PE3101P: Decision Theory and Social Choice

In-class Test Practice Paper

Time allowed: 90 minutes. Total: 100 marks.

- Section A: answer all ten multiple-choice questions. Each is worth 4 marks (40 marks in total).
- Section B: Answer at least three of the four questions. Your best three marks will be counted. Each is worth 20 marks (60 marks in total).

Section A

Answer all 10 questions. Each is worth 4 marks. Choose one answer, A–D, for each question.

1. An agent considers exactly two mutually exclusive hypotheses, H_1 and H_2. Her prior odds for H_1 over H_2 are 1:3, and the Bayes factor for H_1 over H_2 supplied by evidence E is 2. She conditionalises on E. Which statement is correct?

A. The evidence favours H_2; after updating, H_1 is more probable than H_2.
B. The evidence favours H_1; after updating, H_2 is more probable than H_1.
C. The evidence favours H_2; after updating, H_2 is more probable than H_1.
D. The evidence favours H_1; after updating, H_1 is more probable than H_2.

2. An agent has credence 0.2 in A. A bet pays \$20 if A occurs and \$0 otherwise. Under the betting interpretation of credences, which offer is she guaranteed to accept?

A. Sell the bet for \$6. **B.** Sell the bet for \$2.

C. Buy the bet for \$4. **D.** Buy the bet for \$6.

3. The table gives an agent's net monetary gains from three bets, including all prices and payouts. The three states are mutually exclusive and jointly exhaustive. If she accepts all three bets, how much money is she guaranteed to lose?

Net gain (\$)	P	Q	Neither
Bet 1	-18	2	2
Bet 2	2	-18	2
Bet 3	13	13	-7

A. \$7 **B.** \$3

C. \$20 **D.** \$4

4. The table lists dependency hypotheses, their credences, and the utilities each act would produce. What is the agent's credence that A_1 would produce strictly greater utility than A_2?

	K_1	K_2	K_3
Credence	0.2	0.3	0.5
A_1	6	0	8
A_2	4	2	6

A. 0.2 **B.** $\frac{2}{3}$

C. 0.7 **D.** 0.3

Questions 5–6

The agent evaluates prospects by expected utility. She knows the tree and that her preferences will change from the D_1 row to the D_2 row.

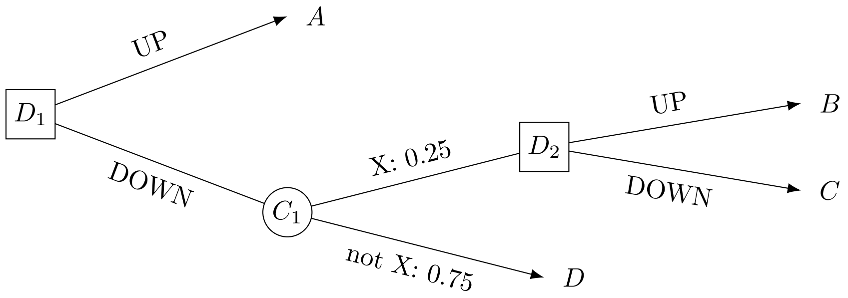

Utility at	A	B	C	D
D_1	2	12	2	0
D_2	2	12	14	0

5. What will an agent following sophisticated choice do at D_1, and why?

A. DOWN: $B[0.25]D$ is better than A using the D_1 row.
B. DOWN: $C[0.25]D$ is better than A using the D_2 row.
C. UP: $C[0.25]D$ is worse than A using the D_1 row.
D. UP: $B[0.25]D$ is worse than A using the D_1 row.

6. Suppose the DOWN action at D_2, giving C, is removed. What will an agent following sophisticated choice do at D_1, and why?

A. DOWN: $C[0.25]D$ is better than A using the D_2 row.
B. UP: $B[0.25]D$ is worse than A using the D_1 row.
C. DOWN: $B[0.25]D$ is better than A using the D_1 row.
D. UP: $C[0.25]D$ is worse than A using the D_1 row.

7. An agent's credences and conditional expectations are

$$\mathrm{Cr}(B) = 0.2, \quad \mathrm{Cr}(\neg B) = 0.8, \quad \mathbb{E}[X \mid B] = 10, \quad \mathbb{E}[X \mid \neg B] = 4.$$

What is $\mathbb{E}[X]$?

A. 2.6 **B.** 5.2
C. 8.8 **D.** 7

8. For distinct certain outcomes A, B, let

$$P = A[0.4]B, \qquad Q = A[0.6]B.$$

What is the probability of A in $P[0.25]Q$?

A. 0.1 **B.** 0.7
C. 0.55 **D.** 0.5

9. Outcomes a, b, c have scores $2, 1, 0$. A prospect is ranked by the highest score among outcomes it gives positive probability; equal scores give indifference.
Here $a \succ b \succ c$, so Continuity requires $b \sim a[p]c$ for some $0 < p < 1$. Why is there no such p?

A. For every $0 < p < 1$, $a \succ a[p]c$.

B. At $p = 0.5$, $a[p]c \sim b$.

C. For every $0 < p < 0.5$, $b \succ a[p]c$.

D. For every $0 < p < 1$, $a[p]c \succ b$.

10. At a later node, an agent can choose B or A^-. The opportunity to keep A for free is no longer available. According to Decision-Tree Separability, what determines the rational choice now?

A. The part of the decision tree still reachable from this node.

B. The plan that would have preserved the initial holding.

C. The number of earlier opportunities the agent declined.

D. The best outcome that was available at any earlier node.

Section B

Question 1 [20 marks]

In this question, you may use without proof any general facts about probabilities proved in the lectures.

An agent's credences satisfy the probability axioms. Events A and B satisfy

$$\mathrm{Cr}(A) = 0.6, \qquad \mathrm{Cr}(B) = 0.5, \qquad \mathrm{Cr}(A \cap B) = x.$$

(a) (i) Calculate $\mathrm{Cr}(A^c)$. [2]

(ii) Calculate $\mathrm{Cr}(B^c)$. [2]

(iii) Calculate the maximum and minimum possible values of x. [4]

(b) (i) Express $\mathrm{Cr}(A \mid B)$ and $\mathrm{Cr}(B \mid A)$ in terms of x. [4]

(ii) Determine the value of x if A and B are probabilistically independent. Show your working. [2]

(c) Now take $x = 0.2$. The agent learns that $A \cup B$ occurred, and learns nothing further. They update by conditionalisation.

(i) Calculate their new credences in A, B and $A \cap B$. [4]

(ii) Are A and B probabilistically independent after this update? Justify your answer using the new credences. [2]

Question 2 [20 marks]

(a) In a decision table, the states are mutually exclusive and jointly exhaustive, and each state together with an act fixes the resulting outcome. To use the *unconditional* state credences as weights, what independence condition on the states and acts is required by CDT, and what independence condition is required by EDT? How can EDT evaluate the acts if the latter condition fails? [4]

(b) A separate decision problem has three available acts and four possible outcomes. How many dependency hypotheses are there? Assume there are no restrictions on which outcome would result from each act. Explain your count. [2]

For the remaining parts, consider a hypothetical smoking-lesion case. A gene influences both whether an agent gets cancer and their inclination to smoke. Smoking itself has *no causal effect* on cancer. It gives a benefit of 2 utility units, whether or not cancer occurs; cancer has a cost of 10 units. These are the only consequences the agent cares about.

Write S for smoking, N for not smoking, and C for cancer. The agent must choose exactly one of S and N. Initially,

$$\mathrm{Cr}(C) = 0.5, \qquad \mathrm{Cr}(C \mid S) = 0.7, \qquad \mathrm{Cr}(C \mid N) = 0.3.$$

	C	C^c
S	-8	2
N	-10	0

(c) Calculate the causal and evidential expected utilities of S and N, and state each theory's recommendation. [4]

Before choosing, the agent receives a test result E about the gene. The test merely supplies information: it does not affect the gene, cancer, or the benefit of smoking. Write Cr_E for the agent's updated credence function after learning E. The new credences are

$$\mathrm{Cr}_E(C) = 0.9, \qquad \mathrm{Cr}_E(C \mid S) = 0.9, \qquad \mathrm{Cr}_E(C \mid N) = 0.9.$$

(d) Calculate the new causal and evidential expected utilities of both acts, and state each theory's recommendation. [4]

(e) Explain why learning E can change the causal expected utilities even though neither the causal effects nor the utilities of the outcomes have changed. [2]

(f) A second agent's preferences are represented by the same utility table. They have the same causal information, but after receiving their own test result have current credences

$$\mathrm{Cr}_2(C) = 0.9, \qquad \mathrm{Cr}_2(C \mid S) = 0.95, \qquad \mathrm{Cr}_2(C \mid N) = 0.85.$$

Deduce $\mathrm{Cr}_2(S)$ and $\mathrm{Cr}_2(N)$. [4]

Question 3 [20 marks]

An agent has a preference relation $\succsim$ over all prospects on the finite set of outcomes $x_1, \ldots, x_n$. Let u assign a real number to each outcome. For a prospect P, write $P(x_i)$ for its probability of outcome x_i, and define

$$\mathrm{EU}_u(P) = \sum_{i=1}^{n} P(x_i)u(x_i).$$

The mixture $P[p]Q$ gives P with probability p and Q otherwise, so its probability of x_i is $pP(x_i) + (1-p)Q(x_i)$.

(a) State what it means for u to represent the agent's preferences by expected utility. Use EU_u in your answer.

[2]

For all remaining parts, assume that u represents the agent's preferences. You may assume any relevant facts about the real numbers without proof. In the remaining questions, you will show that when the agent's preferences are representable in this way, they must satisfy the four axioms of expected utility theory.

(b) Let $v(x_i) = \alpha u(x_i) + \beta$ for every outcome, where $\alpha > 0$ and $\beta \in \mathbb{R}$. You may use without proof the identity

$$\mathrm{EU}_v(P) = \alpha\,\mathrm{EU}_u(P) + \beta.$$

Show that v also represents the agent's preferences.

[4]

(c) **Transitivity.** Suppose $P \succsim Q$ and $Q \succsim R$. Show that $P \succsim R$.

[2]

(d) **Completeness.** Show that, for any prospects P, Q, either $P \succsim Q$ or $Q \succsim P$ (or both).

[2]

(e) **Independence.** You may use without proof the identity

$$\mathrm{EU}_u(P[p]R) = p\,\mathrm{EU}_u(P) + (1-p)\,\mathrm{EU}_u(R).$$

Show that, for $0 < p < 1$,

$$P \succsim Q \quad \text{if and only if} \quad P[p]R \succsim Q[p]R.$$

[5]

(f) **Continuity.** Suppose $P \succ Q \succ R$. Show that there is a $p \in (0,1)$ such that $Q \sim P[p]R$.

Hint: equate the expected utilities and then solve for p. Show that your value lies strictly between 0 and 1, and explain why it gives the required indifference.

[5]

Question 4 [20 marks]

The agent faces the decision tree below. Cash transfers are shown on the branches. Only the actions shown are available.

The agent knows the whole tree and the following utility change from the start. They learn nothing about E before D_2; its probability and their credence in it remain $\frac{1}{2}$. Here N^- is N with a net cash loss of \$1. The table gives the complete final utilities, including all cash transfers; each dollar lost reduces utility by one unit. Only the utility of X changes:

	X	N	N^-
At D_1	8	0	-1
At D_2	-4	0	-1

The agent evaluates prospects by expected utility.

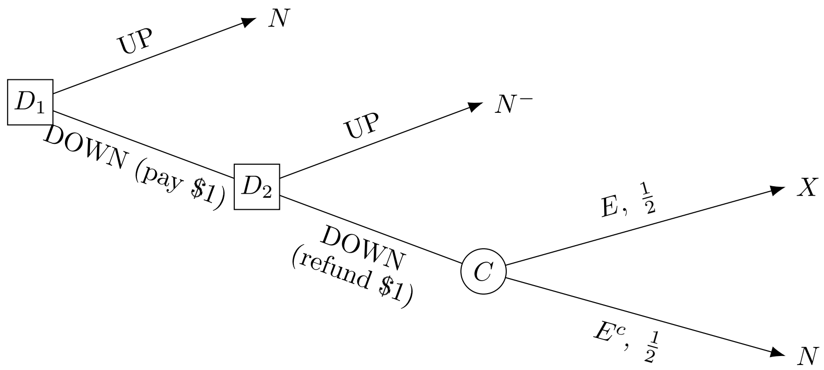

Naive choice selects a plan using current preferences and chooses afresh later. Resolute choice follows the plan selected initially. Sophisticated choice predicts later choices and evaluates their consequences using current preferences.

(a) What does a naive agent initially plan to do at D_1 and at D_2 if reached? Justify the plan using the initial expected utilities. [2]

(b) What will the naive agent actually choose at D_2 if reached? Show the comparison using the utilities at D_2. Determine their final outcome and net cash change. Are they guaranteed to be worse off than choosing UP at D_1? Justify your answer. [4]

(c) What would a resolute agent do? Compare their final result with choosing UP at D_1, first if E occurs and then if E^c occurs. Use the D_1 utilities. Is the resolute agent guaranteed to be worse off? [2]

(d) (i) What would a sophisticated agent choose at D_1? Justify your answer using their prediction of the later choice. [2]

(ii) Could changing only the cost of DOWN at D_1 induce a sophisticated agent to choose DOWN at D_1 and then UP at D_2 at a loss? Explain. [2]

Question 4 (continued)

(e) Consider a separate agent and the separate tree below. Their preferences remain fixed:

$$A \succ B, \qquad B \succ C, \qquad C \succ A.$$

When selecting a plan from a menu with all three outcomes available, the agent selects a plan leading to A. With only two distinct outcomes available, they select according to the stated pairwise preference.

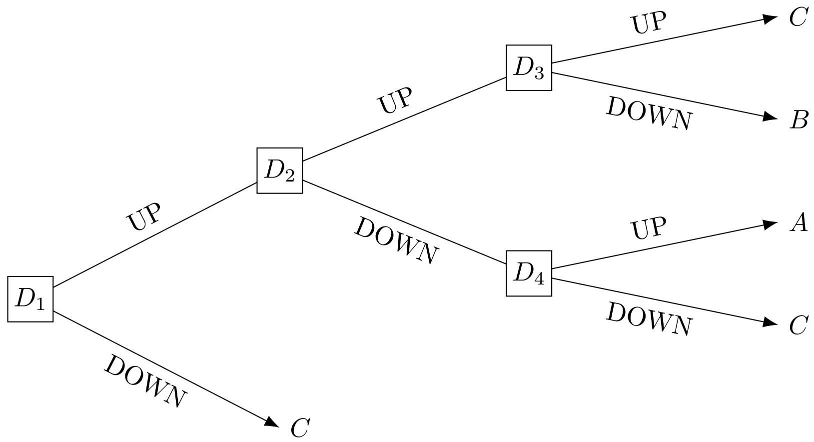

For each procedure below, state the sequence of UP/DOWN choices the agent makes, naming the decision nodes visited, and give the final outcome. You need not specify choices at nodes the agent does not reach.

(i) Naive choice. [3]

(ii) Sophisticated choice. [3]

(iii) Resolute choice. [2]