Assigned: Thursday, June 26
Due: Thursday, July 3
Reading: Ross Chapters 3.1-3.3, 4.1-4.8
Noncredit Exercises: (Do not turn these in) Ross Chapter 4: Problems 2,3,7,8,12,20,35,40,41,43,44,49,70; Theoretical Exrcises: 9,17,18,25,26.

1. Let **X** denote a binomial random variable with parameters (N, p). What is the probability that **X** is an odd integer? Remember that 0 is an even integer.
[Hint: What is $(x+y)^N - (x-y)^N$?]

2. Seven people hold reservations for travel in a 4-passenger limousine from Champaign to St. Louis. The number of persons who actually show up to travel can be modeled as a binomial random variable **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?

3. A long message is divided into L packets of N bits each (including headers, address, CRC bits, tail, flags, etc. etc. etc.), and transmitted over a channel with bit error probability p. If the CRC detects that a packet is received in error, the packet transmission is repeated. *But,* if a packet has been transmitted a *total* of five times, and has not been received correctly even on the fifth try, it is not transmitted again, and is deemed to be lost.

(a) What is the probability that a packet is received correctly?

(b) Let $\mathbf{X}_i$ denote the number of times that the i-th packet is transmitted over the channel. What are the possible values of $\mathbf{X}_i$? What is the pmf of $\mathbf{X}_i$? (Be careful about $P\{\mathbf{X}_i = 5\}$!!).

(c) What is the average number of times that the i-th packet is transmitted? i.e., what is $E[\mathbf{X}_i]$?

(d) What is the probability that all L packets are received successfully = P{entire message is received successfully}?

4. Let **X** denote a Poisson random variable with unknown parameter λ. Suppose that the event $\{\mathbf{X} = k\}$ occurs.

(a) What is the maximum–likelihood estimate of λ? That is, what value of λ maximizes the probability of the observed event $\{\mathbf{X} = k\}$?

(b) Consider a binomial random variable **Y** with parameters (N, p) where the parameter p is unknown. If the event $\{\mathbf{Y} = k\}$ is observed (e.g. heads occurs k times on N tosses of a biased coin with P(Heads) = p), then we showed in class that $\hat{p} = k/N$ is the maximum–likelihood estimate of p. Since for large N and small p, the binomial random variable **Y** can be approximated by a Poisson random variable **X** with parameter λ = Np, it would seem reasonable that the maximum-likelihood estimate of λ would be $\hat{\lambda} = N\hat{p} = k$. Does your answer to part (a) give this result?

5. [The Once and Future King] You are trying to persuade an empty-headed monarch that you can foresee 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(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?

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 power to estimate probabilities from experimental data, and the king, who flunked out of ECE 313, is duly impressed. He tells you that he is going to toss the coin 1000 times and that you are to estimate P(Heads) = p.

(d) Heads occurred for the first time on the 12th toss. You consider announcing the value of 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?

(e) You decide that maybe 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?

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

(g) 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 (f) the same as the estimate in part (e)?

6. This problem on conditional probability has three ***unrelated*** parts.

(a) If $P(A|B) = 0.3$, $P(A^c|B^c) = 0.4$, and $P(B) = 0.7$, find $P(A|B^c)$, $P(A)$, and $P(B|A)$.

(b) If $P(E) = 1/4$, $P(F|E) = 1/2$, and $P(E|F) = 1/3$, find $P(F)$.

(c) If $P(G) = P(H) = 2/3$, show that $P(G|H) \geq 1/2$.

7. Let **X** denote a *negative* binomial (or Pascal) random variable with parameters (r, p). Then, **X** counts the number of trials required to observe r successes where the probability of success on any trial is p. Given that $\mathbf{X} = n$, what is the conditional probability that the i-th trial resulted in a success? To avoid trivialities, assume that $n > r$ and also that $n > i$.

8. Suppose that 105 passengers hold reservations for a 100–seat flight from Chicago to Champaign. The number of passengers showing up for the flight can be modeled as a binomial random variable **X** with parameters (105,0.9).

(a) Find the probability that all passengers who show up get seats, i.e. find $P\{\mathbf{X} \leq 100\}$.

(b) Explain why the number of no-shows can be modeled as a Poisson random variable **Y**, and compute the value of the parameter λ.

(c) Compute the probability that all passengers who show up get seats based on this Poisson model, i.e find $P\{\mathbf{Y} \geq 5\}$, and compare to the "more exact" answer of part (a).

Now suppose that 15 passengers are arriving in Chicago on a connecting flight that is on time with probability 1/3 and late with probability 2/3. If the connecting flight is on time, all 15 passengers show up for the flight to Champaign (no one stops off in a bar for a drink!); else they all are not there. The other 90 passengers each decide independently with probability 0.9, as before, to show up for the flight.

(d) Given that the connecting flight is on time, what is the (conditional) probability that all passengers who show up get seats? Repeat for the case when the connecting flight is late.

9. Monty Hall, the host of the TV game show "Let's Make A Deal™", shows you three curtains. One curtain conceals a valuable prize, while the other two conceal junk. All three curtains are equally likely to conceal the prize. He offers you the following "deal": pick a curtain, and you can have whatever is behind it. When you pick a curtain, instead of giving you your just deserts, Monty (who knows where the prize is) opens one of the remaining curtains to show you that there is junk behind it, and offers the following "new, improved deal™" : you can either stick with your original choice, or switch to the remaining (unopened) curtain. Amidst the deafening roars of "Stand pat" and "Switch, you idiot" from the crowd, Monty points out that previously your chances of winning were 1/3. Now, since you know that the prize is behind one of the two unopened curtains, your chances of winning have increased to 1/2, and thus the new improved deal is indeed better. Use the theorem of total probability to determine

(a) the probability of winning if you always switch.

(b) the probability of winning if you would rather fight than switch.

(c) whether Monty is correct in asserting that if you choose randomly betweeen the two unopened curtains, you have a probability of winning of 1/2.
Note: Everybody knows that the rules of the game are that Monty always opens one of the two unchosen curtains and he always offers the "new improved deal," i.e. he never opens a curtain to reveal the prize (saying "Oops, you lose; go back to your seat")

10. At the County Fair, you see a man sitting at a table and rapidly rolling a pea between three walnut shells. "Step right up, me bucko, and try your luck! The hand is quicker than the eye!" he says, and hides the pea under one of the shells. You have no idea which shell is covering the pea, but you point to one shell at random and bet that the pea is under it. The man picks up one of the shells that you didn't choose, and shows you that the pea is not underneath that shell. He asks if you would like to switch your bet to the other unchosen shell. Should you accept the offer? Why or why not? How does this game differ from the one analyzed in Problem 9?