Homework 5
36-325/725
due Friday Oct 5

(1) Let $X_1, \ldots, X_n \sim \text{Uniform}(0,1)$ and let $Y_n = \max\{X_1, \ldots, X_n\}$. Find $E(Y_n)$. Hint 1: first find the cdf of Y_n. Hint 2: $Y_n \leq y$ if and only if $X_i \leq y$ for $i = 1, \ldots, n$.

(2) A particle starts at the origin of the real line and moves along the line in jumps of one unit. For each jump the probability is p that the particle will jump one unit to the left and the probability is $1-p$ that the particle will jump one unit to the right. Let X_n be the position of the particle after n units. Find $E(X_n)$ and $Var(X_n)$. (This is known as a random walk.)

(3) A fair coin is tossed until a head is obtained. What is the expected number of tosses that will be required?

(4) Let X be a discrete random variable. Let $Y = r(X)$. By definition, $E(Y) = \sum_y y f_Y(y)$. In class, we learned the lazy statistician rule: $E(Y) = \sum_x r(x) f_X(x)$. Prove that $\sum_y y f_Y(y) = \sum_x r(x) f_X(x)$.

(5) Let X be a continuous random variable with cdf F and suppose that $P(X > 0) = 1$ and that $E(X)$ exists. Show that $E(X) = \int_0^\infty P(X > x) dx$.

Hint: Consider integrating by parts. The following fact is helpful: if $E(X)$ exists then

$\lim_{x \to \infty} x[1 - F(x)] = 0$.

(6) Let $X_1, X_2, \ldots, X_n$ be $N(0,1)$ random variables and let $\overline{X}_n = n^{-1} \sum_{i=1}^n X_i$. Plot $\overline{X}_n$ versus n for $n = 1, \ldots, 10,000$. Repeat for $X_1, X_2, \ldots, X_n \sim \text{Cauchy}$. (Recall that the Cauchy has density $f(x) = 1/(\pi(1 + x^2))$.) Explain why there is such a difference.

(7) Let $X \sim N(0,1)$ and let $Y = e^X$. Find $E(Y)$ and $Var(Y)$.

(8) A simulated stock market. Let $Y_1, Y_2, \ldots$ be independent random variables such that $P(Y_i = 1) = P(Y_i = -1) = 1/2$. Let $X_n = \sum_{i=1}^n Y_i$. Think of $Y_i = 1$ as "the stock price increased by one dollar", $Y_i = -1$ as "the stock price decreased by one dollar" and X_n as the value of the stock on day n.

(8a) Find $E(X_n)$ and $Var(X_n)$. (Hint: see question 2).

(8b) Simulate X_n and plot X_n versus n for $n = 1, 2, \ldots, 10,000$. An easy way to do this is:

```
n <- 10000
y <- rbinom(n,1,.5)
y <- 2*y-1   ### do you see what this does?
x <- cumsum(y)/(1:n)
help(cumsum)
```

Repeat the whole simulation several times. Notice two things. First, it's easy to "see" patterns in the sequence even though it is random. Second, you will find that the four runs look very different even though they were generated the same way. How do the calculations in (8a) explain the second observation? The moral of the story: when you lose lots of money in the stock market one day, remember that I warned you.