Part A Simulation and Statistical Programming Hilary 2014
Problem Sheet 3, due 9am Monday week 7 at SPR1.

1. (a) Give a Metropolis-Hastings algorithm to sample according to the Gamma probability density function,
$$\pi(x) \propto x^{\alpha-1} \exp(-\beta x), \quad x > 0$$
with parameters $\alpha, \beta > 0$. Use the proposal distribution $Y \sim \text{Exp}(\beta)$.

   (b) Write an R function implementing your MCMC algorithm. Your function should take as input values for α and β and a number n of steps and return as output a realization $X_1, X_2, ..., X_n$ of a Markov chain targeting $\pi(x)$. State briefly how you checked your code.

2. MCMC for Bayesian inference (first two parts were an exam Q in 2009)

   (a) Give a Metropolis-Hastings Markov chain Monte Carlo algorithm simulating $X \sim \text{Binomial}(n, r)$ for n trials and success probability r.

   (b) Suppose the success probability for X is random, with $\Pr(R = r) = p(r)$ given by
$$p(r) = \begin{cases} r & \text{for } r \in \{1/2, 1/4, 1/8, ...\}, \text{ and} \\ 0 & \text{for } r \text{ otherwise.} \end{cases}$$
An observed value $X = x$ of the Binomial variable in part (a) is generated by simulating $R \sim p$ to get $R = r^*$ say, and then $X \sim \text{Binomial}(n, r^*)$ as before. Specify a Metropolis-Hastings Markov chain Monte Carlo algorithm simulating a Markov chain, $\{R_t\}_{t=0,1,2,...}$ with equilibrium probability mass function $R_t \xrightarrow{D} p(r|x)$ where
$$p(r|x) \propto \pi(x|r)p(r)$$
is called the posterior distribution for r given data x.

   (c) Write an R function implementing your MH MCMC algorithm with target distribution $p(r|x)$. Suppose $n = 10$ and we observe $x = 0$. Run your MCMC and estimate the mode of $p(r|x)$ over values of r.

3. (The random-scan Gibbs sampler) Suppose $p(x)$ is the pmf of some multivariate random variable $X \in \mathbb{Z}^m$, so $x = (x_1, x_2, ..., x_m)$. Let
$$x_{-i} = (x_1, x_2, ..., x_{i-1}, x_{i+1}, ..., x_m)$$
be the vector of components with x_i omitted. Consider a Metropolis-Hastings algorithm simulating a Markov chain $\{X^{(t)}\}_{t=0}^{\infty}$, with $X^{(t)} \to p$. Here is an update scheme: suppose $X^{(t)} = x$ and we propose a candidate y by choosing i at random from $1, 2, ..., m$, simulating y_i from the conditional distribution $X_i|X_{-i} = x_{-i}$ and setting
$$y = (x_1, x_2, ..., x_{i-1}, y_i, x_{i+1}, ..., x_m).$$

   (a) Show that if y and x differ at a single index i then the pmf for$Y = y|X = x$ is
$$q(y|x) = (1/m)p(y_i|x_{-i}).$$
Give a formula for $p(y_i|x_{-i})$ in terms of $p(y)$.

   (b) Show that the acceptance probability $\alpha(y|x)$ is equal to one.

   (c) Give a Gibbs sampler for the multinomial probability mass function,
$$\pi(y, z) = \frac{n!}{y!z!(n-y-z)!} p^y q^z (1-p-q)^{n-y-z}, \quad 0 \le y + z \le n$$
with parameters n, p, q, where $p, q \ge 0$, $p + q \le 1$ and $n \in \{1, 2, 3, \ldots\}$. [Hint: show that $Z|Y = y \sim \text{Binomial}(n - y, q/(1 - p))$]