Expectation, covariance, binomial, Poisson

18.600 Problem Set 4, due March 17

Welcome to your fourth 18.600 problem set! The interesting topics we have discussed in lecture include the linearity of expectation, the bilinearity of covariance, and the notion of *utility* as used in economics. (Under certain "rationality" assumptions everyone has a utility function whose expectation they seek to maximize.) We will see in this problem set how these ideas play a role in some important (though perhaps overly simplistic) theories from finance (MPT and CAPM). We'll also have a few problems about binomial and Poisson random variables (keep thinking about those!) and a chance to learn about Siegel's paradox.

Please stop by my weekly office hours (2-249, Wednesday 3 to 5) for discussion.

A. FROM TEXTBOOK CHAPTER FOUR:

1. Problem 23: You have $1000, and a certain commodity presently sells $2 per ounce. Suppose that after one week the commodity will sell for either $1 or $4 an ounce, with these two possibilities being equally likely.

   (a) If your objective is to maximize the expected amount of money that you possess at the end of the week, what strategy should you employ?

   (b) If your objective is to maximize the expected amount of the commodity that you possess at the end of the week, what strategy should you employ?

   Remark: Look up Siegel's paradox. It's pretty interesting.
   http://mindyourdecisions.com/blog/2012/03/15/siegels-paradox-about-exchange-rates/

2. Theoretical Exercise 13: Let X be a binomial random variable with parameters (n, p). What value of p maximizes $P\{X = k\}, k = 0, 1, 2, \ldots, n$? This is an example of a statistical method used to estimate p when a binomial (n, p) random variable is observed to equal k. If we assume that n is known, then we estimate p by choosing that value of p which maximizes $P\{X = k\}$. This is known as the *method of maximum likelihood estimation*.

3. Theoretical Exercise 19: Show that if X is a Poisson random variable with parameter λ, then
$$E[X^n] = \lambda E[(X + 1)^{n-1}].$$
Now use this result to compute $E[X^3]$.

B. Suppose that during each given minute there is a 10^{-6} probability that there is an accident at a particular intersection (independently of all other minutes). Using the approximation of

$500,000$ minutes per year, we expect to see .5 accidents per year on average. One year somebody proposes to install a new kind of stoplight to reduce accidents. You believe *a priori* that there is a $1/4$ chance that the new stoplight is *effective*, in which case it will reduce the accident rate by fifty percent, and a $3/4$ chance it will have no effect. The new stoplight is installed and during the next two years there are no accidents. Using Poisson approximations, compute your updated estimate of the probability that the light is effective.

Remark: It is often hard to tell whether preventative measures against rare events are having an effect. With twenty years of data we might be more confident, but by that point accident rates may have changed for other reasons (e.g., self driving cars). On the other hand the $k!$ in the Poisson denominator means that *large* numbers are *extremely* unlikely. If we suddenly see 10 accidents in one year, we should seriously question our assumption that the number is Poisson with $\lambda = 1/2$ or $\lambda = 1/4$.

C. Your friend plans to tie one end of a carbon nanotube rope to the moon and use the other to pull a system of power generators around the earth. Your friend somehow convinces you that with probability $p = 10^{-12}$ this scheme will solve the world's energy problems and make your friend af profit of $\$10^{12}$. Your friend needs a dollar for postage to submit a patent application and promises you the entire future profit in exchange for the dollar. How much net profit do you expect to make (ignoring interest, taxes, etc.)? What is the *variance* (in dollar-squared units) of your profit? How about the standard deviation (in dollar units)?

D. Larry the Very Subprime Lender gives loans of size $10,000. In 25 percent of cases, the borrower pays back the loan quickly with no interest or fees. In 50 percent of cases, the borrower disappears (moves away, declares bankruptcy, dies) without paying anything. In 25 percent of cases, the borrower pays back the loan slowly and — after years of ballooning interest payments, hefty fees, etc. — pays Larry a total of $100,000. However, in this scenario, Larry has to give $60,000 to third parties (repo services, foreclosure lawyers, eviction teams, bill collectors, etc.) in order to get the borrower to pay the $100,000. Compute the following:

(a) The expectation and variance of the *net* amount of profit Larry makes from each loan (after subtracting collection expenses and the initial $10,000 outlay).

(b) The expectation and variance of the *net* amount a given borrower ends up paying (i.e., amount paid minus amount borrowed).

Note: You might have some ethical concerns with Larry's business model. You can google payday loans (read the Wikipedia article, say) if you want a more realistic discussion of the legal and moral issues involved in lending to populations with high default risk.

E. Define the covariance $\text{Cov}(X, Y) = E[XY] - E[X]E[Y]$.

1. Check that $\text{Cov}(X, X) = \text{Var}(X)$, that $\text{Cov}(X, Y) = \text{Cov}(Y, X)$, and that $\text{Cov}(\cdot, \cdot)$ is a bilinear function of its arguments. That is, if one fixes one argument then it is a linear

function of the other. For example, if we fix the second argument then for real constants a and b we have $\text{Cov}(aX + bY, Z) = a\text{Cov}(X, Z) + b\text{Cov}(Y, Z)$.

2. If $\text{Cov}(X_i, X_j) = ij$, find $\text{Cov}(X_1 - X_2, X_3 - 2X_4)$.

3. If $\text{Cov}(X_i, X_j) = ij$, find $\text{Var}(X_1 + 2X_2 + 3X_3)$.

F. Suppose n people throw their hats in a pile, the hats are randomly shuffled and returned, one to each person. Let N be the number of people who get their own hat. Let M be the number of people who are part of a two person pair (a, b) where a gets b's hat b gets a's hat. In other words, M is the number of people who don't get their own hat but do get the hat *of the person who got their hat*. (You can observe that M is even with probability one.) Compute $E[M]$, $E[N]$, $E[MN]$ and $\text{Cov}[M, N]$. **Hint:** use indicator variables and linearity of expectation.

G. Instead of maximizing her expected wealth $E[W]$, Jill maximizes $E[U(W)]$ where $U(x) = -(x - x_0)^2$ and x_0 is a large positive number. That is, Jill has a *quadratic utility function*. (It may seem odd that Jill's utility declines with wealth once wealth exceeds x_0. Let us assume x_0 is large enough so that this is unlikely.) Jill currently has W_0 dollars. You propose to sample a random variable X (with mean μ and variance σ^2) and to give her X dollars (she will lose money if X is negative) so that her new wealth becomes $W = W_0 + X$.

1. Show that $E[U(W)]$ depends on μ and σ^2 (but not on any other information about the probability distribution of X) and compute $E[U(W)]$ as a function of x_0, W_0, μ, σ^2.

2. Show that given μ, Jill would prefer for σ^2 to be as small as possible. (One sometimes refers to σ as *risk* and says that Jill is *risk averse*.)

3. Suppose that $X = \sum_{i=1}^n a_i X_i$ where a_i are fixed constants and the X_i are random variables with $E[X_i] = \mu_i$ and $\text{Cov}[X_i, X_j] = \sigma_{ij}$. Show that in this case $E[U(W)]$ depends only on the μ_i and the σ_{ij} (but not on any other information about the joint probability distributions of the X_i) and compute $E[U(W)]$. Hint: first compute the mean and variance of X.

Remark: We conclude (assuming quadratic utility) that portfolio builders care *only* about expectations and covariances of items in their portfolio. This idea underlies the (1990 Nobel Prize Winning) *Modern Portfolio Theory* (MPT) and *Capital Asset Pricing Model* (CAPM). Before these theories, it was believed that when the *variance* of an asset return is high, the *expected* return should be higher as well (the *risk premium*) because otherwise people wouldn't buy risky assets. MPT and CAPM predict that one gets a risk premium for *systemic risk* (the part of the variance explained by correlation with the *market portfolio*, defined to be the sum total of all risky assets) but not for *idiosyncratic risk* (explosure to which can be reduced by

diversification). These theories also predict that everyone's optimal investment strategy is to put some (investor-dependent) fraction of their money in a risk free asset and the remainder in the market portfolio (which we think of as a giant index fund). You can google MPT and CAPM to read about how well or poorly these theories match reality.

Remark: People often say utility functions should be strictly concave (negative second derivative) to explain risk aversion... but is that necessarily true? Here is a naive story about charitable giving. Suppose your utility function is given by your own health/comfort plus a constant c times the sum of the health/comfort of all other humans on the planet. For example, if $c = .01$, then you are mostly selfish, but you would be willing to give up a comfort for yourself if it would enable more than 100 strangers to enjoy the same comfort. You'd give up your life if you could save more than 100 other lives. Utilitarians might theorize that it is a good to thing that $c > 0$ (so that we help others when we can make a big difference) but maybe also a good thing that $c < 1$ (since a little selfishness might be efficient in practice). As you acquire more money, there may come some point at which you believe that the marginal value of another dollar to you (in added health/comfort) is *less than* c times the amount a dollar donated to a global charity with relatively high expected impact (like those profiled at givewell.org) would increase health/comfort for others. After that point, in principle you should donate *all* of your additional money to charity. If this is indeed your plan, then your utility function might be very close to linear for a long time after that point, since the amount of good you do in a huge global effort is roughly linear in the amount you give.

Remark: Some economists say that in reality charitable giving should be modeled as a consumptive good (that happens to have positive externality — google *"warm glow giving"*) that has to compete with other consumptive goods among even the very wealthy. This point of view might predict actual behavior better than what I sketched above.

Remark: A "rational" person (in the economic sense) has a utility function and a subjective probability measure, and makes decisions in order to optimize expected utility. A group of people, each of whom is rational in this sense, may be decidely *irrational* as a group. Arrow's impossibility theorem (look it up) states that (under any reasonable voting scheme) a democratic group of people may prefer A to B and B to C and C to A. Political parties, companies, and entire countries can all be "irrational" to a greater extent than their individual members.