STA 2201S 2001 Assignment 4

Quiz on Thursday Febuary 8th. Bring printouts to the quiz.

1. As we saw in class, a natural way to construct a Monte Carlo estimate of a probability p is to simulate a collection of *i.i.d.* random variables $Y_1, \ldots, Y_m$ with $Pr\{Y_i = 1\} = p$ and $Pr\{Y_i = 0\} = 1 - p$ for $i = 1, \ldots, m$. The Monte Carlo approxiation of p is then $\overline{Y}_m$. Any time we do this, we should accompany our estimate with a confidence interval, to indicate the proabble accuracy of our approximation. Our 99% confidence interval will be $\overline{Y}_m$ plus or minus a "99% margin of error." Using this vocabulary,

   (a) Find the Monte Carlo sample size for this *i.i.d.* Bernoulli case so that the 99% margin of error is *guaranteed* to be less than 0.001. I get a number in excess of 1.6 million. Ouch.

   (b) What Monte Carlo sample size is required if the 99% margin of error is to be at most 0.01?

   (c) Suppose we are interested in estimating a power that is to be 0.80 or more. Now what Monte Carlo sample size is required if the 99% margin of error is to be at most 0.01?

   Bring your printout.

2. Sometimes, Monte Carlo estimates of p are based on simulating *i.i.d.* $Y_1, \ldots, Y_m$ that are not Bernoulli, but still satisfy $E[Y_i] = p$. As before, we use $\overline{Y}_m$ as our Monte Carlo approximation, and it converges by the Strong Law of Large Numbers. The quantities $Y_1, \ldots, Y_m$ are usually conditional probabilities, so that $0 \leq Y_i \leq 1$. Show that the required Monte Carlo sample sizes you calculated in Question One also apply to this case, as follows.

   Let Y be a random variable that is arbitrary except that $Pr\{0 \leq Y \leq 1\} = 1$ and $E[Y] = \theta$. Prove $Var(Y) \leq \theta(1 - \theta)$.

3. Let the random variable X have density

$$f(x) = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(x-u)^2} \frac{1}{\pi(1 + u^2)} \, du.$$

   Approximate $Pr\{X < 0\}$ by Monte Carlo. Accompany your estimate by a large-sample confidence interval of width no more than 0.02. Bring your printout.

   Some comments: This density is a location mixture of normals, in which the mixing distribution is standard Cauchy. It's symmetric about zero, so the answer is exactly one-half. I did the computer part two different ways, the *i.i.d.* Bernoulli way, and another way that turned out to be more efficient. In each case it took only a few lines of code.