Assigned: Wednesday, September 15
Due: Wednesday, September 22

Reading Assignment: This problem set covers the material in Chapter 4, Sections 4.1 − 4.6 (pp. 122 − 149), and 4.8 (pp. 158 − 166) of Ross, and Chapter 8, Sections 8.1 − 8.3 (pp. 261-269) of Yates and Goodman, which you can find on reserve at Grainger. You are encouraged to attempt the self-test problems $\#8, 10, 15, 16, 17$ and 18 on pp. 185 − 186 of Ross; their solutions can be found in Appendix B.

The problems you will turn in, and which will be graded are the following:

16. Problem 47, p. 176, Ross

17. Problem 67, p. 178, Ross

18. Let $\mathbf{X}$ denote a binomial random variable with parameters (N, p). What is the probability that $\mathbf{X}$ is an odd integer? Remember that 0 is an even integer.
Hint: What is $(x+y)^N - (x-y)^N$?

19. Seven people hold reservations for travel in a 4-passenger shuttle from Champaign to O'Hare. The number of persons who actually show up to travel can be modeled as a binomial random variable $\mathbf{X}$ with parameters $(7, \frac{1}{2})$. If more than 4 people show up, only the first 4 get to go, and the rest are left behind. What is the average number of passengers who are left behind?

20. You are trying to persuade an empty-headed king that you can see the future. You offer to forecast what happens on repeated independent tosses of a biased coin of the realm which you happen to know has $P(\text{heads}) = 0.11$.

    (a) The skeptical king asks you to predict the number of heads that will occur on the next 1000 tosses and promises to execute you if your guess is wrong, just to make it more interesting. Which number should you predict and why? What is the probability that the 1000 coin tosses *do* result in the number of heads you predicted?

    (b) You luck out and guess right in part (a). The next day, the king asks you to predict how many tosses will be required to observe the next Head. Which number should you predict and why? What is the probability that a Head *does* occur for the first time on the toss you predicted?

    (c) Since you guessed right twice in a row, the king is thinking that you can indeed see into the future, and assigns a harder problem: predict the number of tosses required to observe a Head for the 105th time. Which number should you predict and why? What is the probability that a Head *does* occur for the 105th time on the toss you predicted?

21. Courtiers jealous of your growing fame substitute a coin bearing an image of the king's father. Fortunately, this is observed by your trusty sidekick who tells you that the coin to be used tomorrow is different. Naturally you are reluctant to make further predictions about the coin. To forestall further requests for amazing demonstrations of your powers, you tell the king that you have the powers to estimate probabilities from experimental data, and the king, who flunked out ECE 413, is duly impressed.

He tells you that he is going to toss the coin 1000 times and that you are to estimate $P(\text{Heads}) = p$.

(a) Head occurred for the first time on the 12th toss. You consider announcing the value p right away (without waiting for the 1000 tosses to be completed). What is the maximum-likelihood estimate of the value of p? That is, what value of p maximizes the probability of a Head occurring for the first time on the 12th toss?

(b) You decide may be it is best to wait for the results of some more tosses before deciding on your estimate of p. The 300th head occurred on the 994th toss. What value of p maximizes the probability of a Head occurring for the 300th time on the 994th toss?

(c) You sensibly decide to wait out the last 6 tosses also, and all of these results in Tails. What is your estimate of the value of p after 1000 tosses?

(d) You already knew after 994 tosses that 300 heads occurred. The last 6 tosses did not result in Heads and thus conveyed no information about p. So why isn't the maximum-likelihood estimate of p in part (c) the same as the estimate in part (b)?

22. When we perform an experiment, event A occurs with probability $P(A) = 0.01$. In this problem, we estimate $P(A)$ using R_n, the relative frequency of A over n independent trials. How many trials n are needed so that the probability R_n differs from $P(A)$ by more than 0.001 is less than 0.01?