Assigned: Wednesday, February 13
Due: Wednesday, February 20
Reading: Ross, Chapters 3 and 4
Noncredit Exercises:
Chapter 4 Problems: 49, 51, 53, 57-59, 70; Theoretical Exercises: 16-18, 25, 26.
Chapter 3 Problems: 1, 2, 5, 10, 12, 16, 31, 38, 39, 44; Exercises at the end of this problem set.
Problems:

1. The Sirrah Poll wishes to assess what percentage of voters believe Governor Shrub's claim that he was shocked, shocked to hear that the word *Rats* was used subliminally in one of his TV commercials. A random sample of N voters is asked for opinions. Assume that each voter decides the matter independently, that is, the number of voters believing the Governor can be modeled as a binomial random variable $\mathbf{X}$ with parameters (N,p) where p is the probability that a randomly chosen voter believes the Governor. The Poll knows the value of $\mathbf{X}$ that it obtained, and it wishes to *estimate* the value of p and report this.

(a) What is the maximum-likelihood estimate of p? Call this $\hat{p}$. Note that $\hat{p}$ is a *function* of $\mathbf{X}$.

(b) The Poll wants to be *fairly sure* that its estimate $\hat{p}$ has a margin of error of at most 2%. The Poll thus wants to have the following inequality hold:

$$P\{|\hat{p} - p| \geq 0.02\} \leq 0.05.$$

The Chicago Tribune then reports on its front page that the Sirrah Poll had found that $100\hat{p}$% of the people believed Governor Shrub, and that the margin of error of the poll is ±2%. Use the Chebyshev inequality to determine the value of N that will guarantee a confidence interval of width 0.04 with a confidence level of 95%.

2. ["Let your communication be Yea Yea, Nay Nay, for whatsoever is more than this cometh of evil"]
The bits transmitted over a communication link are in error with probability p independently of each other. Since p is too large, a simple *repetition code* is used in which each data bit is sent n times (i.e. sent repeatedly) over the link. Assume that n is odd. The receiver examines the n received bits and decides that the transmitter was trying to send a data bit 0 or 1 according as the majority of the n received bits received is 0 or 1.

(a) Let $\mathbf{X}$ denote the number of errors in the n received bits corresponding to a data bit. What kind of random variable is $\mathbf{X}$? What is the mean of $\mathbf{X}$? What is its variance?

(b) The receiver decision is in error if and only if the event $\{\mathbf{X} > n/2\}$ occurs. Express the probability that the receiver decision is in error as a function of n and p, and compute its numerical value for the case $n = 3$, $p = 10^{-2}$. Compared to just transmitting the data bit once and receiving it with probability of error p, does the repetition code reduce the probability that the receiver decision is in error? Since there a'int no such thing as a free lunch, what is being given up in thus reducing the bit error probability?

(c) You are asked to choose a value of n for the repetition code such that the probability that the receiver decision is in error is at most 10^{-5}. *Estimate* $P\{\mathbf{X} > n/2\}$ using the Chebyshev inequality and use this to find the smallest value of n that meets the desired specification.

3. ["I'm leaving on a jet plane, Don't know when I'll be back again…"] 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 $\mathbf{X}$ with parameters (105,0.9), that is, each is deciding independently of the others (with probability 0.9) to show up. Note that this simple model does not take into account things such as a family either all showing up or all failing to show.

(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 $\mathbf{Y}$, 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)

(d) Now suppose that 15 passengers are arriving 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 remaining 90 passengers decide independently as before. 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.

4. ["I am from Iowa; I only work in outer space…"] Each box of Cornies, the breakfast of silver medalists, contains one picture, which is of Luke Skywalker with probability 2/3 and of Darth Vader with probability 1/3, independently of which picture is in any other box of Cornies. Little Jimmy Kirk of Cedar Rapids, Iowa, asks his mother to buy boxes of Cornies until he has at least one picture of both beings, and his mother agrees to do so.

(a) What is the minimum number of boxes of Cornies that Mrs Kirk must buy?

(b) Let $\mathbf{X}$ denote the number of boxes of Cornies Mrs Kirk purchases until such time as Jimmy has acquired at least one picture of each of the two entities. What is the pmf of $\mathbf{X}$? Verify that the total probability mass specified by your pmf does equal 1.

(c) What is the conditional pmf of $\mathbf{X}$ given that the first box contained a picture of Luke? What is the conditional pmf of $\mathbf{X}$ given that the first box contained a picture of Darth?

(d) Use the theorem of total probability to compute the unconditional pmf of $\mathbf{X}$ from the conditional pmfs found in part (c). Do you get the same answer as in part (b)? Why not?

(e) Find the mean and variance of $\mathbf{X}$. [Hint: it might be easier to determine the conditional means and variances and then combine them to obtain the unconditional mean and variance]

5. 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 $P(G|H) \geq 1/2$.

6. An urn contains 6 red balls and 4 green balls. A ball is drawn at random from the urn, and then another ball is drawn at random without the first one being replaced in the urn.

(a) Let R_1 denote the event that the first ball is red and R_2 the event that the second ball is red. What is $P(R_1)$? What is $P(R_2|R_1)$? What is $P(R_2|R_1^c)$?

(b) Use the numbers found in part (a) to compute $P(R_2)$. Does the answer surprise you?

(c) Repeat parts (a) and (b) for an experiment in which after the first ball is drawn, *it is put back into the urn along with 3 more balls of the same color before the second ball is drawn.*

Noncredit Exercises:

1. ["One man's meat is another man's Poisson"] Let $\mathbf{X}$ 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 $\mathbf{Y}$ 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 $\mathbf{Y}$ can be approximated by a Poisson random variable $\mathbf{X}$ with parameter $\lambda = Np$, 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?

2. Let **X** denote a Poisson random variable with parameter λ.

(a) Show that P{**X** is even} = exp(–λ)cosh(λ). Don't forget that 0 is an even integer!

(b) In Problem 2 of Problem Set #4, you proved (I hope!) that the probability that a binomial random variable with parameters (n,p) has even value is $[1+(1-2p)^n]/2$. For large n and small p, show that $[1+(1-2p)^n]/2 \approx \exp(-np)\cosh(np)$, which is consistent with part (a).

3. 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; return to your seat")

4. 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 3?