the N(0,1) and discard all draws that are less than a. This method may be effective if a is small relative to μ, but can get costly otherwise.

a - In the example above, what is the average number of simulations from N(0,1) required for each acceptance?

b - A potential instrumental or proposal distribution in this case is the translated exponential distribution, with pdf

$$g_\alpha(z) = \alpha \exp[-\alpha(z - a)] \; z \geq a$$

for some value of α.

b.1 - Derive the upper bound on the ratio $p/g_\alpha(z)$. Hint: There are two upper bounds on the ratio, depending on whether $\alpha > a$ or not.

b.2 - Find the value of α that minimizes the upper bound in each case.

b.3 - Given your results in part (b.2), propose an algorithm to sample values from a truncated normal distribution using the rejection sampling method. List all of the steps that are needed to carry out your algorithm.

Problem 2

Consider a simple linear regression model $y_i = \alpha + \beta x_i + e_i$, where $e_i \sim \text{N}(0, \sigma^2)$, and the x_i are known and fixed. Assume that the errors are independent.

Let the prior distributions for α and β be: $\text{N}(\alpha_0, \sigma_a^2)$ and $\text{N}(\beta_0, \sigma_b^2)$, respectively, with parameters fixed and known. For σ^2, let the prior distribution be an inverted scaled χ^2 distribution with known parameters.

We are interested in estimating the posterior distributions of α, β, and σ^2.

a - Write down the full probability model for this problem: likelihood, priors, and joint posterior.

b - Write down the full conditional distributions for the parameters in the model, to implement a Gibbs sampler for this problem.

c - In words, explain how you would go about implementing the Gibbs sampler

d - Write down model portion of the WinBUGS code for this problem. Do no attempt to actually run the program, since you do not have any data.