MA676: Bayesian Methods – Exercise Sheet 3

1. Code (write computer programme) the Metropolis algorithm for obtaining samples from (i) $N(5, 1.5^2)$, (ii) Gamma($\alpha = 0.5, \beta = 1$). Study the sensitivity of the algorithm with respect to the chosen proposal scaling. Generalise your programme which is then able to generate from any univariate distribution with two parameters. In this case the Metropolis routine should take a function name (which evaluates the density function) as one of its argument.

2. Suppose that $x_1, \ldots, x_n$ are i.i.d. observations from a Bernoulli distribution with mean θ. A logistic normal prior distribution is proposed for θ (a normal distribution for $\log \frac{\theta}{1-\theta}$). Show that if the prior mean and variance for $\log \frac{\theta}{1-\theta}$ are 0 and 1 respectively then the prior density function for θ is

$$\pi(\theta) \;=\; \frac{1}{\sqrt{2\pi}\theta(1-\theta)} \exp\left(-\tfrac{1}{2}\left(\log \tfrac{\theta}{1-\theta}\right)^2\right)$$

As this prior distribution is not conjugate, the Bayes estimator $E(\theta|x_1, \ldots, x_n)$ is not directly available. It is proposed to estimate it using a Monte Carlo sample generated by the Metropolis-Hastings method. One possible algorithm involves generating proposals from the prior distribution, independently of the current observation. Suppose $n = 10$, $\sum_{i=1}^{10} x_i = 8$. Write a Splus programme and run it to obtain the posterior mean.

3. Assume that $X_1, X_2, \ldots, X_n$ are independent identically distributed $N(\theta, 1)$ observations. Suppose that the prior distribution for θ is Cauchy with density

$$\pi(\theta) \;= \frac{1}{\pi}\frac{1}{1+\theta^2} \qquad -\infty < \theta < \infty.$$

Derive, upto a constant of proportionality, the posterior density of θ. Suppose that the importance sampling distribution is the prior distributions given above. Obtain the acceptance probability for the rejection method and the Metropolis-Hastings independence sampler. Suppose that $n = 10$ and $\bar{x} = 1.5$. Code the two methods and find the Bayes estimate for θ under squared error loss.

4. Assume that $X_1, X_2, \ldots, X_n$ are independent identically distributed $N(\theta, \sigma^2)$ observations. Suppose that the joint prior distribution for θ and σ^2 is

$$\pi(\theta, \sigma^2) \;= \frac{1}{\sigma^2}.$$

   (a) Derive, upto a constant of proportionality, the joint posterior density of θ and σ^2.

   (b) Derive the conditional posterior distributions of θ given σ^2 and σ^2 given θ.

   (c) Derive the marginal posterior density of θ.

   (d) Write a Splus programme for Gibbs sampling from the joint posterior distribution of θ and σ^2. Hence obtain the estimates of $E(\theta|x_1, \ldots, x_n)$ and $\mathrm{Var}(\theta|x_1, \ldots, x_n)$. For your own data set verify that the estimates are close to the true values.