The purpose of these problems is to use simulation to produce estimates of certain quantities, also guaranteeing in a particular sense the accuracy of those estimates. You will also use recursive methods to update statistics, to maximize the run-time efficiency of the code.

1. Write a simulation program whose goal is to estimate the expected value of the maximum of 10 i.i.d. continuous uniform (0,10) random variables. Construct a graph that shows how the variance of your estimate decreases as the sample size increases. Use recursive methods to update your values of the mean and variance. Include your code with your answer, as well as the graph that was produced in your code.

Extra credit: What is the exact CDF of the maximum of 10 i.i.d. continuous uniform (0,10) random variables? What are its mean and variance?

2. Write a simulation that estimates the value of π. Do this as follows: generate independent *x*- and *y*-coordinates that are continuous uniform (0,1) distributed; thus, the ordered pairs (x, y) fall uniformly on a square. Then ask, for each of those points, if it falls within the circle centered at (0.5, 0.5) and with radius 0.5. This should happen with probability $\pi/4$ (think about the area of the circle vs. the area of the square). You can use the fraction of times that the points fall within the circle, over a very large sample, to construct an estimate of $\pi/4$. Produce an estimate of π whose standard deviation is less than 10^{-4}. Your answer should include your estimate, and the number of iterations necessary.

Extra credit: As your simulation for problem 2 is proceeding, produce an estimate for the amount of time remaining (at each increment of 100,000 iterations). Keep these estimates in memory, and compare against the actual amounts of time that it took from each of these points to complete the simulation. When were the estimates most accurate? Use a scatterplot to show this, and include a 45° line in the plot to show what a perfect predictor would have produced.

3. (Adapted from Ross, "Simulation", 4th edition). Suppose you generate a sequence of independent continuous uniform (0,1) random variables, and stop the first time one of them is less than its predecessor. Clearly, it takes at least two numbers for this to happen, and there is no guaranteed maximum. What is the expected number of random variables necessary for this to happen?
    a. If M is the random variable indicating the number of random uniforms that were necessary until the first one that was less than its predecessor, make a cogent argument for why M does not have the geometric distribution.
    b. Show that $\mathrm{P}\{M > n\} = \frac{1}{n!}$ for integer $n \geq 0$.

c. For a discrete random variable that is strictly non-negative, another way to compute its mean is: $\mathrm{E}[M] = \sum_{n=0}^{\infty} \mathrm{P}\{M > n\}$. Use this fact to show that $\mathrm{E}[M] = e$.
d. Simulate this process, and use that simulation to estimate e. Continue until the standard deviation of your estimate is less than 10^{-4}. Use recursive methods to help with the termination criterion. Include your code, your estimate of e, and the number of runs it took to satisfy the criterion.