Logic, Reasoning, and Persuasion 07; Midterm 2

Monday, December 8th, 2025 — Written by Adrian Liu

- 100 points (20 + 40 + 10 + 10 + 15 + 5).
- Please write your answers *in the bluebook*. Nothing written on this exam paper will be graded. Incorrect answers get no points unless work is shown.
- You are allowed to use the formula sheet plus a double-sided sheet of paper for the exam. You can use a calculator, but you should not need it. The test is otherwise closed-book, closed-notes, and no use of computers.
- You can leave answers as fractions if you like: no need to reduce to simplest form or a decimal. But they should be numbers!

Honor Code

Please Write out the Rutgers Honor Code on your Bluebook: *On my honor, I have neither received nor given any unauthorized assistance on this examination (assignment).*

1 | Formal Practice (20 points)

Suppose you have three hypotheses, H_1, H_2, H_3. You know *exactly one* of those hypotheses is true.

(1.1) (5 points) *What is $Pr(H_1) + Pr(H_2) + Pr(H_3)$? Do not use numbers from below.*

→ **Solution:** By Partitionality (on the sheet), the probability is $\boxed{1}$ (even if you know nothing else about the probabilities!)

Right now, your estimations are that $Pr(H_1) = 0.2$, $Pr(H_2) = 0.3$, $Pr(H_3) = 0.5$. You're wondering whether some statement A is true. You know that

$$Pr(A \mid H_1) = 1/2, Pr(A \mid H_2) = 2/3, Pr(A \mid H_3) = 3/5.$$

$$\begin{aligned} Pr(A \mid H_1) &= 1/2 \\ Pr(A \mid H_2) &= 2/3 \\ Pr(A \mid H_3) &= 2/5 \end{aligned}$$

(1.2) (5 points) *What is $Pr(A)$? Use the law of total probability.*

→ **Solution:** By total probability,

$$\begin{aligned} Pr(A) &= Pr(A \mid H_1)Pr(H_1) + Pr(A \mid H_2)Pr(H_2) + Pr(A \mid H_3)Pr(H_3) \\ &= 1/2 \cdot 0.2 + 2/3 \cdot 0.3 + 2/5 \cdot 0.5 = 0.1 + 0.2 + 0.2 = \boxed{0.5} \end{aligned}$$

(1.3) (5 points) *Suppose you learn that A is in fact true. If you used conditionalization to update your opinions, what would your new estimates about the hypotheses H_1, H_2, H_3 be? Use Conditionalization and Bayes' Theorem.*

(a) $Pr_A(H_1) =$

$$\begin{aligned} Pr_A(H_1) &= Pr(H_1 \mid A) && \text{(By Conditionalization)} \\ &= Pr(H_1)\frac{Pr(A \mid H_1)}{Pr(A)} && \text{(By Bayes' Theorem)} \\ &= 0.2 \cdot \frac{0.5}{0.5} = \boxed{0.2}. \end{aligned}$$

(b) $Pr_A(H_2) =$

$$\begin{aligned} Pr_A(H_2) &= Pr(H_2 \mid A) && \text{(By Conditionalization)} \\ &= Pr(H_2)\frac{Pr(A \mid H_2)}{Pr(A)} && \text{(By Bayes' Theorem)} \\ &= 0.3 \cdot \frac{2/3}{0.5} = 0.2/0.5 = \boxed{0.4}. \end{aligned}$$

(c) $Pr_A(H_3) =$

$$\begin{aligned} Pr_A(H_3) &= Pr(H_3 \mid A) && \text{(By Conditionalization)} \\ &= Pr(H_3)\frac{Pr(A \mid H_3)}{Pr(A)} && \text{(By Bayes' Theorem)} \\ &= 0.5 \cdot \frac{2/5}{0.5} = 2/5 = \boxed{0.4}. \end{aligned}$$

(1.4) (5 points) *Evidence:*

(a) *Is A evidence for H_1?*
(b) *Is A evidence for H_2?*
(c) *Is A evidence for H_3?*
(d) *Is H_1 evidence for A?*
(e) *Is H_2 evidence for A?*
(f) *Is H_3 evidence for A?*

→ **(a):** No. By the evidence lemma, A is evidence for H_1 if $Pr(H_1 \mid A) > Pr(H_1)$, A is evidence against H_1 if $Pr(H_1 \mid A) < Pr(H_1)$, and A is independent of H_1 if $Pr(H_1 \mid A) = Pr(H_1)$. Since $Pr(H_1 \mid A) = 0.2$ and $Pr(H_1) = 0.2$, A is $\boxed{\text{independent of}}$ H_1

→ **(b):** Yes. Since $Pr(H_2 \mid A) = 0.4$ and $Pr(H_2) = 0.3$, we have $Pr(H_2 \mid A) > Pr(H_2)$, so A is $\boxed{\text{evidence for}}$ H_2.

→ **(c):** No. Since $Pr(H_3 \mid A) = 0.4$ and $Pr(H_1) = 0.5$, we have $Pr(H_3 \mid A) < Pr(H_3)$, so A is $\boxed{\text{evidence against}}$ H_3.

→ **(d)-(f):** By the evidence Lemma, If A is evidence for H then H is evidence for A. So we don't have to do the calculations again. H_1 is $\boxed{\text{independent of}}$ A, H_2 is $\boxed{\text{evidence for}}$ A, and H_3 is $\boxed{\text{evidence against}}$ A.

2 | Food Insecurity (40 points)

You're a social worker at a school trying to determine whether a student is facing food insecurity at home. The student is sometimes disruptive in class and is listless in class. Their grades are not great. Based on this, you think there's a 50% chance that they face food insecurity at home, explaining these symptoms. Students often will not tell you outright if they are facing food insecurity. However, you can have them take a diagnostic with the school nurse. The diagnostic's accuracy is as follows:

1. If a student faces food insecurity at home, there is a 80% chance the diagnostic comes back positive (+) and a 20% chance that it comes back negative (−).
2. If a student does **not** face food insecurity, there is a 40% chance the diagnostic comes back positive (+) and a 60% chance that it comes back negative (−).

(2.1) (5 points) *Let F be the statement "the student is facing food insecurity". Let $+$ be the statement "the diagnostic comes back positive". Write out all the probabilities you already know based on the problem statement.*

(a) $Pr(F) = 0.5$
(b) $Pr(\neg F) = 0.5$
(c) $Pr(+ \mid F) = 0.8$
(d) $Pr(- \mid F) = 0.2$
(e) $Pr(+ \mid \neg F) = 0.6$
(f) $Pr(- \mid \neg F) = 0.4.$

(2.2) (5 points) *You've ordered the diagnostic, but haven't seen the results yet. How likely do you think it is that diagnostic comes back positive?*

$$\begin{aligned} Pr(+) &= Pr(+ \mid F) \cdot Pr(F) + Pr(+ \mid \neg F) \cdot Pr(\neg F) \qquad \text{(Total Probability)} \\ &= 0.8 * 0.5 + 0.4 * 0.5 = \boxed{0.6.} \end{aligned}$$

(2.3) (5 points) *How likely do you think it is that the student faces food insecurity, supposing that their diagnostic is positive?*

$$Pr(F \mid +) = Pr(F)\frac{Pr(+ \mid F)}{Pr(+)} = 0.5\frac{0.8}{0.6} = \boxed{2/3 \approx 0.67.} \qquad \text{(Bayes)}$$

(2.4) (5 points) *How likely do you think it is that the student faces food insecurity, supposing that their diagnostic is negative?*

$$Pr(F \mid -) = Pr(F)\frac{Pr(- \mid F)}{Pr(-)} = Pr(F)\frac{Pr(- \mid F)}{1 - Pr(+)} = 0.5\frac{0.2}{0.4} = \boxed{0.25.} \quad \text{(Bayes)}$$

In this school district, around 25% of students face food insecurity.

(2.5) (5 points) *What is the probability that a random student in the school district would test positive on this diagnostic?*

Let F_R be the probability that a random student in the school district has food insecurity. Then $F_R = 0.25$.

$$\begin{aligned} Pr(+) &= Pr(+ \mid F_R)Pr(F_R) + Pr(+ \mid \neg F_R)Pr(\neg F_R) && \text{(Total Probability)} \\ &= 0.8 \cdot 0.25 + 0.4 \cdot 0.75 = \boxed{0.5.} \end{aligned}$$

(2.6) (5 points) *School statistics show that around 70% of diagnostics run in this district come back positive. If you did the math right, your answer to the previous question is lower than this. How could this be?*

About 70% of students test positive. But if you ran the diagnostic randomly on students, you would expect about 50% to test positive. This must mean that the students actually administered the diagnostic are more likely to face food insecurity than the general population of students in the district. This makes sense, if we consider the counselors would tend to run the diagnostic on students who they already suspect are more likely to be facing food insecurity.

(In fact, I calculated the 70% by supposing that, on average, students for whom counselors had diagnostics ran were already 70% likely to be facing food insecurity on average. Let F_S be the probability that some student who got the diagnostic actually faces food insecurity. Then, if we use total probability,

$$\begin{aligned} Pr(+) &= Pr(+ \mid F_S)Pr(F_S) + Pr(+ \mid \neg F_S)Pr(\neg F_S) && \text{(Total Probability)} \\ &= 0.8 \cdot 0.7 + 0.4 \cdot 0.3 = 0.68 \approx 70\%.) \end{aligned}$$

The diagnostic comes back **positive**.

(2.7) (5 points) *Are the diagnostic results evidence for or against the hypothesis that the student faces food insecurity? Use the Evidence Lemma.*
The diagnostic results, because they've come back positive, are evidence $\boxed{\textit{for}}$ the hypothesis that the student faces food insecurity. The probability that the student faces food insecurity, *given* that they diagnostic is positive, is higher than the prior probability that the student faces food insecurity.

$$Pr(F \mid +) = 0.67 > 0.5 = Pr(F).$$

(2.8) (5 points) *What should your new estimate be for how likely the student is to be facing food insecurity?*

$$Pr_+(F) = Pr(F \mid +) = \boxed{0.67.}$$

3 | Peer Disagreement (10 points)

You and your friend Xavier have a bet about whether the Lions will make the playoffs.

1. You think that either both ~~of them~~ the Lions and the Chiefs will make the playoffs, or neither of them will. According to your estimates, there's a 50% chance they both make the playoffs, and a 50% chance that neither of them make the playoffs.
2. Xavier thinks that *exactly one* of the teams will make the playoffs. There's a 50% chance that the Lions make the playoffs but the Chiefs don't, and a 50% chance that the Chiefs make the playoffs but the Lions don't.

(3.1) (5 points) *How likely are the Lions to make the playoffs according to you? How likely are the Lions to make the playoffs according to Xavier?*

(a) You think there's a 50% chance that both the Lions and the Chiefs make the playoffs, and a 50% chance that neither make the playoffs. If both make the playoffs, then the Lions make the playoffs. If neither makes the playoffs, then the Lions do not make the playoffs. So you think there's a $\boxed{50\%}$ chance the Lions make the playoffs.

(b) Xavier thinks there's a 50% chance that the Lions make the playoffs but the Chiefs don't, and a 50% chance that the Chiefs make the playoffs but the Lions don't. So Xavier also thinks there's a $\boxed{50\%}$ chance the Lions make the playoffs.

You and Xavier agree that if the Colts were to win next week, this would make it *less likely* that the Chiefs make the playoffs. Thus, learning that the Colts win next week would be evidence against the Chiefs making the playoffs. In particular, given that the Colts win, you and Xavier both think there would be only a 30% chance that the Chiefs make the playoffs.

(3.2) (5 points between 3.2 and 3.3) *Suppose the Colts win. Does this increase or decrease your confidence that the Lions will make the playoffs?*

→ This *decreases* your confidence. If the Colts win, then the Chiefs are less likely to make the playoffs. And so, based on your opinions, this means the Lions are also $\boxed{\text{less}}$ likely to make the playoffs (since either both of them make the playoffs or neither does.)

(3.3) (5 points between 3.2 and 3.3) *Suppose the Colts win. Does this increase or decrease Xavier's confidence that the Lions will make the playoffs?*

→ This *increases* Xavier's confidence. If the Colts win, then the Chiefs are less likely to make the playoffs. And so, based on Xavier's opinions, this means the Lions are $\boxed{\text{more}}$ likely to make the playoffs, since Xavier thinks that the Lions make the playoffs exactly if the Chiefs do *not*.

4 | Bookbags (10 points)

This is a variant of the problem from Monday used to illustrate Hindsight Bias:

> There are two bookbags, one containing 600 red and 300 blue chips, the other containing 300 red and 600 blue. Take one of the bags. Now, you sample, randomly, with replacement after each chip. You draw three times and get this sequence:
>
> blue blue blue
>
> What is the probability that you chose the bookbag with mostly blue chips (600 blue, 300 red)?

Assume that the samplings are *independent*: each sample does not affect the probabilities of the other samples. This means that you can multiply the probabilities of two sample outcomes to get the probability that both occurred:

> If A and B are independent, then $Pr(A\&B) = \Pr(A) \cdot \Pr(B)$.

For example, $Pr(\text{1st is red AND 2nd is blue}) = Pr(\text{1st is red}) \cdot Pr(\text{2nd is blue})$.

(4.1) (1 exit-quiz point): ***without doing any calculations***, *estimate the probability that you chose the bookbag with mostly blue chips (600 blue, 300 red), to the nearest 5%. Just give your guess:*

→ My guess was around 75%, for what it's worth.

(4.2) (5 points) *Solve rigorously for the probability that you chose the bookbag with mostly blue chips, given that you drew three straight blue chips.*

Solution: Let S be the statement "you get the sequence blue, blue, red, blue, red", B be the statement "you chose the bag with mostly blue chips," and R be the statement "you chose the bag with mostly red chips".

What is the problem asking? The problem statement says we've sampled and gotten a sequence of five (S), and is asking how likely it is that we chose the bookbag with the mostly blue chips (B). That is, we learn that S is true. So in the setup we learn something, like in this diagram:

$$Pr(B) \longrightarrow \text{learn } S \longrightarrow Pr_S(B).$$

The problem is asking us what $Pr_S(B)$ is. We'll assume that we do conditionalization, so that $Pr_S(B) = Pr(B \mid S)$. So the problem is really asking us what $Pr(B \mid S)$ is.

1. **Setting Up**. Given the setup of the case, we already know the following probabilities:

   $$Pr(B) = 0.5 \qquad\qquad Pr(R) = 0.5.$$

   We can also solve for the probability that we would get the sequnce S from the mostly-blue bag or the mostly-red bag. Let a number followed by a colon and a color (like 1:blue) mean that that number draw was that color. Since the samples are independent, we have

$$\begin{aligned} Pr(S \mid B) &= Pr(1\text{:blue}|B)Pr(2\text{:blue}|B)Pr(3\text{:blue}|B) \\ &= 2/3 \cdot 2/3 \cdot 2/3 = 8/27 \\ Pr(S \mid R) &= Pr(1\text{:blue}|R)Pr(2\text{:blue}|R)Pr(3\text{:blue}|R) \\ &= 1/3 \cdot 1/3 \cdot 1/3 = 1/27 \end{aligned}$$

So we have

$$\begin{aligned} Pr(B) &= 1/2 & Pr(R) &= 1/2. \\ Pr(S \mid B) &= 8/27 & Pr(S \mid R) &= 1/27. \end{aligned}$$

2. **Solving for** $Pr(B \mid S)$. Should we use Bayes' Theorem or Total Probability to figure out $Pr(B \mid S)$? Because this is a *conditional* probability, with a bar in the middle, we should use **Bayes' Theorem**.

   We write out Bayes' Theorem with the right statement letters:

   $$Pr(B \mid S) = Pr(B)\frac{Pr(S \mid B)}{Pr(S)}.$$

   Now, we already know (from our list above) that $Pr(B) = 1/2$ and $Pr(S \mid B) = 1/27$. So we can substitute these in:

   $$Pr(B \mid S) = 1/2\frac{1/27}{Pr(S)}.$$

   But we don't know what $Pr(S)$ is, so we'll have to calculate it.

3. **Finding** $Pr(B \mid S)$. Should we use Bayes' Theorem or Total Probability to figure out $Pr(B \mid S)$? Because this is a *unconditional* probability, with *no* bar in the middle, we should use **Total Probability**.

   What are the two ways in which S could be true? Well, we could either have gotten the sequence from the mostly blue bag or from the mostly red bag. So the two hypotheses, exactly one of which is true, are B and R. Let's write out total probability with the right statement letters:

   $$Pr(S) = Pr(S \mid B)Pr(B) + Pr(S \mid R)Pr(R).$$

   And we see from above that we know all these values already, so let's substitute them in:

   $$Pr(S) = 8/27 \cdot 1/2 + 1/27 \cdot 1/2 = 8/54 + 1/54 = 9/54 = 1/6.$$

4. **Finding** $Pr(B \mid S)$**, resumed**: Now that we have $Pr(S) = 1/6$ we can substi-

tute it back into Bayes' Theorem:

$$Pr(B \mid S) = 1/2 \frac{8/27}{1/6} = 4/27 \cdot 6 = 24/27 = 8/9.$$

So our final answer is $\boxed{Pr(B \mid S) = 8/9 \approx 89\%.}$

(4.3) (5 points) *What is the probability that the next chip you draw out of the bookbag you chose is blue?*

Solution: Because $Pr_S(B) = Pr(B \mid S) = 8/9$, we know that $Pr_S(R) = 1 - Pr_S(B) = 1 - 8/9 = 1/9$. We can use total probability to calculate the probability that the next chip will be blue:

$$\begin{aligned} Pr_S(\text{4:blue}) &= Pr_S(\text{4:blue} \mid B)Pr_S(B) + Pr_S(\text{4:blue} \mid R)Pr_S(R) \\ &= 2/3 \cdot 8/9 + 1/3 \cdot 1/9 = 16/27 + 1/27 = \boxed{17/27 \approx 64\%.} \end{aligned}$$

5 | GPT Checker (15 points)

You submit a paper that you've written entirely by yourself to Canvas. Canvas has an integrated AI checker that is 90% accurate. This means that:

1. If the text was AI generated, the checker would mark it as AI generated 90% of the time.
2. If the text was not AI generated, the checker would mark it as AI generated 10% of the time.

The AI checker results are shown to the professor, not you.

(5.1) (5 points) *You know for sure that you didn't use AI: $Pr(AI) = 0$. How likely is it that the AI checker will mark your essay as AI generated?*

→ By total probability,

$$\begin{aligned} Pr(Flag) &= Pr(Flag \mid AI)Pr(AI) + Pr(Flag \mid \neg AI)Pr(\neg AI) \\ &= 0.9 \cdot 0 + 0.1 \cdot 1 = \boxed{0.1.} \end{aligned}$$

Your professor reaches out to you. She says that she knows you've been a responsible student throughout the class, and she previously thought there was only around a 10% chance that you would use AI to write your essay. But the AI checker is very confident, so she wants to talk to you about what might have happened.

Let AI be the statement that you used AI to write the essay. Let $Flag$ be the statement that the essay was marked as AI generated.

(5.2) (5 points) *After learning that the integrated AI checker has marked it as AI generated, how sure should you be that you didn't use AI?*

→ Previously, you were already sure you didn't use AI: So $Pr(\neg AI) = 1$. So by Bayes' Theorem:

$$Pr(\neg AI \mid Flag) = Pr(\neg AI)\frac{Pr(Flag \mid \neg AI)}{Pr(Flag)} = 1 \cdot \frac{0.1}{0.1} = \boxed{1.}$$

This makes sense. If you already *know* you didn't use AI, then no AI checker could convince you otherwise. (More generally, when you have prior opinion 1 or 0, Conditionalizing with Bayes' Theorem can *never* change your opinion.)

The teacher says that the evidence is really hard to deny: it suggests that there's a 90% chance you used AI to generate your essay. Luckily, you've taken PHIL101, so you know that this isn't correct.

(5.3) (5 points) *How likely should your professor now think it is that you used AI to generate your essay, given that your essay got flagged? Remember, she said she previously thought there was a 10% chance that you would use AI to write your paper.*

→ Previously, the teacher was 10% confident that you would use AI. So $Pr^T(AI) = 0.1$, where I'm using Pr^T to denote your teacher's probabilities. So by Bayes' Theorem:

$$Pr^T(AI \mid Flag) = Pr^T(AI)\frac{Pr^T(Flag \mid AI)}{Pr^T(Flag)} = 0.1 \cdot \frac{0.9}{Pr^T(Flag)}.$$

We don't know $Pr^T(Flag)$ though. Here again we use total probability to figure it out. By total probability,

$$\begin{aligned} Pr^T(Flag) &= Pr^T(Flag \mid AI)Pr^T(AI) + Pr^T(Flag \mid \neg AI)Pr^T(\neg AI) \\ &= 0.9 \cdot 0.1 + 0.1 \cdot 0.9 = 0.18. \end{aligned}$$

Substituting this back into Bayes' Theorem above, we have

$$Pr^T(AI \mid Flag) = 0.1 \cdot \frac{0.9}{Pr^T(Flag)} = 0.1 \cdot \frac{0.9}{0.18} = \boxed{0.5}.$$

6 | Reflecting on the Course (5 points)

Full points for thoughtful reflection. No "correct answer" here.