Problem Set 3

Due: Start of class on Tuesday, March 4.

To help you prepare for the quiz, solutions will be posted shortly after lecture on Tuesday, March 4. As a result, no late assignments can be accepted.

Problem 1. Prove that if event A is independent of event B, then A is independent of $\overline{B}$.

Solution. We want to show that A is independent of $\overline{B}$; that is:

$$\Pr\left\{A \cap \overline{B}\right\} \;=\; \Pr\{A\} \cdot \Pr\left\{\overline{B}\right\}$$

Let's begin with the expression on the right and transform it into the expression on the left:

$$\begin{aligned}
\Pr\{A\} \cdot \Pr\left\{\overline{B}\right\} &= \Pr\{A\} \cdot (1 - \Pr\{B\}) \\
&= \Pr\{A\} - \Pr\{A\} \cdot \Pr\{B\} \\
&= \Pr\{A\} - \Pr\{A \cap B\} \\
&= \Pr\left\{A \cap \overline{B}\right\}
\end{aligned}$$

The first step uses the fact that the probabilities of complementary event sum to 1. The second step uses only algebra. In the third step, we use the fact that events A and B are independent. The final step uses the identity:

$$\Pr\{A\} \;=\; \Pr\{A \cap B\} + \Pr\left\{A \cap \overline{B}\right\}$$

This, in turn, follows from the sum rule, since the events $A \cap B$ and $A \cap \overline{B}$ are disjoint and their union is A.

Problem 2. Two teams of 6.042 students, the Amazings and the Unfair Dice, get together for a rowdy game of Rock, Scissor, Paper. The rules are as follows. In secret, each team selects rock, scissor, or paper. Then both teams announce their selections simultaneously. If both teams selected the same object, then the game is a draw. Otherwise, the team with the superior object wins, where superiority is determined according to the following rules:

- Rock smashes scissor, so rock is superior to scissor.
- Scissor cuts paper, so scissor is superior to paper.
- Paper covers rock, so paper is superior to rock.

Naturally enough, the Unfair Dice employ a strategy involving slightly skewed probabilities. They pick rock with probability 0.4, scissor with probability 0.2, and paper with probability 0.4. The Amazings like to "run with scissors"; they pick rock with probability 0.2, scissors with probability 0.7, and paper with probability 0.1. Assume that the teams select objects independently.

(a) What is the probability that the Amazings defeat the Unfair Dice, given that the outcome is not a draw?

Solution. Since the Amazings can choose three objects and the Unfair Dice can choose three objects, there are nine outcomes. We can denote these RR, RS, RP, SR, ... The first letter indicates the object chosen by the Amazings and the second indicates the object chosen by the Unfair Dice. Since the teams select their objects independently, the probability of an outcome (such as RP) is the product of the probability that the Amazings choose rock and the Unfair Dice choose paper ($0.2 \cdot 0.4 = 0.08$). Using these observations, we can compute the desired probability as follows:

$$\begin{aligned}
&\Pr\{\text{Amazings win} \mid \text{no draw}\} \\
&\quad= \frac{\Pr\{\text{Amazings win} \cap \text{no draw}\}}{\Pr\{\text{no draw}\}} \\
&\quad= \frac{\Pr\{\{RS, SP, PR\}\}}{\Pr\{\{RS, SP, PR, SR, PS, RP\}\}} \\
&\quad= \frac{0.2 \cdot 0.2 + 0.7 \cdot 0.4 + 0.1 \cdot 0.4}{0.2 \cdot 0.2 + 0.7 \cdot 0.4 + 0.1 \cdot 0.4 + 0.7 \cdot 0.4 + 0.1 \cdot 0.2 + 0.2 \cdot 0.4} \\
&\quad= \frac{18}{37} \\
&\quad= 0.486\ldots
\end{aligned}$$

(b) Tutor Sam Daitch boldly announces that one of *his* sections, the Roosters, can "mop the floor, clean the oven, and shovel the front walk" with the Unfair Dice. Assume that the Unfair Dice keep their earlier strategy. With what probability should the Roosters select each of rock, scissor, and paper so as to maximize the probability that they win, given that the outcome is not a draw?

Solution. Let's introduce some variables to capture the range of strategies available to the Roosters.

- Let r be the probability that the Roosters pick rock.

- Let s be the probability that the Roosters pick scissor.
- Let p be the probability that the Roosters pick paper.

Of course, we must have $r + s + p = 1$, since the Roosters must pick *some* object. Now let's get an expression for the probability that the Roosters win, given that the game is not a draw. (Here the first letter in an outcome indicates the Rooster's selection and the second indicates the Unfair Dice selection.)

$$\begin{aligned}
&\Pr\{\text{Roosters win} \mid \text{no draw}\} \\
&\quad = \frac{\Pr\{\text{Roosters win} \cap \text{no draw}\}}{\Pr\{\text{no draw}\}} \\
&\quad = \frac{\Pr\{\{RS, SP, PR\}\}}{\Pr\{\{RS, SP, PR, SR, PS, RP\}\}} \\
&\quad = \frac{r \cdot 0.2 + s \cdot 0.4 + p \cdot 0.4}{r \cdot 0.2 + s \cdot 0.4 + p \cdot 0.4 + s \cdot 0.4 + p \cdot 0.2 + r \cdot 0.4} \\
&\quad = \frac{0.2r + 0.4s + 0.4p}{0.6r + 0.8s + 0.6p}
\end{aligned}$$

The last expression is maximized by setting $p = 1, r = s = 0$. That is, the Roosters should always player paper. By doing so, the Roosters defeat the Unfair Dice with probability $\frac{0.4}{0.6} = \frac{2}{3}$.

Note on maximization: In the solution above, we needed to find the maximum of:

$$\frac{0.2r + 0.4s + 0.4p}{0.6r + 0.8s + 0.6p}$$

subject to the constraints that $r + s + p = 1$ and that all of r, s, and p must be nonnegative. Miniproblems like this, requiring interpretation and manipulation of expressions, come up all the time in applied mathematics. In this case, we blithely asserted that the maximum is $\frac{2}{3}$, which is achieved when $p = 1$ and $r = s = 0$. How did we get this answer?

There are general methods for solving such problems, but we'll neither teach them nor expect you to know them in 6.042. Rather, we used an ad hoc technique based on a simple observation: if we make one variable big, the value of the whole fraction approaches the ratio of the coefficients of that variable. For example, if we make r big, then the fraction approaches the ratio $0.2/0.6 = 1/3$. If we make s big, the fraction approaches $0.4/0.8 = 1/2$. However, if we make p big, the fraction approaches $0.6/0.8 = 2/3$, which is larger than the other two ratios. Thus, in order to maximize the fraction, we should make p as large as possible relative to the other two variables. On this basis, we chose $p = 1$ and $r = s = 0$.

Alternatively, this result can be obtained by rewriting the fraction using the constraint $r + s + p = 1$.

$$\begin{aligned}\frac{0.2r + 0.4s + 0.4p}{0.6r + 0.8s + 0.6p} &= \frac{r + 2s + 2p}{3r + 4s + 3p} \\ &= \frac{r + 2s + 2p + 2 - 2(r + s + p)}{3r + 4s + 3p + 3 - 3(r + s + p)} \\ &= \frac{2 - r}{3 + s}\end{aligned}$$

This last expression is decreasing in both r and s. Thus, to maximize it we should choose $r = s = 0$ and, consequently, $p = 1$.

Problem 3. Prove or disprove each of the following assertions.

(a) For all events A, B, and C, if A is independent of B and A is independent of C, then A is independent of $B \cap C$.

Solution. This is false, as seen in the two coin example in lecture. Suppose that we flip two fair, independent coins.

- Let A be the event that the first coin shows heads.
- Let B be the event that the second coin shows heads.
- Let C be the event that the coins agree; that is, they both show heads or both show tails.

As we saw in lecture, event A is independent of events B and C individually. However, event A is independent of $B \cap C$ if and only if:

$$\Pr\{A \cap (B \cap C)\} = \Pr\{A\} \cap \Pr\{B \cap C\}$$

The expression on the left is $\frac{1}{4}$, since the event $A \cap B \cap C$ consists of the single outcome in which both coins are heads. However, the expression on the right is $\frac{1}{2} \cdot \frac{1}{4} = \frac{1}{8}$. Since equality does not hold, A is not independent of $B \cap C$.

(b) For all events A, B, and C, if A is independent of B and A is independent of C, then A is independent of $B \cup C$.

Solution. This is also false. Define events A, B, and C as before. Event A is independent of event $B \cup C$ if and only if:

$$\Pr\{A \cap (B \cup C)\} = \Pr\{A\} \cdot \Pr\{B \cup C\}$$

The probability on the left is $\frac{1}{4}$, since the event $A \cap (B \cup C)$ occurs only when both coins are heads. On the other hand, the probability on the right is $\frac{1}{2} \cdot \frac{3}{4} = \frac{3}{8}$. (The second probability is $\frac{3}{4}$, because the event $B \cup C$ occurs unless the first coin is heads and the second is tails.) Therefore, event A is not independent of event $B \cup C$.

Problem 4. Suppose that you flip three fair, mutually independent coins. Define the following events:

- Let A be the event that the first coin is heads.
- Let B be the event that the second coin is heads.
- Let C be the event that the third coin is heads.
- Let D be the event that an even number of coins are heads.

(a) Are these events pairwise independent?

Solution. The sample space for this experiment consists of eight, equally-probably outcomes:

$$\{H, T\}^3 = \{HHH, HHT, HTH, HTT, THH, \ldots\}$$

Each sequence of three symbols specifies the outcomes of the three coin tosses.

The first three events (A, B, and C) are pairwise independent, since they are mutually independent. All that remains is to check that each of these is independent of D. Since they are symmetric, it suffices to check just one of the three, say A:

$$\begin{aligned}
\Pr\{A\} &= \frac{1}{2} \\
\Pr\{D\} &= \Pr\{\{HHT, HTH, THH, TTT\}\} \\
&= \frac{1}{2} \\
\Pr\{A \cap D\} &= \Pr\{\{HHT, HTH\}\} \\
&= \frac{1}{4}
\end{aligned}$$

Therefore, $\Pr\{A \cap D\} = \Pr\{A\} \cdot \Pr\{D\}$, and so these events are independent. We conclude that all four events are pairwise independent.

(b) Are these events three-way independent? That is, does

$$\Pr\{X \cap Y \cap Z\} = \Pr\{X\} \cdot \Pr\{Y\} \cdot \Pr\{Z\}$$

always hold when X, Y, and Z are distinct events drawn from the set $\{A, B, C, D\}$?

Solution. Because the coin tosses are mutually independent, we know:

$$\Pr\{A \cap B \cap C\} = \Pr\{A\} \cdot \Pr\{B\} \cdot \Pr\{C\}$$

What remains is to check that equality holds for the other subsets of three events: $\{A, B, D\}$, $\{A, C, D\}$, and $\{B, C, D\}$. Since these three are symmetric, it suffices to check only one, say the first.

$$\begin{aligned} \Pr\{A \cap B \cap D\} &= \Pr\{\{HHT\}\} \\ &= \frac{1}{8} \end{aligned}$$

Since this is equal to $\Pr\{A\} \cdot \Pr\{B\} \cdot \Pr\{D\}$, these three events are independent. Adding this to our earlier observations, we conclude that all four events are three-way independent.

(c) Are these events mutually independent?

Solution. No, because:

$$\Pr\{A \cap B \cap C \cap D\} \neq \Pr\{A\} \cdot \Pr\{B\} \cdot \Pr\{C\} \cdot \Pr\{D\}$$

The probability on the left is zero, but the product on the right is equal to $\frac{1}{16}$.

Problem 5. Eight people go out for dinner at a restaurant and sit down at a round table. Everyone orders a different meal, and soon the waiter returns with those eight dishes. Unfortunately, he has largely forgotten who ordered what. He manages only to serve *three* of the eight people what they ordered; each of the other five diners gets someone else's meal. One diner suggests that they spin the table and see if chance can do better than the befuddled waiter. If they spin the table to a new orientation, selected uniformly at random, what is the expected number of people that end up with the correct meal in front of them?

Solution. The sample space consists of the eight possible orientations of the table, which are equally likely. Let the random variable R denote the number of people that receive the correct meal after the table is spun. More specifically, let R_i be an indicator random variable for the event that the i-th person gets the right meal. Then we we have:

$$R = R_1 + R_2 + \ldots + R_8$$

Taking the expected value of both sides and applying linearity of expectation gives:

$$\begin{aligned} \text{Ex}[R] &= \text{Ex}[R_1 + R_2 + \ldots + R_8] \\ &= \text{Ex}[R_1] + \text{Ex}[R_2] + \ldots + \text{Ex}[R_8] \end{aligned}$$

Note that $\Pr\{R_i\} = \frac{1}{8}$, since there is exactly one orientation of the table that aligns the i-th person with his or her meeal. Also, recall that the expectation of an indicator is equal to the probability that the indicator is 1. Therefore, we have:

$$\begin{aligned} \text{Ex}[R] &= \text{Ex}[R_1] + \text{Ex}[R_2] + \ldots + \text{Ex}[R_8] \\ &= \Pr\{R_1 = 1\} + \Pr\{R_2 = 1\} + \ldots + \Pr\{R_1 = 1\} \\ &= \frac{1}{8} + \frac{1}{8} + \ldots + \frac{1}{8} \\ &= 1 \end{aligned}$$

This is *exactly* the argument used for the hat check problem used in lecture. The information about three diners receiving the correct dish initially was irrelevant!

Problem 6. Let R be a random variable, and let E be an event with nonzero probability. The *conditional expectation* of R, given event E, is denoted $\text{Ex}[R \mid E]$ and defined by:

$$\text{Ex}[R \mid E] = \sum_{x \in \text{Range}(R)} x \cdot \Pr\{R = x \mid E\}$$

(a) Compute the expected value of the number rolled on a fair, six-sided die, given that the outcome is even.

Solution. Let the random variable R be the number rolled, and let E be the event that the number rolled is even.

$$\begin{aligned} \text{Ex}[R \mid E] &= \sum_{x=1}^{6} x \cdot \Pr\{R = x \mid E\} \\ &= 1 \cdot 0 + 2 \cdot \frac{1}{3} + 3 \cdot 0 + 4 \cdot \frac{1}{3} + 5 \cdot 0 + 6 \cdot \frac{1}{3} \\ &= 4 \end{aligned}$$

(b) Compute the expected value of the number rolled on a fair, six-sided die, given that the outcome is odd.

Solution. Define R and E as before.

$$\begin{aligned}
\text{Ex}\left[R \mid \overline{E}\right] &= \sum_{x=1}^{6} x \cdot \Pr\left\{R = x \mid \overline{E}\right\} \\
&= 1 \cdot \frac{1}{3} + 2 \cdot 0 + 3 \cdot \frac{1}{3} + 4 \cdot 0 + 5 \cdot \frac{1}{3} + 6 \cdot 0 \\
&= 3
\end{aligned}$$

(c) Let R be a random variable, and let E be an event such that $0 < \Pr\{E\} < 1$. Prove the following identity:

$$\text{Ex}[R] = \text{Ex}[R \mid E] \cdot \Pr\{E\} + \text{Ex}\left[R \mid \overline{E}\right] \cdot \Pr\left\{\overline{E}\right\}$$

This is called the *Total Expectation Law*. It is analogous to the Total Probability Law; it lets you break down the calculation of a complicated expectation into cases based on whether or not the event E occurs.

Solution. Let S be the sample space. We can transform the left side into the right as follows:

$$\begin{aligned}
&\text{Ex}[R] \\
&\quad= \sum_{s \in S} R(s) \cdot \Pr\{s\} \\
&\quad= \sum_{s \in S} R(s) \cdot \left(\Pr\{s \mid E\} \cdot \Pr\{E\} + \Pr\left\{s \mid \overline{E}\right\} \cdot \Pr\left\{\overline{E}\right\}\right) \\
&\quad= \left(\sum_{s \in S} R(s) \cdot \Pr\{s \mid E\} \cdot \Pr\{E\}\right) + \left(\sum_{s \in S} R(s) \cdot \Pr\left\{s \mid \overline{E}\right\} \cdot \Pr\left\{\overline{E}\right\}\right) \\
&\quad= \Pr\{E\} \cdot \left(\sum_{s \in S} R(s) \cdot \Pr\{s \mid E\}\right) + \Pr\left\{\overline{E}\right\} \cdot \left(\sum_{s \in S} R(s) \cdot \Pr\left\{s \mid \overline{E}\right\}\right) \\
&\quad= \Pr\{E\} \cdot \text{Ex}[R \mid E] + \Pr\left\{\overline{E}\right\} \cdot \text{Ex}\left[R \mid \overline{E}\right]
\end{aligned}$$

The first step uses the definition of expectation. In the second step, we apply the law of total probability. Next, we split the one sum into two and pull out the constants $\Pr\{E\}$ and $\Pr\left\{\overline{E}\right\}$. Finally, we apply the definition of conditional expectation.

(d) Define the random variable R using the following procedure. First, flip a fair coin. If the result is heads, then roll a fair, six-sided die and let R equal the number that comes up. If the result is tails, then roll two fair, independent, six-sided dice and let R be the sum of the numbers that come up. What is the expected value of R?

Solution. Recall that the expectation of a fair die is $\frac{7}{2}$. From the total expectation law and linearity of expectation, we have:

$$\begin{aligned}
\text{Ex}[R] &= \frac{1}{2} \cdot \frac{7}{2} + \frac{1}{2} \cdot \left(\frac{7}{2} + \frac{7}{2} \right) \\
&= \frac{21}{4}
\end{aligned}$$

Problem 7. In a certain card game, each card has a point value.

- Numbered cards in the range 2 to 9 are worth five points each.
- The card numbered 10 and the face cards (jack, queen, king) are worth ten points each.
- Aces are worth fifteen points each.

(a) Suppose that you shuffle a 52-card deck. What is the expected total point value of the top three cards?

Solution. Let the random variable X be the total point value of the top three cards. Then we can write:

$$X = X_1 + X_2 + X_3$$

Here the random variables X_1, X_2, and X_3 are the point values of the first, second, and third cards. By the definition of expectation:

$$\begin{aligned}
\text{Ex}[X_i] &= \sum_{r \in X_i(S)} r \cdot \Pr\{X_i = r\} \\
&= 5 \cdot \frac{8}{13} + 10 \cdot \frac{4}{13} + 15 \cdot \frac{1}{13} \\
&= \frac{95}{13}
\end{aligned}$$

Now we can solve the problem by taking the expected value of both sides of our original equation and then using linearity of expectation:

$$\begin{aligned}
\text{Ex}[X] &= \text{Ex}[X_1 + X_2 + X_3] \\
&= \text{Ex}[X_1] + \text{Ex}[X_2] + \text{Ex}[X_3] \\
&= \frac{95}{13} + \frac{95}{13} + \frac{95}{13} \\
&= \frac{285}{13}
\end{aligned}$$

(b) Suppose that you throw out all the red cards and shuffle the remaining 26-card, all-black deck. Now what is the expected total point value of the top three cards? (Note that drawing three aces, for example, is now impossible!)

Solution. The expected point value is the same as before, since expected point value of a single card is unchanged. Nothing in our solution assumed a 52 card deck.

Problem 8. The Mostly Harmless tutorial section gets together to play a board game. The game consists of ten rounds. In each round, each player rolls two dice, takes the sum of the results, and then advances his or her gamepiece by that many squares. Reportedly, one of the players— and we're not saying who— likes to whine. *A lot.* In fact, whenever this player rolls a pair of ones, he complains so much that the other players graciously allow him to reroll until he gets something better. At the end of a game, how far ahead of a normal, civil player can the whiner expect to be?

Solution. First, let's determine how far the whiner advances in a single round. Let the random variable D be sum of two fair, independent dice, and let P be the event that both rolls are 1's. Our goal is to compute $\text{Ex}\left[D \mid \overline{P}\right]$. We can do this using the total expectation law:

$$\begin{aligned}
\text{Ex}[D] &= \text{Ex}[D \mid P] \cdot \Pr\{P\} + \text{Ex}\left[D \mid \overline{P}\right] \cdot \Pr\left\{\overline{P}\right\} \\
\text{Ex}\left[D \mid \overline{P}\right] &= \frac{\text{Ex}[D] - \text{Ex}[D \mid P] \cdot \Pr\{P\}}{\Pr\left\{\overline{P}\right\}} \\
&= \frac{7 - 2 \cdot \frac{1}{36}}{\frac{35}{36}} \\
&= \frac{50}{7}
\end{aligned}$$

Here we've used the fact that $\text{Ex}[D]$, the expected sum of two dice, is $\frac{7}{2} + \frac{7}{2} = 7$.

Turning to the game as a whole, let the random variables $C_1, C_2, \ldots, C_{10}$ denote the rolls of a civil player, and let $W_1, W_2, \ldots, W_{10}$ denote the rolls of the whiner. The answer now follows from linearity of expectation:

$$\begin{aligned}
\operatorname{Ex}\left[\sum_{i=1}^{10} W_i - \sum_{i=1}^{10} C_i\right] &= \sum_{i=1}^{10} \operatorname{Ex}[W_i] - \operatorname{Ex}[C_i] \\
&= 10 \cdot \frac{50}{7} - 10 \cdot 7 \\
&= \frac{10}{7}
\end{aligned}$$

Thus, on average, the whiner comes out more than one square ahead overall.