Part A Simulation and Statistical Programming HT19

Problem Sheet 2 – due Week 4 Thursday 10am

1. Suppose X is a discrete random variable taking values $X \in \{1, 2, ..., m\}$ with probability mass function (pmf) $p(i) = \Pr(X = i)$. Let $q(i) = 1/m$ be the pmf of the uniform distribution on $\{1, 2, ..., m\}$. Give a rejection algorithm simulating $X \sim p$ using proposals Y distributed according to q. Calculate the expected number of simulations $Y \sim q$ per returned value of X if $p = (0.5, 0.25, 0.125, 0.125)$.

2. Let $Y \sim q$ with probability density function (pdf) $q(x) \propto \exp(-|x|)$ for $x \in \mathbb{R}$. Let $X \sim \mathrm{N}(0, 1)$ be a standard normal random variable, with pdf $p(x) \propto \exp(-x^2/2)$.

   (a) Find M to bound $p(x)/q(x)$ for all real x.

   (b) Give a rejection algorithm simulating X using q as the proposal pdf.

   (c) Can we simulate $Y \sim q$ by rejection using p as the proposal pdf?

3. Consider a discrete random variable $X \in \{1, 2, \ldots\}$ with probability mass function

$$p(x; s) = \frac{1}{\zeta(s)} \frac{1}{x^s}, \qquad \text{for } x = 1, 2, 3, ...$$

   where $s > 1$.

   (a) The normalising constant $\zeta(s)$ is hard to calculate. However, when $s = 2$ we do have $\zeta(2) = \pi^2/6$. Give an algorithm to simulate $Y \sim p(y; 2)$ by inversion.

   (b) Implement your inversion algorithm as an R function. Your function should take as input an integer $n > 0$ and return as output n iid realisations of $Y \sim p(y; 2)$. Say briefly how you checked your code.

   (c) Give a rejection algorithm simulating X with pmf $p(x; s)$ for $s > 2$, using the rejection algorithm and draws from $Y \sim q$ where the proposal is $q(y) = p(y; 2)$. You will need to derive the upper bound $M' \geq \tilde{p}(x; s)/\tilde{q}(x)$ for all x.

   (d) Compute the expected number of simulations of $Y \sim q$ for each simulated X in the previous part question, giving your answer in terms of $\zeta(s)$.

   (e) Implement your algorithm as an R function. Your function should take as input s and return as output $X \sim p(x; s)$ and the number of trials N it took to simulate X.

4. Suppose $X \sim N(0, \sigma^2)$ is a Gaussian random variable with mean 0 and variance σ^2. We want to estimate $\mu_\phi = \mathbb{E}(\phi(X))$ for some function ϕ :

$\mathbb{R} \to \mathbb{R}$ such that $\phi(X)$ has finite mean and variance. Suppose we have iid samples $Y_1, ..., Y_n$ with $Y_i \sim N(0,1), i = 1, 2, ..., n$. We consider the following two estimators for μ_ϕ:

$$\widehat{\theta}_{1,n} = \frac{1}{n} \sum_{i=1}^{n} \phi(\sigma Y_i)$$

and

$$\widehat{\theta}_{2,n} = \frac{1}{n\sigma} \sum_{i=1}^{n} \exp\left[-Y_i^2 \left(\frac{1}{2\sigma^2} - \frac{1}{2}\right)\right] \phi(Y_i).$$

(a) Show that $\widehat{\theta}_{1,n}$ and $\widehat{\theta}_{2,n}$ are unbiased and give the expression of their variances.

(b) What range of values must σ be in for $\widehat{\theta}_{2,n}$ to have finite variance? Can you give a weaker condition if it is known that $\int_{-\infty}^{\infty} \phi^2(x)\, dx < \infty$?

(c) Why might we prefer $\widehat{\theta}_{2,n}$ to $\widehat{\theta}_{1,n}$, for some values of σ^2 and functions ϕ? (Hint: consider estimating $\mathbb{P}(X > 1)$ with $\sigma \ll 1$).