PE3101P: Decision Theory and Social Choice

In-class Test Practice Paper — Solutions

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.

Solution. **B.** A Bayes factor greater than 1 favours H_1 over H_2. But posterior odds equal prior odds multiplied by the Bayes factor:

$$\frac{\mathrm{Cr}(H_1 \mid E)}{\mathrm{Cr}(H_2 \mid E)} = \frac{1}{3} \times 2 = \frac{2}{3}.$$

The posterior odds are $2:3$, so H_2 remains more probable. Evidence favouring H_1 need not make it the more probable hypothesis.

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.

Solution. **A.** The fair price is stake times credence: $\$20 \times 0.2 = \$4$. Selling for \$6 is above the fair price, so acceptance is guaranteed by the betting interpretation. Selling for \$2 or buying for \$6 is not favourable by that rule. Buying at exactly \$4 is at the fair price, where acceptance is not guaranteed.

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

Solution. **B: $3.** Add the three net gains within each state:

$$P: -18 + 2 + 13 = -16 + 13 = -3,$$
$$Q: 2 - 18 + 13 = -16 + 13 = -3,$$
$$\text{Neither}: 2 + 2 - 7 = 4 - 7 = -3.$$

The states exhaust the possibilities, and each gives a net loss of $3.

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

Solution. **C:** 0.7**.** Under K_1, act A_1 gives $6 > 4$; under K_2, it gives $0 < 2$; under K_3, it gives $8 > 6$. Strictly greater utility therefore occurs exactly under K_1 or K_3. Adding their credences gives $0.2 + 0.5 = 0.7$.

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.

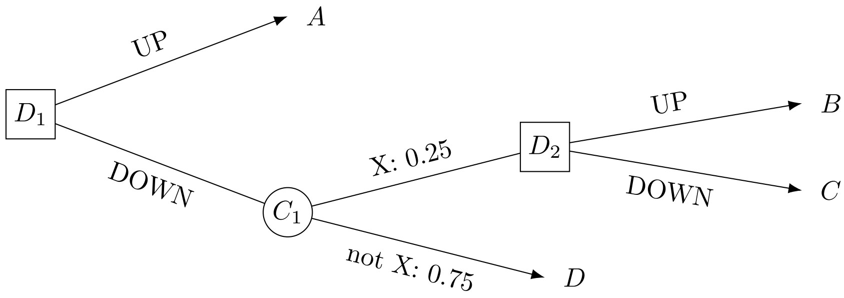

Utility at	A	B	C	D
D_1	2	12	2	0
D_2	2	12	14	0

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.

Solution. **C.** At D_2, the agent chooses DOWN, giving C, because $14 > 12$ using the D_2 row. Sophisticated choice predicts this. At D_1, DOWN therefore gives the prospect $C[0.25]D$, evaluated using the D_1 row:

$$\text{EU}_1(C[0.25]D) = 0.25(2) + 0.75(0) = 0.5.$$

UP gives A, with utility 2. Since $2 > 0.5$, she chooses UP.

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.

Solution. **C.** Removing C leaves B as the only outcome available at D_2. Thus DOWN at D_1 now leads to $B[0.25]D$. Its initial expected utility is

$$\text{EU}_1(B[0.25]D) = 0.25(12) + 0.75(0) = 3.$$

UP still gives A with utility 2. Since $3 > 2$, she chooses DOWN.

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

Solution. **B: 5.2.** By the law of total expectation,

$$\begin{aligned}\mathbb{E}[X] &= \mathrm{Cr}(B)\mathbb{E}[X \mid B] + \mathrm{Cr}(\neg B)\mathbb{E}[X \mid \neg B] \\ &= 0.2(10) + 0.8(4) \\ &= 2 + 3.2 = 5.2.\end{aligned}$$

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

Solution. **C: 0.55.** The outer mixture chooses P with probability 0.25 and Q with probability 0.75. Outcome A has probability 0.4 under P and 0.6 under Q, so its total probability is

$$0.25(0.4) + 0.75(0.6) = 0.10 + 0.45 = 0.55.$$

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

Solution. **D.** For every $0 < p < 1$, the prospect $a[p]c$ gives positive probability to a, whose score is 2. Its highest possible score is therefore 2, whereas certain b has score 1. Hence $a[p]c \succ b$ for every such p, so no interior mixture is indifferent to 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.

Solution. **A.** Decision-Tree Separability requires that the agent evaluates the remaining decision tree as though the tree had started at the current node. Only B and A^- remain available; keeping A for free is no longer an option.

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]

Solution. The complement rule gives $\mathrm{Cr}(A^c) = 1 - \mathrm{Cr}(A) = 1 - 0.6 = 0.4$.

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

Solution. Likewise, $\mathrm{Cr}(B^c) = 1 - \mathrm{Cr}(B) = 1 - 0.5 = 0.5$.

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

Solution. Since $A \cap B \subset A$ and $A \cap B \subset B$, monotonicity gives $x \leq 0.6$ and $x \leq 0.5$. Thus $x \leq 0.5$.

By inclusion–exclusion,

$$\mathrm{Cr}(A \cup B) = \mathrm{Cr}(A) + \mathrm{Cr}(B) - \mathrm{Cr}(A \cap B) = 0.6 + 0.5 - x = 1.1 - x.$$

A credence cannot exceed 1, so $1.1 - x \leq 1$. Adding x to both sides gives $1.1 \leq 1 + x$. Subtracting 1 from both sides gives $0.1 \leq x$.

Both endpoints are possible. The four disjoint regions have credences

	$A \cap B$	$A \setminus B$	$B \setminus A$	$(A \cup B)^c$
General	x	$0.6 - x$	$0.5 - x$	$x - 0.1$
$x = 0.1$	0.1	0.5	0.4	0
$x = 0.5$	0.5	0.1	0	0.4

Each endpoint row is non-negative, sums to 1, and gives the required marginals. Therefore the **minimum is** 0.1 **and the maximum is** 0.5.

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

Solution. By the definition of conditional credence,

$$\mathrm{Cr}(A \mid B) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(B)} = \frac{x}{0.5} = 2x,$$

and

$$\mathrm{Cr}(B \mid A) = \frac{\mathrm{Cr}(B \cap A)}{\mathrm{Cr}(A)} = \frac{x}{0.6} = \frac{5x}{3}.$$

The intersections are the same event, but the denominators differ.

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

Solution. Independence requires

$$x = \mathrm{Cr}(A \cap B) = \mathrm{Cr}(A)\,\mathrm{Cr}(B) = 0.6(0.5) = 0.3.$$

Alternatively, using the conditional-probability definition, independence requires $\mathrm{Cr}(A \mid B) = \mathrm{Cr}(A)$. Substituting part (b)(i) gives $2x = 0.6$, so division by 2 again gives $\boxed{x = 0.3}$.

(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]

Solution. Write $U = A \cup B$. With $x = 0.2$, inclusion–exclusion gives

$$\mathrm{Cr}(U) = 0.6 + 0.5 - 0.2 = 0.9.$$

Conditionalisation sets $\mathrm{Cr}_U(X) = \mathrm{Cr}(X \mid U) = \mathrm{Cr}(X \cap U)/\,\mathrm{Cr}(U)$. Each of A, B and $A \cap B$ is a subset of U, so intersecting it with U leaves it unchanged. Thus

$$\begin{aligned} \mathrm{Cr}_U(A) &= \frac{0.6}{0.9} = \frac{6}{9} = \frac{2}{3}, \\ \mathrm{Cr}_U(B) &= \frac{0.5}{0.9} = \frac{5}{9}, \\ \mathrm{Cr}_U(A \cap B) &= \frac{0.2}{0.9} = \frac{2}{9}. \end{aligned}$$

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

Solution. **No.** Independence would require the new intersection credence to equal the product of the new marginal credences. But

$$\mathrm{Cr}_U(A)\,\mathrm{Cr}_U(B) = \frac{2}{3} \cdot \frac{5}{9} = \frac{10}{27}, \qquad \mathrm{Cr}_U(A \cap B) = \frac{2}{9} = \frac{6}{27}.$$

Since $10/27 \neq 6/27$, independence fails after the update.

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]

Solution. For **causal decision theory**, the states must be *causally independent* of the choice of act: choosing a different act would not change which state obtains. This lets the unconditional state credences serve as the causal weights.

For **evidential decision theory**, the states must be *probabilistically independent* of the acts: for each state s and act a, $\mathrm{Cr}(s \mid a) = \mathrm{Cr}(s)$ whenever the conditional is defined. Only then can its conditional state weights be replaced by unconditional ones.

If probabilistic independence fails, evidential decision theory can still use the table by weighting each act's utilities with the state credences *conditional on that act*:

$$\mathrm{EU}_e(a) = \sum_s \mathrm{Cr}(s \mid a)u(\text{outcome of } a \text{ in } s).$$

(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]

Solution. $4^3 = 64$**.** A dependency hypothesis specifies what would result from *each* act. There are 4 choices for the first act's outcome, independently 4 for the second, and 4 for the third. Thus there are $4 \times 4 \times 4 = 64$ complete assignments.

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]

Solution. Smoking has no causal effect on cancer, so causal decision theory uses the same unconditional cancer credence 0.5 for both acts:

$$\mathrm{EU}_c(S) = 0.5(-8) + 0.5(2) = -4 + 1 = -3,$$
$$\mathrm{EU}_c(N) = 0.5(-10) + 0.5(0) = -5 + 0 = -5.$$

Since $-3 > -5$, **causal decision theory recommends smoking**.

Evidential decision theory uses the cancer credence conditional on each act. Conditional credences satisfy the probability axioms when the unconditional credences do, so the complement rule also applies to them. The complementary credences are therefore $1 - 0.7 = 0.3$ given S and $1 - 0.3 = 0.7$ given N. Hence

$$\mathrm{EU}_e(S) = 0.7(-8) + 0.3(2) = -5.6 + 0.6 = -5,$$
$$\mathrm{EU}_e(N) = 0.3(-10) + 0.7(0) = -3 + 0 = -3.$$

Since $-3 > -5$, **evidential decision theory recommends not smoking**.

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]

Solution. The updated cancer credence is 0.9 both unconditionally and conditional on either act; its complement is 0.1. Both theories therefore use the same weights:

$$\begin{aligned}\mathrm{EU}_{c,E}(S) = \mathrm{EU}_{e,E}(S) &= 0.9(-8) + 0.1(2) = -7.2 + 0.2 = -7,\\ \mathrm{EU}_{c,E}(N) = \mathrm{EU}_{e,E}(N) &= 0.9(-10) + 0.1(0) = -9 + 0 = -9.\end{aligned}$$

Since $-7 > -9$, **both theories recommend smoking**.

(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]

Solution. Causal expected utility depends on the agent's credences as well as on the causal consequences and their utilities. The test changes how likely the agent considers cancer, from 0.5 to 0.9. It therefore increases the weight on the worse outcome under *both* acts, even though neither act causes cancer. In this case both causal expected utilities fall by $10(0.9 - 0.5) = 4$; smoking retains its advantage of 2.

More generally, if one act strongly statewise dominates another on causally independent states of nature, its causal expected utility is greater whatever the agent's credences over those states. Each state favours the dominating act, so any probability-weighted average does too. Here smoking gives 2 more utility in each state.

(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]

Solution. Let $s = \mathrm{Cr}_2(S)$. Since the agent chooses exactly one of S, N, they are mutually exclusive and exhaustive, so $\mathrm{Cr}_2(N) = 1 - s$. The law of total probability gives

$$\mathrm{Cr}_2(C) = \mathrm{Cr}_2(C \mid S)\,\mathrm{Cr}_2(S) + \mathrm{Cr}_2(C \mid N)\,\mathrm{Cr}_2(N).$$

Substitute the stated credences and solve:

$$\begin{aligned} 0.95s + 0.85(1 - s) &= 0.9,\\ 0.95s + 0.85 - 0.85s &= 0.9,\\ 0.10s + 0.85 &= 0.9,\\ 0.10s &= 0.05,\\ 10s &= 5,\\ s &= 0.5. \end{aligned}$$

Thus $\boxed{\mathrm{Cr}_2(S) = 0.5}$ and $\boxed{\mathrm{Cr}_2(N) = 1 - 0.5 = 0.5}$.

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]

Solution. For *every* pair of prospects P, Q,

$$P \succsim Q \quad \iff \quad \mathrm{EU}_u(P) \geq \mathrm{EU}_u(Q).$$

Thus the preference comparison and the expected-utility comparison agree in both directions.

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]

Solution. We must show that, for all P, Q, $P \succsim Q$ if and only if $\mathrm{EU}_v(P) \geq \mathrm{EU}_v(Q)$. The given identity applies to both prospects, so

$$\begin{aligned}
\mathrm{EU}_v(P) \geq \mathrm{EU}_v(Q) &\iff \alpha\mathrm{EU}_u(P) + \beta \geq \alpha\mathrm{EU}_u(Q) + \beta \\
&\iff \alpha\mathrm{EU}_u(P) \geq \alpha\mathrm{EU}_u(Q) \\
&\iff \mathrm{EU}_u(P) \geq \mathrm{EU}_u(Q) \\
&\iff P \succsim Q.
\end{aligned}$$

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

[2]

Solution. Since $P \succsim Q$, representation gives $\mathrm{EU}_u(P) \geq \mathrm{EU}_u(Q)$. Since $Q \succsim R$, representation gives $\mathrm{EU}_u(Q) \geq \mathrm{EU}_u(R)$. By transitivity of the numerical relation $\geq$, we infer $\mathrm{EU}_u(P) \geq \mathrm{EU}_u(R)$. Representation then gives $P \succsim R$, as required.

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

[2]

Solution. The expected utilities are real numbers. Any two real numbers are comparable, so

$$\mathrm{EU}_u(P) \geq \mathrm{EU}_u(Q) \quad \text{or} \quad \mathrm{EU}_u(Q) \geq \mathrm{EU}_u(P).$$

Representation turns the first inequality into $P \succsim Q$ and the second into $Q \succsim P$. If the expected utilities are equal, both comparisons hold. Hence the preference relation is complete.

(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]

Solution. Apply the mixture identity separately to the two prospects:

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

By representation, $P[p]R \succsim Q[p]R$ holds exactly when

$$p\mathrm{EU}_u(P) + (1-p)\mathrm{EU}_u(R) \geq p\mathrm{EU}_u(Q) + (1-p)\mathrm{EU}_u(R).$$

Subtract $(1-p)\mathrm{EU}_u(R)$ from both sides. This is equivalent to

$$p\mathrm{EU}_u(P) \geq p\mathrm{EU}_u(Q).$$

Since $p > 0$, divide both sides by p without reversing the inequality. The result is equivalent to

$$\mathrm{EU}_u(P) \geq \mathrm{EU}_u(Q),$$

which, by representation, is equivalent to $P \succsim Q$. All the steps are reversible, establishing both directions of Independence.

(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]

Solution. Write $a = \mathrm{EU}_u(P)$, $b = \mathrm{EU}_u(Q)$ and $c = \mathrm{EU}_u(R)$. Representation implies $a > b > c$.

We want $Q \sim P[p]R$. Indifference is equivalent to equal expected utility, so the indifference relation holds if and only if:

$$\begin{aligned} \mathrm{EU}_u(Q) &= p\mathrm{EU}_u(P) + (1-p)\mathrm{EU}_u(R), \\ b &= pa + (1-p)c, \\ b &= pa + c - pc, \\ b &= c + p(a-c), \\ b - c &= p(a-c), \\ p(a-c) &= b - c, \\ p &= \frac{b-c}{a-c}. \end{aligned}$$

Since $a > b > c$, we have $0 < b - c < a - c$. Dividing by $a - c > 0$ gives

$$0 < \frac{b-c}{a-c} < 1.$$

Thus the value found for p lies in $(0, 1)$. Since the rearrangements above are reversible, it gives $\mathrm{EU}_u(Q) = \mathrm{EU}_u(P[p]R)$, hence $Q \sim P[p]R$, as required.

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.

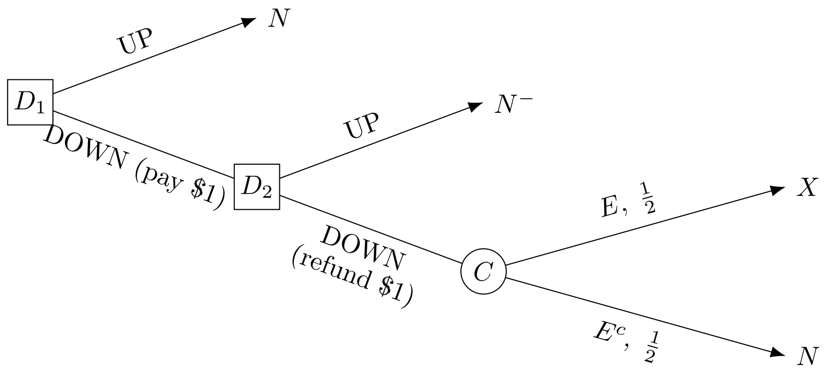

Naive choice selects a plan using current preferences and chooses afresh later. Resolute choice follows the plan selected initially. Sophisticated choice predicts later choices and evaluates their consequences using current preferences.

(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]

Solution. Using the D_1 row, UP at D_1 gives N with utility 0. DOWN at D_1 followed by UP at D_2 gives N^- with utility -1. DOWN at both nodes gives X or N, each with probability 1/2, so its initial expected utility is

$$\tfrac{1}{2}(8) + \tfrac{1}{2}(0) = 4 + 0 = 4.$$

Since $4 > 0 > -1$, the initial plan is **DOWN at D_1, then DOWN at D_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]

Solution. At D_2, DOWN gives X with utility -4 if E occurs and N with utility 0 if E^c occurs. Its expected utility is

$$\tfrac{1}{2}(-4) + \tfrac{1}{2}(0) = -2 + 0 = -2.$$

UP gives N^- with utility -1. Since $-1 > -2$, the naive agent chooses **UP at D_2**.

The final outcome is N^-: DOWN at D_1 cost \$1, and UP at D_2 gives no refund. The net cash change is $-\$1$. UP at D_1 would have given N with utility 0. Thus the agent is **guaranteed to be worse off**, by one utility unit using either row.

(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]

Solution. The resolute agent follows the initial plan: **DOWN at D_1, then DOWN at D_2**. The \$1 is refunded at D_2.

- If E occurs, the result is X, with initial utility 8. UP at D_1 gives 0, so the result is better by 8.
- If E^c occurs, the result is N, with initial utility 0, the same as UP at D_1.

The agent is **guaranteed not to be worse off**: using the D_1 utilities, the result is better if E occurs and equally good otherwise.

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

Solution. The agent predicts UP at D_2, since its later utility -1 exceeds the expected utility -2 of DOWN there. DOWN at D_1 therefore leads to N^- with initial utility -1, while UP at D_1 gives N with utility 0. Since $0 > -1$, the sophisticated agent chooses **UP at D_1**.

(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]

Solution. **No.** Let DOWN at D_1 cost $b > 0$. If the agent predicts UP at D_2, DOWN at D_1 leads to N with a cash loss of b. Its initial utility is

$$u_1(N) - b = 0 - b = -b < 0.$$

UP at D_1 gives N with utility 0, so the sophisticated agent chooses UP there. If the cost is zero or negative, DOWN at D_1 followed by UP at D_2 does not produce a loss.

(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.

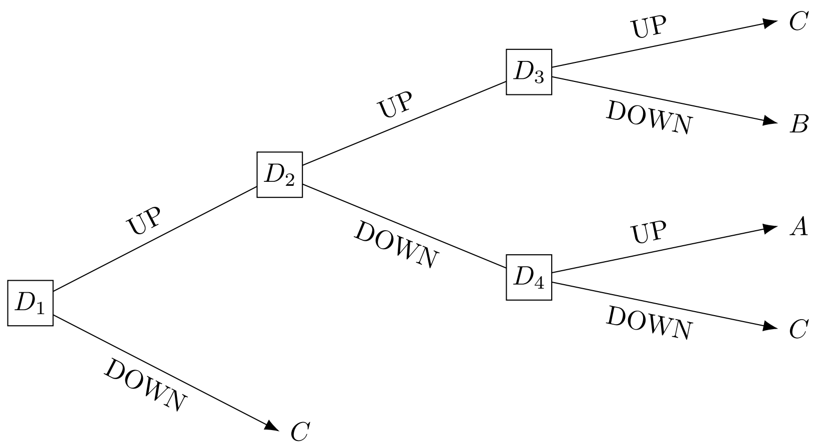

For each procedure below, state the sequence of UP/DOWN choices the agent makes, naming the decision nodes visited, and give the final outcome. You need not specify choices at nodes the agent does not reach.

(i) Naive choice. [3]

Solution. At D_1, all three outcomes are available through complete plans. The stipulated three-outcome rule selects a plan leading to the sole A, so the agent takes **UP to** D_2. At D_2, all three outcomes remain available, so the agent again plans for A and takes **DOWN to** D_4. At D_4, the available outcomes are only A and C. Choosing afresh, the agent uses $C \succ A$ and takes **DOWN**, giving C.

$$\boxed{D_1 : \text{UP}, \quad D_2 : \text{DOWN}, \quad D_4 : \text{DOWN}; \quad C.}$$

(ii) Sophisticated choice. [3]

Solution. Work backwards. At D_3, the comparison is C versus B; since $B \succ C$, the agent would choose DOWN and get B. At D_4, the comparison is A versus C; since $C \succ A$, the agent would choose DOWN and get C.

At D_2, the agent therefore compares the predicted B from UP with the predicted C from DOWN. Since $B \succ C$, they choose UP. At D_1, UP now predicts B, while DOWN gives C immediately; again $B \succ C$, so they choose UP.

$$\boxed{D_1 : \text{UP}, \quad D_2 : \text{UP}, \quad D_3 : \text{DOWN}; \quad B.}$$

The reasoning about D_4 is needed to make the prediction at D_2, although D_4 is not actually visited.

(iii) Resolute choice. [2]

Solution. Initially all three outcomes are available, so the stipulated rule selects a plan leading to A. Reaching A requires UP at D_1, DOWN at D_2, and UP at D_4. Resolute choice carries out that initial plan, including its instruction at D_4 rather than choosing afresh there.

$$\boxed{D_1 : \text{UP}, \quad D_2 : \text{DOWN}, \quad D_4 : \text{UP}; \quad A.}$$