MIT OpenCourseWare

18.440: Probability and Random Variables
Problem Set 4

Travis McVoy

February 5, 2024

Creative Commons License Statement

Correctness Statement

While there is always a chance of errors in any body of the work, everything I achieve here will be self taught—which means the risk of errors is perhaps higher than usual. If, for some reason, one finds themself on my page tempted to use my work, please first check that it is indeed correct. If you find an error, email me and I'll fix it :)

Note: The majority of the below problems are from *A First Course in Probability* 8th ed. by Sheldon Ross.

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?

   Solution.

   a) I will assume our problem is discrete—that we are only allowed to buy an integer number of ounces of the commodity. Also, being that this chapter is on random variables, I assume we are supposed to use random variables when solving the problem. However, I think random variables hinder communication, so I'll do both and let you be the judge.

   Random variables solution:

   Let's first consider the expected dollar gain/loss on the second week. That is, if we buy an ounce on the first week and sell it the second week, what do we expect to happen? Let S be a random variable that denotes the selling price of the commodity on the second week (so S takes on two values: 1 and 4, each with probability $1/2$). Now, define a random variable R that denotes the transaction result. Observe that R is a function of S given by $S - 2$. It then follows that $E[R]$ is the expected money we posses for each transaction during the second week. We can now solve the problem. If we buy n ounces on the first week, we will have

   $$1000 + nE[R] = 1000 + n\left(\frac{1}{2}(4-2) + \frac{1}{2}(1-2)\right) = 1000 + \frac{n}{2}(2-1) = 1000 + \frac{n}{2}$$

   dollars on the second week when we sell the commodity. Clearly, the optimal strategy is to just buy everything on the first week, sell everything on the second.

   Solution without random variables:

   If an ounce sells for one dollar next week, we will lose a dollar for every ounce. On the other hand, if it sells for four, we will gain two dollars for every ounce. The expected value per ounce is then

   $$\frac{1}{2}(2-1) = 0.5 \text{ dollars.}$$

   Since the expected value per ounce is positive, I think we maximize the expected amount by going all in and buying 500 ounces.

   Remark. Be careful about the wording in this problem. I believe the problem is effectively asking, "What's the maximum value you could have," which is different from minimizing your loss (shoutout machine learning). If you want to minimize your loss, and you play this week 1, week 2 game only once, the clear strategy is to not buy anything because you don't risk losing anything. If, on the other hand, you have the opportunity to play this game over and over again, then, as we said before, you buy as much as you can during week 1 and sell on week 2.

You will lose here and there, but the more you play, the more you will profit in the long run (on average). We can see this buy running a simulation in Python (see `buySell.py` on code page). I graphed the results after each iteration for several samples, then created a dotplot of the value obtained in the end for 100 samples:

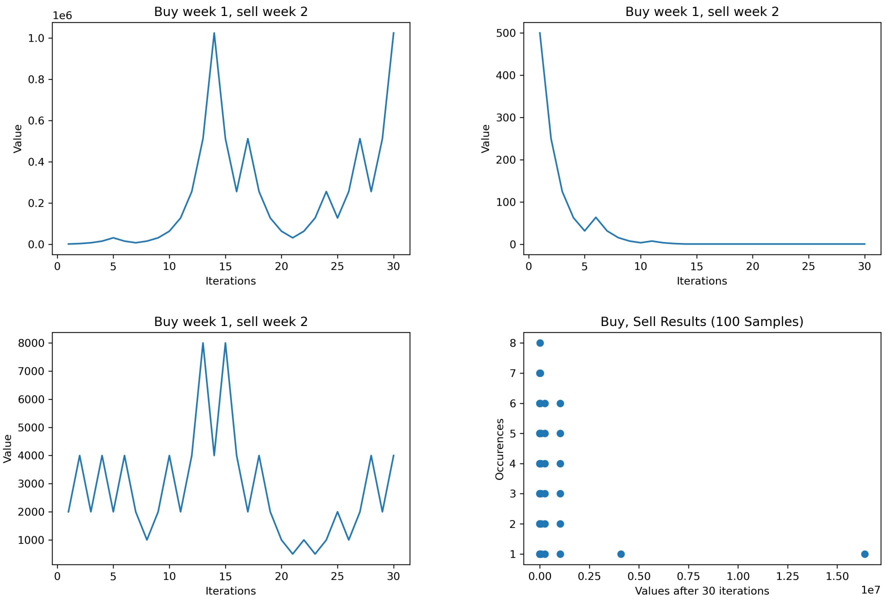

Notice the different scales of the y axes across each of the three samples and the x axis on the dotplot. Clearly, it isn't a given that we do really well. If we were to keep playing the game, we might make a bunch of money and then lose it all. Therefore, if we were to repeatedly play the game, we might consider stopping after we've made a sizeable profit instead of blindly iterating.

Remark. I need to learn how to remove outliers in histograms and dotplots in Python. It isn't clear if most samples end in a complete loss or in a value that is quite small relative to 4 million because the scale is all messed up.

b) I'm rather confused because it seems to me that we should get the same answer (buy everything on week 1). The expected cost per ounce in the second week is \$2.5 so it seems to me it would be better to just buy everything on week 1. It seems suspicious to me that we would get the same answer to two different questions, so let's try another angle.

Suppose we spend x dollars on the second week and $1000 - x$ on the first week. It then follows that the expected number of ounces owned on the second week is given by

$$\frac{1}{2}(1000 - x) + \frac{1}{2}x + \frac{1}{2} \cdot \frac{x}{4} = 500 - \frac{4x}{8} + \frac{5x}{8} = 500 + \frac{x}{8}.$$

The above is maximized when we make x as big as possible (1000) which means our second method suggests it's best to not buy anything on the first week. I guess this makes sense. If we buy everything on the first week, we will always have 500 ounces. If we wait to buy everything on the second week, sometimes we will have 250 ounces and sometimes we will have 1000 ounces. If we average the two results out, we will get 625 (which agrees with the above). I'm satisifed. Time to move on.

■

2. **Problem 35.** A box contains 5 red and 5 blue marbles. Two marbles are withdrawn randomly. If they are the same color, then you win \$1.10; if they are different colors, you lose a \$1.00. Calculate

   a) the expected value of the amount you win;
   b) the variance of the amount you win.

   Solution.

   a) Let R be a hypergeometric variable that represents the number of red marbles out of the 2 drawn marbles. That is, the probability that there are k red marbles out of the 2 chosen marbles is given by

   $$P(R = k) = \frac{\binom{5}{k}\binom{5}{2-k}}{\binom{10}{2}}.$$

   Next, let W be a random variable as a function of R that denotes how much money we win. There are two ways in which we can win—either two red marbles are drawn or no red marbles are drawn—and one way we can lose (one red, one blue drawn). Therefore, we have

   $$W = g(R) = \begin{cases} 1.1 & \text{if } R = 0 \text{ or} R = 2 \\ -1 & \text{if } R = 1 \\ 0 & \text{otherwise.} \end{cases}$$

   We now find the expected value of the amount we win with

   $$\begin{aligned} E[W] = E[g(R)] &= \sum_k g(k)p_R(k) \\ &= 1.1p_R(0) + 1.1p_R(2) - 1p_R(1) \\ &= \frac{22 \cdot 10}{10 \cdot 45} - \frac{25}{45} \\ &= -\frac{1}{15}. \end{aligned}$$

   b) There are several ways we could approach calculating the variance. I will use the fact that $\text{Var}(W) = E[W^2] - E[W]^2$. We already have $E[W]$ so we just need to find $E[W^2]$. Letting $h(W) = W^2$ we have

   $$\begin{aligned} E[W^2] &= \sum_k h(k)p_R(k) \\ &= 1.1^2\big(p_R(0) + p_R(2)\big) + (-1)^2 p_R(1) \\ &= \frac{11^2 \cdot 20}{10^2 \cdot 45} + \frac{25}{45} \\ &= \frac{242 + 250}{450} = \frac{492}{450} = \frac{82}{75}. \end{aligned}$$

To conclude, we find that

$$\begin{aligned}\mathrm{Var}(W) &= \frac{82}{75} - \left(-\frac{1}{15}\right)^2 \\ &= \frac{246-1}{225} = \frac{49}{45}.\end{aligned}$$

At this point, we have answered the problem. However, plots are fun so let's generate some plots.

Fig. 1a-c: Winnings vs. Time

I first simulated the marbles game and plotted are winnings vs. time. We can see three samples, each containing 300 iterations. Since the expected winnings are $-1/15$ per iteration, we would expect to have -20 dollars in the end (assume we start with no money...for some reason). Mostly, I found that we do indeed end around -20, but it's possible to end higher or lower:

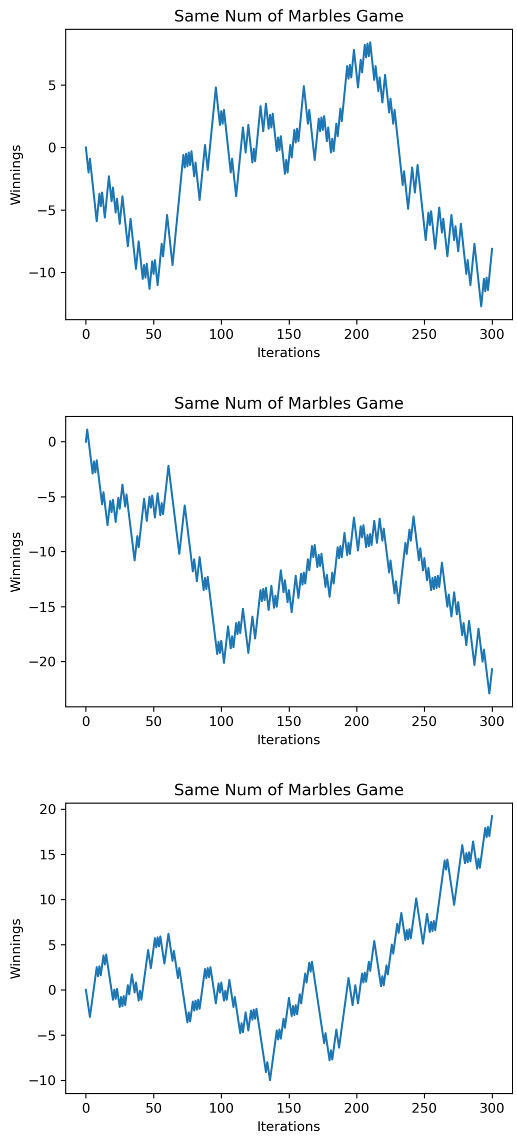

Fig. 2: Distribution of Means

Next, I tried to gain some understanding of variance, but I wasn't quite sure what to plot. I first tried playing the game 300 times, getting the average winnings for that sample, and then repeating for 500 samples, but I'm not sure the results I got make sense:

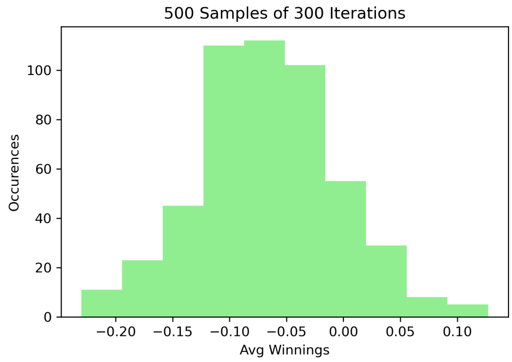

While we do see that most values are around $-0.067 \approx -1/15$, I don't think our plot matches what we found for the variance. The variance we found is slightly bigger than 1, so $\sigma \approx 1$ so shouldn't there be values that are -0.067 ± 1? Or is that not how it works? I guess not, because then the average would be bigger. So maybe our results make sense. This is where independent study fails. I'd really like to talk to a professor/TA and get a better understanding. Oh well.

Fig 3a-c: Distribution of Final Winnings

For my last plot, I again took 500 samples of 300 iterations, but this time, I took our final winnings from each sample. We see that -20 is most common, which is what we'd expect. What's interesting is that most of the values lie within -20 ± 20. My question is, how does this relate to the variance? Does it relate to the variance? For instance, if our variance happened to be roughly instead of roughly 1 and we still had $\mu = -20$, would we see most values falling between -20 ± 40? I don't think so, because that seems to apply some sort of linearity between variance and the mean. All the more reason to take a mathematical statistics course. Here are some plots:

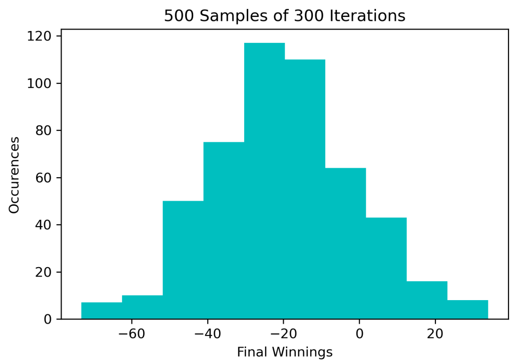

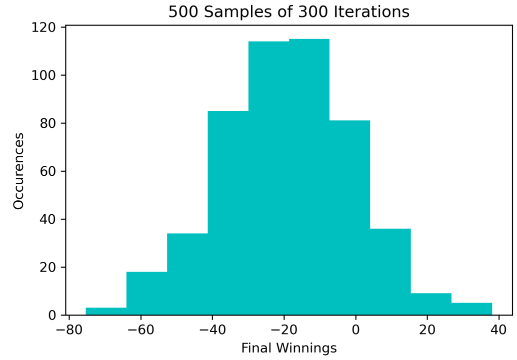

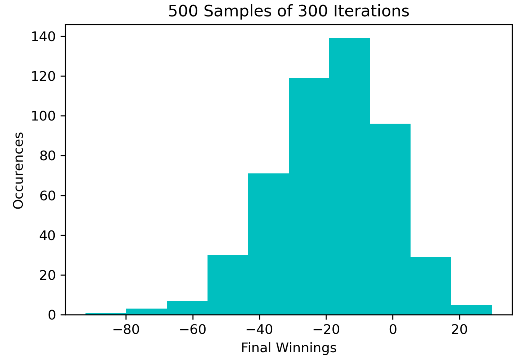

Edit. I talked to Greg (professor on campus) about my results. He plugged $(300 \cdot 1.089)^{1/2}$ into a calculator and got roughly 18, which matches our plots. I don't really understand why, but he also said I'm sort of jumping ahead by several chapters. I guess we'll see all the above in more detail when we talk about the central limit theorem.

3. **Problem 50.** Suppose that a biased coin that lands on heads with probability p is flipped 10 times. Given that a total of 6 heads results, find the conditional probability that the first 3 outcomes are

   a) h, t, t

   b) t, h, t

Solution.

a) Let X be binomial random variable denoting the number of heads in 10 flips. That is,

$$P(X = k) = \binom{10}{k} p^k (1-p)^{10-k}.$$

Now define the event E_1 to be the occurence of htt. We wish to find $P(E_1|X = 6)$. Using the definition of conditional probability, we have

$$\begin{aligned} P(E_1|X = 6) &= \frac{P(E_1 \cap X = 6)}{P(X = 6)} \\ &= \frac{P(E_1)P(X = 6|E_1)}{P(X = 6)}. \end{aligned}$$

We can find $P(X = 6|E_1)$ using a random variable Y that denotes the number of heads that occur in 7 coin flips. Since only 1 head occured in the first 3 flips, we need $Y = 5$ heads to occur in the remaining 7 in order for there to be 6 total. Thus, we have

$$\begin{aligned} P(E_1|X=6) &= fracP(E_1)P(X=6|E_1)P(X=6) \\ &= \frac{P(E_1)P(Y=5)}{P(X=6)} \\ &= \frac{p(1-p)^2\binom{7}{5}p^5(1-p)^2}{\binom{10}{6}p^6(1-p)^4} \\ &= \frac{\binom{7}{5}p^6(1-p)^4}{\binom{10}{6}p^6(1-p)^4} \\ &= \frac{\binom{7}{5}}{\binom{10}{6}} = \frac{21}{210} = \frac{1}{10}. \end{aligned}$$

b) Unless I'm missing something, this is basically the same problem. You could define an event E_2 given by tht, but you'd still need 5 out of remaining 7, so the answer is still 1/10.

4. **Problem 61.** The probability of being dealt a full house in a hand of poker is approximately 0.0014. Find an approximation for the probability that, in 1000 hands of poker, you will be dealt at least 2 full houses. (NOTE: the MIT pset claims this is problem 57. In my text (8th ed), it's problem 61. Go figure.)

   Solution. We need to be careful. The words "at least" complicate the problem. I think we can use a cumulative distribution function. In particular, we know that if X is a random variable and s, t are defined values of X, then

$$P(s \le X \le t) = P(X \le t) - P(X \le s-1).$$

   Clearly, we are to use a Poisson random variable. If it isn't clear, try doing it with a binomial. You'll need to calculate $(1 - 0.0014)^{999}$ and $(1 - 0.0014)^{1000}$, which are certainly not easy to find. Therefore, we let X be a Poisson random variable and $\lambda = 1000 \cdot 0.0014 = 1.4$ and we have

$$\begin{aligned} P(2 \le X \le 1000) &= 1 - P(X=1) - P(X=0) \\ &= 1 - \frac{1.4}{e^{1.4}} - \frac{1}{e^{1.4}} \\ &\approx 0.4082. \end{aligned}$$

5. **Theoretical Exericse 13.** Let X be a binomial random variable with parameters (n, p). What value of p maximizes $P(X = k)$ for $k = 0, 1, 2, \dots, n$? This is an example of a statisitcal method used to estimate p when a binomial (n, p) random variable is observed to equal k. If we assume that n is unknown, 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.*

Solution. There might be a clever way of doing this, but I'm just going to use calculus. We want to find the maximum of the function

$$f(p) = \binom{n}{k} p^k (1-p)^{n-k}.$$

We first take the derivative. For convience, I'll denote $\binom{n}{k}$ as the constant c. We now have

$$f'(p) = ckp^{k-1}(1-p)^{n-k} - (n-k)cp^k(1-p)^{n-k-1}$$

We find the critical points by letting $f'(p) = 0$ which implies

$$(n-k)cp^k(1-p)^{n-k-1} = ckp^{k-1}(1-p)^{n-k}.$$

Obviously we can immediately drop c. From here, we will assume $0 < p < 1$ for if $p = 1$ or $p = 0$, then $f(p)$ is always zero. With our assumtion, we have

$$\begin{aligned}
(n-k)p^k(1-p)^{n-k-1} &= kp^{k-1}(1-p)^{n-k} \\
(n-k)p &= k(1-p) \\
np - kp &= k - kp \\
np &= k \\
p &= \frac{k}{n}.
\end{aligned}$$

At this point, we are actually done (sort of). Normally, we'd take the second derivative and blah blah blah. In this case, we can do better. It isn't possible for p to be negative, and we know $0 \leq p \leq 1$. Therefore, if k/n is a critical point, it must be a maximum. A bit hand-waivy, but oh well. I'm not turning any of this in, so I'll cut corners every now and then. Just to be safe, I guess we can consider the case where $k = 0$ and $k = n$, as that seems to violate our assumption. If $k = 0$, then

$$f(p) = (1-p)^n$$

which is clearly maximized at $p = 0 = k/n$. If $k = n$, then

$$f(p) = p^k(1-p)^0 = p^k$$

which is clearly maximized at $p = 1 = k/n$. Although, this seems sketchy. Aren't we effectively applying 0^0, which is undefined? I graphed everything, so I'm convinced my final answer is correct, but I'm starting to think my argument is garbage. Anyway, I'm not doing this to test my ability to prove things. With more time, I could definitely clean up my argument, but I'd rather move on to other problems.

■

6. **Theoretical Exercise 19.** Show that if X is Poisson random variable with parameter λ, then

$$E[X^n] = \lambda E[(X+1)^{n-1}].$$

Then use this result to compute $E[X^3]$.

Solution. We can use a similar trick as in example 2.7 from Bertsekas's *Introduction to Probability.* Namely, we write

$$\begin{aligned}
E[X^n] &= \sum_{k=0}^{\infty} k^n e^{-\lambda} \frac{\lambda^k}{k!} \\
&= \sum_{k=1}^{\infty} k^n e^{-\lambda} \frac{\lambda^k}{k!} && [k = 0 \text{ term is zero}] \\
&= \lambda \sum_{k=1}^{\infty} k^{n-1} e^{-\lambda} \frac{\lambda^{k-1}}{(k-1)!} && \left[\frac{k^n}{k!} = \frac{k^{n-1}}{(k-1)!}\right] \\
&= \lambda \sum_{m=0}^{\infty} (m+1)^{n-1} e^{-\lambda} \frac{\lambda^m}{m!} && [m = k-1] \\
&= \lambda E[(X+1)^{n-1}]
\end{aligned}$$

as desired.

We now handle $E[X^3]$. We first expand $(X+1)^2$, as I don't see why that won't work. I really wanted there to be a nice, recursive solution, but that runs into problems. Suppose we try it. That is, define $Y_k = X + k$. Could we write

$$E[X] = \lambda E[Y_1^{n-1}] = \lambda^2 E[Y_2^{n-2}] = ... = \lambda^k E[Y_k^{n-k}] \quad ?$$

No! The random variable [1] Y is *a function* of a Poisson variable, but is not a Poisson variable itself. Therefore, the expectation trick does not necessarily hold (and in fact, it doesn't if you try). Therefore, we must expand $(X+1)^2$ and go from there:

$$\begin{aligned}
E[X^3] &= \lambda E[(X+1)^2] \\
&= \lambda E[X^2 + 2X + 1] \\
&= \lambda E[X^2] + 2\lambda E[X] + \lambda \\
&= \lambda^2 E[X+1] + 2\lambda E[X] + \lambda \\
&= \lambda^3 + \lambda^2 + 2\lambda^2 + \lambda && [\mu = \lambda \text{ for a Poisson random var}] \\
&= \lambda^3 + 3\lambda^2 + \lambda.
\end{aligned}$$

7. Define the covariance $\text{Cov}(XY) = E[XY] - E[X]E[Y]$...and then do some problems.

   a) 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 binlinear 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 $\text{Cov}(aX + bY, Z) = a\text{Cov}(X, Z) + b\text{Cov}(Y, Z)$.

[1] Also, we haven't verified that Y_k has a valid, PMF, so that's possibly another major flaw. Oh well.

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

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

Solution.

a) We first have

$$\text{Cov}(X, X) = E[XX] - E[X]E[X] = E[X^2] - E[X]^2$$

so indeed, $\text{Cov}(X, X) = \text{Var}(X)$. Next, we see that

$$\text{Cov}(Y, X) = E[YX] - E[Y]E[X]$$

but $E[Y]$ and $E[X]$ are real numbers, so $E[Y]E[X] = E[X]E[Y]$. I looked up joint distributions in Bertsekas and found that we can represent $E[YX]$ with nested sums (assuming each are discrete, and I think we require independence as well, but I can't remember). At any rate, we'd use that to argue $E[XY] = E[YX]$. Finally, we need to show that $\text{Cov}(aX+bY, Z) = a\text{Cov}(X, Z)+b\text{Cov}(Y, Z)$, which I believe will follow immediately from the linearity of expectation:

$$\begin{aligned}\text{Cov}(aX + bY, Z) &= E[aXZ + bYZ] - E[aX + bY]E[Z] \\ &= aE[XZ] + bE[YZ] - aE[X]E[Z] - bE[Y]E[Z] \\ &= a\text{Cov}(X, Z) + b\text{Cov}(Y, Z)\end{aligned}$$

as desired. NOTE: I made use of the results in section 4.9 of Ross's *A First Course in Probability*

b) We again will use linearity. First, write $\text{Cov}(X_1 - X_2, X_3 - 2X_4)$ as

$$E[(X_1 - X_2)(X_3 - 2X_4)] - E[X_1 - X_2]E[X_3 - 2X_4].$$

Now, notice that we can write $E[(X_1 - X_2)(X_3 - 2X_4)]$ as

$$E[X_1X_3] - 2E[X_1X_4] - E[X_2X_3] + 2E[X_2X_4]$$

and we can write $-E[X_1 - X_2]E[X_3 - 2X_4]$ as

$$-E[X_1]E[X_3] + 2E[X_1]E[X_4] + E[X_2]E[X_3] - 2E[X_2]E[X_4].$$

Grouping like terms together, we can rewrite all of the above as

$$\text{Cov}(X_1, X_3) - 2\text{Cov}(X_1, X_4) - \text{Cov}(X_2, X_3) + 2\text{Cov}(X_2, X_4).$$

Alternatively, if we used the last result from (a) we could immediately see that the above is correct. Finally, from the definition in the problem, we find the final answer to be $3 - 8 - 6 + 16 = 5$.

c) At this point, it's busy work. I'm moving on.

8. 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's 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 so that her new wealth is $W = W_0 + X$.

   a) 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, and σ^2

   b)

   c) Omitted.

   Solution. (Edit: I think my work might be wrong)

   a) Let's start with expanding the quadratic:

   $$\begin{aligned} E[U(W)] &= E[-(W_0 + X - x_0)^2] \\ &= E[-(W_0^2 + X^2 + x_0^2 + 2W_0X + -2W_0x_0 - 2x_0X)]. \end{aligned}$$

   From here, we can do quite a bit of clean up due to linearity of expectation. Namely, we write the above as

   $$-E[X^2] - 2W_0E[X] - 2x_0E[X] - W_0^2 - x_0^2 + 2W_0x_0$$

   which simplifies to

   $$-E[X^2] - 2W_0\mu - 2x_0\mu - W_0^2 - x_0^2 + 2W_0x_0.$$

   From here, we're almost done, but we need to get rid of the $-E[X^2]$ term. We recall that $\sigma^2 = E[X^2] - E[X]^2$ so if we add and subtract μ^2, we can simplify the expression. In particular observe that $-E[X^2] + \mu^2 - \mu^2$ factors to be come

   $$-(E[X^2] - mu^2) - mu^2 = -\sigma^2 - \mu^2$$

   so the expression

   $$-E[X^2] - (2W_0 + 2x_0)\mu - (W_0 - x_0)^2$$

   is equivalent to

   $$-\sigma^2 - \mu^2 - (2W_0 + 2x_0)\mu - (W_0 - x_0)^2$$

   and we're done. Or at least, I think we are. The expression above seems to me like it's negative, unless μ is negative and the $-(2W_0 + 2x_0)\mu$ term is sufficiently large. I probably dropped a minus sign somewhere. Normally, I'd try and find it, but I need to move on to continuous random varialbes if I'm going to finish the course by the desired date.

   b) Omitted.

c) Omitted.

9. Omitted (though I should probably come back to this at some point; I think it might relate to the previous problem).