Monte Carlo Methods - Spring 16 - Homework 4

1. (based on problems in Owen) For the nominal density we let $p(x)$ be the standard normal density. We want to use importance sampling to compute the expected value of $f(x)$ with respect to $p(x)$ where $f(x)$ is defined below. We take the sampling density $q(x)$ to be normal with variance 1 and mean θ. So

$$q(x,\theta) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2}(x-\theta)^2)$$

(a) Let $f(x) = exp(kx)$. Find the optimal θ where optimal means that it minimizes the variance of the importance sampling algorithm using $q(x, \theta)$ Can you do any better with a sampling density that is not of the assumed form?

(b) Now let $f(x) = \exp(-(x-10)^2)$. Find the optimal θ for our one parameter family of sampling densities. Now consider all possible sampling densities. What is the optimal one? How does the variance of the importance sampling algorithm for this optimal q compare with the variance of the importance sampling algorithm using $q(x, \theta)$ with the optimal choice of θ?

2. Problem 9.10 in chapter 9 of Owen's notes.

3. In class we looked at the following example. Z has a standard normal distribution. We want to compute $P(Z > 4)$. We let

$$p(x,\theta) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2}(x-\theta)^2)$$

so the nominal density $p(x)$ is $p(x, 0)$.

(a) Suppose we take $\theta = 4$ for our sampling density. Using importance sampling, compute (approximately) the probability $P(Z > 4)$. Find the variance of the estimator and use it to give a 95% confidence interval for your answer.

(b) Now suppose we want to reduce the variance even further. In class I discussed doing a pilot MC run with some reference value θ_r for the sampling θ and using this single run to estimate the variance for other values of θ. Using a reference $\theta_r = 4$, estimate the variance for other θ and find the optimal choice of θ.

4. Consider the random walk exit example from section 6.4 in the notes. In class (and the notes) I took the sampling density to be a random walk with steps having a normal density with variance 1 and mean θ with $\theta = 0$. Now consider sampling densities where the mean of the step distribution θ is a parameter. Find the optimal θ. It is possible to do this by brute force: for many different values of θ you can do a MC run to estimate the variance. But ideally you should be able to estimate the optimal θ with a single MC run.