CAUSAL THINKING (MATH-352)

Graded homework

General information and guidelines: the finest enumerated item[1] in each question will be marked on a scale of $0-2$ points, indicating an incorrect, partially correct and completely correct answer, respectively (half-points are not given). The deadline is set for **November 18th, 2024**. You are asked to give answers and figures (no code) in a PDF document named `scipernumber_surname_name.pdf` and upload it on Moodle before the deadline.

Please, note that we do not accept homework assignments written by hand. However, DAGs and SWIGs may be drawn by hand and photographed.

Exercise 1 (DAG warm-up)**.** Suppose that we have a data-generating mechanism that follows the DAG $\mathcal{G}$ in Figure 1. We assume the DAG in Figure 1 is faithful to the true law P; in other words, for any disjoint set of nodes A, B, and C, $A \perp\!\!\!\perp B|C$ implies $(A \perp\!\!\!\perp B|C)_{\mathcal{G}}$, or, equivalently, for any disjoint set of nodes A, B, and C, $(A \not\perp\!\!\!\perp B|C)_{\mathcal{G}}$ implies $A \not\perp\!\!\!\perp B|C$.

For the following statements, state if they are true or false. Justify each answer.

Questions:

(1) $X_7 \perp\!\!\!\perp X_2|X_3, X_6$.
(2) $X_1 \not\perp\!\!\!\perp |X_2|X_6$.
(3) $X_7 \perp\!\!\!\perp X_5|X_4$.
(4) $X_4 \perp\!\!\!\perp X_2|X_5$.
(5) $X_1 \perp\!\!\!\perp X_2|X_5$.
(6) $X_6 \perp\!\!\!\perp X_2|X_1, X_3, X_5$.

Solutions:

(1) False, $X_7 \leftarrow X_4 \leftarrow X_1 \rightarrow X_3 \leftarrow X_2$ is open as we condition on X_3.

[1] If a question contains subitems such as (a)-(i)–(iii), then such subitems each receive at most 2 points. If a question does not contain subitems, for example a question with items (a)-(c), then each item receives at most 2 points.

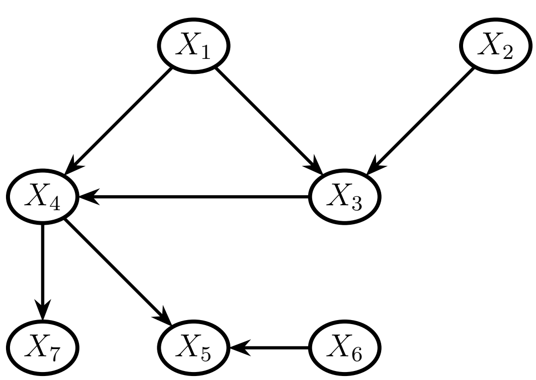

Figure 1

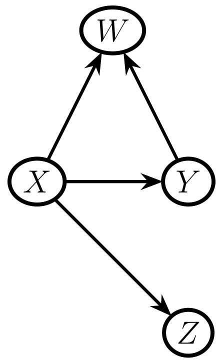

FIGURE 2

(2) False, $X_1 \perp\!\!\!\perp X_2 | X_6$ as all paths from X_1 to X_2 are blocked.
(3) True, all paths from X_7 to X_5 are blocked conditional on X_4.
(4) False, $X_4 \leftarrow X_3 \leftarrow X_2$ is an open path.
(5) False, X_5 is a descendent of collider X_4 on the path $X_1 \rightarrow X_4 \leftarrow X_3 \leftarrow X_2$.
(6) True, all paths from X_6 to X_2 are blocked conditional on X_1, X_3, X_5.

Exercise 2 (SEM warm-up)**.** Let a causal SEM for variables (X, Y, Z, W) and mutually independent, mean zero noise (U_X, U_Y, U_Z, U_W) be given by

$$X = U_X$$
$$Z = -2X + U_Z$$
$$Y = 2X + U_Y$$
$$W = -2X + Y + U_W$$

In this exercise, the causal effect of, say, Z on W is defined as $\mathbb{E}[W^{z=1}] - \mathbb{E}[W^{z=0}]$.

Questions:

(1) Draw the DAG $\mathcal{G}$ representing this data-generating mechanism and the SWIT $\mathcal{G}(x)$.
Tip: you may want to use the 'quiver' website to easily draw, and then export your code in LaTeX `https://q.uiver.app/`.

For the following statements, say if they are true or false and give a justification.

(2) The causal effect of X on W is zero.
(3) The causal effect of Z on Y is strictly negative.
(4) The distribution induced by the SEM is faithful with respect to the DAG.

Solutions:

(1) The DAG is given in Figure 2.
(2) True, as $W = U_Y + U_W$.
(3) False, it is 0, as Z does not directly cause Y.
(4) False, as $W = U_Y + U_W$, hence $W \perp\!\!\!\perp X$ and $W \perp\!\!\!\perp Z$, which is not encoded in the DAG.

Exercise 3. For the following statements, say if they are true or false. Justify your answers.

Questions:

(1) Consider estimating the total causal effect of A on Y in a causal DAG. If A has no parents, then $Y^a \perp\!\!\!\perp A$.
(2) Adjusting for "pre-treatment" variables, i.e. variables describing the subject before treatment, can introduce bias in the estimate of the causal effect of A on Y.

Solutions:

(1) This is true, A is independent of all other variables in the SWIG $\mathcal{G}(a)$, in particular Y^a.
(2) This is true, if the pre-treatment variables are a collider in a backdoor path, adjusting for them can introduce bias.

Exercise 4 (Leveraging the preferences of a physicians)**.** Let $A \in \{0, 1\}$ denote whether a patient takes drug 0 or drug 1. Let $Y \in \{0, 1\}$ be an indicator of patient recovery ($Y = 1$ denotes recovery). A randomized controlled trial reported a 70% recovery rate for each drug. Furthermore, a physician (P_1) found exactly the same recovery rate when reviewing their own patient records.

Questions:

(1) The statement: *"A study found a 70% recovery rate for each drug"* refers to what parameter(s) with respect to our observations of A and Y? Write the statement in formal mathematical notation and justify.

In the remainder of the exercise, we consider two unobserved confounding factors that may affect both treatment and recovery.

The first possible confounder is the patient's socio-economic status (SES), encoded as either low-SES ($S = 0$) or high SES ($S = 1$).

The second unmeasured confounder is the patient's treatment request, which can be influenced by Direct-to-Consumer Advertising (DTCA) in some countries (see `https://en.wikipedia.org/wiki/Direct-to-consumer_advertising` for more details). In particular, a patient may request a treatment ($R = 1$) or not ($R = 0$), which may influence a physician's treatment decision.

(2) Draw the DAG of the data-generating mechanism of (S, R, A, Y). The DAG should include S, R, A and Y only. Justify.

Now, let $D \in \{P_1, P_2\}$ be a random variable indicating whether a patient goes to physician P_1 or a different physician P_2. Suppose that a patient's choice of physician is random and is not directly caused by any other variables.

Physician P_1 makes decisions based on a patients' requests and also their perception of each patient's SES. Specifically, Physician P_1 assigns treatment by the structural equation,

$$A = \begin{cases} 1, & \text{if } (S, R) = (1, 0) \text{ or } (S, R) = (0, 1), \\ 0, & \text{otherwise.} \end{cases}$$

Suppose Physician P_2 is aware of the influences of DTCA, and consciously (though without record) refuses to let patient requests influence the decisions; as such, Physician P_2's treatments can be modeled by the structural equation $A = S$.

Furthermore, suppose $P(R = r, S = s | D = P_i) = P(R = r, S = s) = 0.25$ for all $r \in \{0, 1\}$, $s \in \{0, 1\}$, $i = 1, 2$.

$P(Y=1 \mid A=a', S=s, R=r)$	$(s,r)=(0,0)$	$(s,r)=(0,1)$	$(s,r)=(1,0)$	$(s,r)=(1,1)$
$a'=0$	0.70	0.80	0.60	0.70
$a'=1$	0.90	0.70	0.70	0.50

TABLE 1. True recovery rates, $P(Y=1 \mid A=a', S=s, R=r)$.

	$P(Y^a=1)$	$P(Y=1 \mid A=a, D=P_1)$	$P(Y=1 \mid A=a, D=P_2)$
$a=0$	xx	xx	xx
$a=1$	xx	xx	xx

TABLE 2. FDA experiment, physician P_1 recovery rates, and physician P_2 recovery rates.

	$a'^{P_1}=0$	$a'^{P_1}=1$	$a'^{P_2}=0$	$a'^{P_2}=1$
$a=0$	xx	xx	xx	xx
$a=1$	xx	xx	xx	xx

TABLE 3. Table of $P(Y^a=1 \mid A=a'^{P_i}, D=P_i)$.

The true probabilities of recovery under each confounder state, $P(Y=1 \mid A, S, R)$, are given in Table 1.

(3) Fill in the results of FDA's experimental study and the observations of the two physicians in Table 2.

(4) Imagine that each physician P_i conducts an experiment, randomly assigning treatments $a=0$ and $a=1$, conditional on their intended treatment $A=a'^{P_i}$. Justify that the outcomes of these experiments, $P(Y^a \mid A=a'^{P_i}, D=P_i)$, are identified based on the information you already have and fill in the numbers in Table 3.

Note: The elements of Table 3 are the values of $P(Y^a=1 \mid A=a'^{P_i}, D=P_i)$. The rows represent the values of the different interventions a conducted in the experiment, and the columns represent different values of the physicians' intended treatment A. A superscript, P_i, on a'^{P_i} indicates which physician is conducting the experiment. For example, Table 3's top leftmost entry should be the value of $P(Y^{a=0}=1 \mid A=0, D=P_1)$.

(5) Fill in Table 4 that gives the recovery rates conditional on physicians intent ($A^{D=P_i}$), and explain how you obtained the numbers.

Note: The elements of Table 4 are the values of $P[Y^a=1 \mid A^{D=P_1}, A^{D=P_2}]$. The rows represent the values of the different interventions a, and the columns represent the different possible values of $(A^{D=P_1}, A^{D=P_2})$. For example, Table 4's bottom rightmost entry should be $P(Y^{a=1}=1 \mid A^{D=P_1}=1, A^{D=P_2}=1)$.

Solutions:

(1) The statement refers to $P(Y^{a=1}=1)=P(Y^{a=0}=1)=0.7$. Indeed, an RCT aims to identify the counterfactual distribution.

(2) It is said that S and R are potential ummeasured confounders of the effect of A on Y, hence we draw arrows from S and R to A and Y. We are interested in the effect of A on Y, hence we draw an arrow from A to Y, see Figure 3.

	$(a_1, a_2) = (0, 0)$	$(a_1, a_2) = (0, 1)$	$(a_1, a_2) = (1, 0)$	$(a_1, a_2) = (1, 1)$
$a = 0$	xx	xx	xx	xx
$a = 1$	xx	xx	xx	xx

TABLE 4. Table of expected recovery rates conditional on physicians' intentions, $P[Y^a = 1|A^{D=P_1} = a_1, A^{D=P_2} = a_2]$.

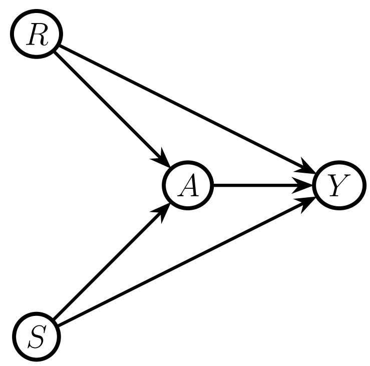

FIGURE 3

	$P(Y^a = 1)$	$P(Y = 1\|A = a, D = P_1)$	$P(Y = 1\|A = a, D = P_2)$
$a = 0$	0.70	0.70	0.75
$a = 1$	0.70	0.70	0.60

TABLE 5. FDA experiment and physician recovery rates.

(3) The results of the FDA experiment are 0.70 as given in the statement. For Physician P_i,

$$
\begin{aligned}
&P(Y = 1|A = a, D = P_i) \\
&= \sum_{r,s\in\{0,1\}} P(Y = 1|A = a, R = r, S = s)P(R = r, S = s|A = a, D = P_i) \\
&= \sum_{r,s\in\{0,1\}} P(Y = 1|A = a, R = r, S = s)\frac{P(A = a|R = r, S = s, D = P_i)P(R = r, S = s)}{\sum_{r',s'\in\{0,1\}} P(A = a|R = r', S = s', D = P_i)P(R = r', S = s')},
\end{aligned}
$$

where the first quality follows from the law of total probability and $Y \perp\!\!\!\perp D|A, R, S$, and the second equality follows from Bayes' law.

Plugging the numbers from Table 1 gives the numbers in Table 5.

	$a'^{P_1}=0$	$a'^{P_1}=1$	$a'^{P_2}=0$	$a'^{P_2}=1$
$a=0$	0.70	0.70	0.75	0.65
$a=1$	0.70	0.70	0.80	0.60

TABLE 6. Result of $P(Y^a = 1|A = a'^{P_i}, D = P_i)$.

	$(a_1,a_2)=(0,0)$	$(a_1,a_2)=(0,1)$	$(a_1,a_2)=(1,0)$	$(a_1,a_2)=(1,1)$
$x=0$	0.70	0.70*	0.80*	0.60
$x=1$	0.90*	0.50	0.70	0.70*

TABLE 7

(4) The expected recovery of P_i's patients is

$$\begin{aligned}
&\mathbb{E}[Y^a|A=a', D=P_i] \\
&= \sum_{r,s} \mathbb{E}[Y^a|A=a', D=P_i, R=r, S=s] P(R=r, S=s|A=a', D=P_i) \\
&= \sum_{r,s} \mathbb{E}[Y^a|A=a', R=r, S=s] \frac{P(A=a'|R=r, S=s, D=P_i)P(R=r, S=s)}{\sum_{r',s'\in\{0,1\}} P(A=a'|R=r', S=s', D=P_i)P(R=r', S=s')} \\
&= \sum_{r,s} \mathbb{E}[Y^a|A=a, R=r, S=s] \frac{P(A=a'|R=r, S=s, D=P_i)P(R=r, S=s)}{\sum_{r',s'\in\{0,1\}} P(A=a'|R=r', S=s', D=P_i)P(R=r', S=s')} \\
&= \sum_{r,s} \mathbb{E}[Y|A=a, R=r, S=s] \frac{P(A=a'|R=r, S=s, D=P_i)P(R=r, S=s)}{\sum_{r',s'\in\{0,1\}} P(A=a'|R=r', S=s', D=P_i)P(R=r', S=s')},
\end{aligned}$$

where the first equality follows from the law of total probability, the second equality follows from Bayes' theorem and $Y^a \perp\!\!\!\perp D|A, R, S$, the third equality follows from $Y^a \perp\!\!\!\perp A|R, S$, and the last follows from consistency, see Table 6.

(5)

$$\begin{aligned}
\mathbb{E}[Y^a|A^{D=P_1}, A^{D=P_2}] &= \mathbb{E}[Y^a|\text{XOR}(S,R), S] \\
&= \mathbb{E}[Y^a|S, R] \\
&= \mathbb{E}[Y|A, S, R],
\end{aligned}$$

where the first equality follows from the statement, the second equality follows from the fact that we know $(\text{XOR}(S,R), S)$[2] if and only if we know (S, R), and the last equality follows from $Y^a \perp\!\!\!\perp A|S, R$, as can be seen e.g. using the backdoor criterion in the DAG.

Hence, using Table 1, we can fill Table 4, see Table 7.

Exercise 5 (A spinning curiosity). Let Y_1 and Y_2 be two outcomes of interest in $\{-1, 1\}$ that do not cause each other, and let X_1 and X_2 be two randomly assigned treatments taking values in $\{0, 1, 2\}$ that can directly affect Y_1 and Y_2. Suppose furthermore, that X_1 and X_2 are randomly assigned independently of each other, that is, $X_1 \perp\!\!\!\perp X_2$. Let $M = I(Y_1 = Y_2)$. Suppose that $\mathbb{E}[M^{x_1,x_2}]$ is known for all x_1, x_2.

[2]Another name for the function describing Physician P_1's treatment assignment is 'XOR'.

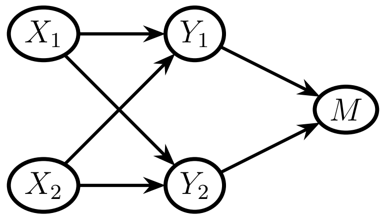

FIGURE 4

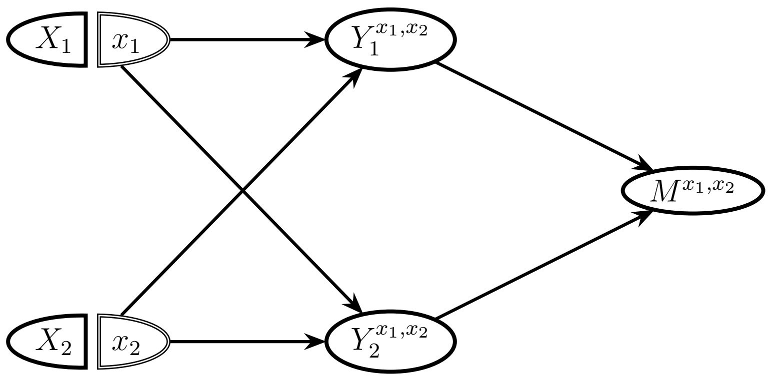

FIGURE 5

We will be interested in the following assumptions:

(A) $Y_1^{x_1,x_2} = Y_1^{x_1}$,
(B) $Y_2^{x_1,x_2} = Y_2^{x_2}$.

Questions:

(1) Without assuming Assumptions (A) and (B), draw the DAG $\mathcal{G}$ representing the data-generating mechanism above. $\mathcal{G}$ should include X_1, X_2, Y_1, Y_2, and M.
(2) Draw the SWIT corresponding to interventions on X_1 and X_2, $\mathcal{G}(x_1, x_2)$.
(3) Show that if for one $i \in \{1, 2\}$, $Y_i^{x_1,x_2} = s$ for some (x_1, x_2) and $Y_i^{x_1',x_2'} = -s$ for all other (x_1', x_2'), then Assumptions (A) and (B) cannot both hold.
(4) Suppose that for $x_1 \in \{0, 1\}, x_2 \in \{0, 2\}$, we have that for some unit $M^{x_1,x_2} = 1$ if and only if $(x_1, x_2) = (1, 2)$. Then, show that if $Y_1^{x_1,x_2}$ does not verify that $Y_1^{x_1',x_2'} = -Y_1^{x_1,x_2}$ for all other x_1', x_2', then $Y_2^{x_1,x_2}$ verifies that $Y_2^{x_1',x_2'} = -Y_2^{x_1,x_2}$ for all other x_1', x_2'.
(5) Show that if $\mathbb{E}[M^{1,2}] - \mathbb{E}[M^{0,2}] - \mathbb{E}[M^{1,0}] - \mathbb{E}[M^{0,0}] > 0$, there must exist a unit with $M(1, 2) = 1, M(0, 2) = M(1, 0) = M(0, 0) = 0$.
(6) Deduce that if $\mathbb{E}[M^{1,2}] - \mathbb{E}[M^{0,2}] - \mathbb{E}[M^{1,0}] - \mathbb{E}[M^{0,0}] > 0$, Assumptions (A) and (B) cannot both hold.

Solutions:

(1) See Figure 4.
(2) See Figure 5.
(3) Suppose without loss of generality that $i = 1$. Take $x_2' \in \{0, 1, 2\}, x_2' \neq x_2$. Then, $Y_1^{x_1,x_2'} \neq Y_1^{x_1,x_2}$ and hence Assumption (A) is violated.

(4) $M(1,2) = 1$ implies that $Y(1,2) = Y(2,1) = s$ for some $s \in \{-1, 1\}$. We show that if $Y_1^{x_1,x_2}$ does not satisfy the hypothesis of the previous question, then $Y_2^{x_1,x_2}$ must.
If $Y_1^{x_1,x_2}$ does not satisfy the hypothesis of the previous question, then either:
   (a) The other three values of $Y_1^{x_1,x_2}$ are also s.
   Then, for $x_1 \in \{0,1\}$ and $x_2 \in \{0,2\}$, as $M^{x_1,x_2} = 0$ for $(x_1, x_2) \neq (1,2)$, we have $Y_2^{x_1,x_2} = -s$ for $(x_1, x_2) \neq (1,2)$, and thus $Y_2^{x_1,x_2}$ satisfies the hypothesis of the previous question.
   (b) Exactly one other value of $Y_1^{x_1,x_2}$ is equal to s and the other two are $-s$.
   Then, because $M^{x_1,x_2} = 0$ for $(x_1, x_2) \neq (1,2)$, we have $Y_2^{x_1,x_2} = -s$ for exactly one value of $(x_1, x_2, x_1 \in \{0,1\}, x_2 \in \{0,2\}$, in which case, $Y_2^{x_1,x_2}$ would satisfy the hypothesis of the previous question.

(5) We argue by contradiction. Suppose that there were no units with $M(1,2) = 1, M(0,2) = M(1,0) = M(0,0) = 0$. Then, for all units, $M^{1,2} - M^{0,2} - M^{1,0} - M^{0,0} \leq 0$, which implies $\mathbb{E}[M^{1,2}] - \mathbb{E}[M^{0,2}] - \mathbb{E}[M^{1,0}] - \mathbb{E}[M^{0,0}] \leq 0$, a contradiction.

(6) This follows immediately from the previous questions.

Exercise 6 (Monty's New Paradox)**.** You are a participant in Monty Hall's new game show. In front of you, there are two doors, Door A and Door B. Monty Hall opens Door A. Behind Door A, you see a check for 1'000 CHF. You cannot see the contents of Door B. Monty tells you the following:

> *"Behind Door A there is a 1'000 CHF check. Door B contains a check of either 0 or 1'000'000 CHF. You can chose to only open Door B or both doors. The content of Door B was chosen by our MontyAI program. Based on the personality test you have taken before coming on the show, MontyAI has put 1'000'000 CHF in Door B if it predicts that you will only open Door B, but has left it empty if it predicts that you will open both doors. As you know, MontyAI has a* 99% *success rate; it can correctly predict if you will open Door B or both doors* 99% *of times."*

We call opening Door B *one-picking* and opening both doors *two-picking.*
Questions:

(1) Give a representation in DAG form of the situation. The following variables must appear: a choice variable A representing the option you choose ($A = 1$ if you open both doors and $A = 0$ if you open Door B), an outcome variable Y representing the money you receive at the end of the game, a variable L representing the result of the personality test you took before the game, and a variable D representing MontyAI's prediction ($D = 1$ if you pick both boxes and $D = 0$ if you pick Box B).

(2) (a) What is the expected factual reward of one-pickers?
   The expected factual reward is the expected reward if we were to observe a large amount of one-pickers playing Monty's new game. The expected factual reward is a parameter of the factual law (L, A, D, Y), there should not be any counterfactuals in your estimand.
   (b) What is the expected factual reward of two-pickers?

We call the action (one-picking or two-picking) with the highest expected factual reward the optimal factual action.

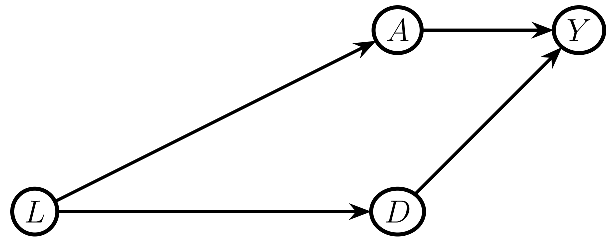

FIGURE 6. DAG representing Monty's new game show.

(3) In the situation described above, where MontyAI has already picked the amount behind Door B, is your counterfactual expected reward higher if you one-pick or two-pick?

The expected counterfactual reward is the expected reward under different interventions on A. The expected counterfactual reward is a parameter of the counterfactual law (L, A, D, Y^a). We call the action with the highest expected counterfactual outcome the optimal counterfactual action.

(4) Do the optimal factual and counterfactual actions change if MontyAI has a 100% success rate? What about 50%?

(5) For which accuracy levels p of MontyAI are the optimal factual and counterfactual actions equal?

(6) If Monty adjusts the cash rewards behind the doors, keeping a strictly positive reward behind Door B, can he ensure that the optimal factual expected reward will be different from the optimal counterfactual expected reward any level of accuracy p of MontyAI?

Solutions:

(1) A influences your outcome Y, and is influenced by your personal characteristics L. D influences your outcome Y and is determined by your characteristics L, see Figure 6 for the corresponding DAG.

(2) The expected reward of one-picking is $\mathbb{E}[Y|A = 0]$. By the law of total expectation and the problem statement, it can be calculated as follows:

$$\begin{aligned}\mathbb{E}[Y|A = 0] &= \mathbb{E}[Y|A = 0, D = 0]P(D = 0|A = 0)\\ &\quad + \mathbb{E}[Y|A = 0, D = 1]P(D = 1|A = 0)\\ &= 1'000'000 \times 0.99 + 0 \times 0.01\\ &= 990'000.\end{aligned}$$

Similarly, the expected reward of two-picking is $\mathbb{E}[Y|A = 1]$, and can be calculated as follows:

$$\begin{aligned}\mathbb{E}[Y|A = 1] &= \mathbb{E}[Y|A = 1, D = 0]P(D = 0|A = 1)\\ &\quad + \mathbb{E}[Y|A = 1, D = 1]P(D = 1|A = 1)\\ &= 1'001'000 \times 0.01 + 1'000 \times 0.99\\ &= 10'010 + 990 = 11'000 < 990'000\end{aligned}$$

Hence, one-pickers receive higher rewards on average than two-pickers.

(3) The expected reward of choosing action $a \in \{0,1\}$ is $\mathbb{E}[Y^a]$. However,

$$\begin{aligned}\mathbb{E}[Y^1] &= \mathbb{E}[\mathbb{E}[Y^1|D]]\\ &= \mathbb{E}[\mathbb{E}[Y^0 + 1'000|D]]\\ &= \mathbb{E}[\mathbb{E}[Y^0|D] + 1'000]\\ &= \mathbb{E}[\mathbb{E}[Y^0|D]] + 1'000\\ &= \mathbb{E}[Y^0] + 1'000 > \mathbb{E}[Y^0].\end{aligned}$$

Hence, our expected reward is always higher if we two-pick.

(4) The reasoning given in the solution to the previous question does not depend on the accuracy of Monty's AI, hence the optimal counterfactual action does not change even if the AI has 100% accuracy. It is easy to see that the optimal factual action also remains the same.

If the AI has 50% accuracy, both approaches show it is better to two-pick.

(5) Replacing the probabilities of rewards by p and $1-p$, we get that the value p at which point it the optimal factual and counterfactual actions are equal is $p = 50.05\%$.

Indeed, the optimal counterfactual action is always to two-pick, and

$$\begin{aligned}\mathbb{E}[Y|A=1] - \mathbb{E}[Y|A=0] &= 1'001'000(1-p) + 1'000p - 1'000'000p\\ &= 1'001'000 - 2'000'000p,\end{aligned}$$

which is equal to zero if and only if $p = 0.5005$.

(6) If $p \neq 50\%$, we can change the cash amounts so that the optimal factual and counterfactual actions are different.

The optimal counterfactual action will always be to two-pick if the reward behind Door B is strictly positive. However, for $p > 50\%$, if we keep the reward behind Door B to be 1'000 CHF (the actual number does not matter as long as it is positive), and let x be the reward behind Door A, then

$$\begin{aligned}\mathbb{E}[Y|A=1] - \mathbb{E}[Y|A=0] &= (x + 1'000)(1-p) + 1'000p - xp\\ &= x + 1'000 - 2xp,\end{aligned}$$

which is smaller than 0 if and only if $x > \frac{1'000}{2p-1}$.

If $p < 50\%$, we can simply do the opposite of the recommendation of MontyAI. However if $p = 50\%$, $\mathbb{E}[Y|A=1] - \mathbb{E}[Y|A=0]$ will always be equal to the reward behind Door B, and hence the optimal factual regime will always be equal to the optimal counterfactual regime.

References