



EXAMINATION PAPER

Examination Session: May/June	Year: 2025	Exam Code: MATH3071-WE01
---	----------------------	------------------------------------

Title: Decision Theory III

Time:	3 hours	
Additional Material provided:		
Materials Permitted:		
Calculators Permitted:	Yes	Models Permitted: Casio FX83 series or FX85 series.

Instructions to Candidates:	<p>Answer all questions.</p> <p>Section A is worth 40% and Section B is worth 60%. Within each section, all questions carry equal marks.</p> <p>Write your answer in the white-covered answer booklet with barcodes.</p> <p>Begin your answer to each question on a new page.</p>	
		Revision:

SECTION A

Q1 A football team manager is considering a player (player X) for the coming season. The manager has the following information:

- At the moment, the team wins 50% of their games and loses 50% (**we will assume there are no draws**)
- The team plays 20 games in a season.
- Winning a game leads to an income of £20,000, and losing a game leads to a loss of £10,000.
- Hiring player X will cost £40,000.
- Player X will either have a good season, in which case the probabilities of winning/losing change to (60%, 40%), or a bad season, in which case the probabilities stay the same.

We denote the probability of player X having a good season by q . We presume that we know q in advance. You may make the following assumptions:

- (i.) The results of any two games are independent.
- (ii.) The manager has no way of assessing player X's performance other than the results of games.

Answer the following questions:

- (a) Draw a decision tree for this problem and find the best decision about whether to hire player X, so as to optimise expected money value in terms of q .
- (b) Suppose that, rather than player X either having a 'good season' or a 'bad season', their inclusion in the team leads to a probability W of the team winning any given match. We want to model W as a random variable, and describe our knowledge of W as a distribution over $[0, 1]$. Discuss the advantages and disadvantages of using a Beta distribution as a prior distribution for W .

Q2 Suppose we have two mutually utility independent attributes X and Y , each of which can take any real-number value. We want to define a utility function over these attributes. We will use the origin $(x_0, y_0) = (0, 0)$

We have the marginal utilities (up to an additive constant)

$$U(x, y_0) = x^2 + x,$$

$$U(x_0, y) = \ln(1 + |y|).$$

Suppose that we have the following indifferences and preferences between rewards (a reward (a, b) is shorthand for the reward $(X, Y) = (a, b)$):

We are indifferent between $(2, 0)$ and $(0, 1)$ and prefer both to $(0, 0)$.

- (a) Suppose that we are also indifferent between $(3, 1)$ and $(5, 3)$. Find a formula for the general utility of reward (x, y) , with $U(x_0, y_0) = U(0, 0) = 0$. Are X and Y complementary or substitutable?
- (b) Suppose that, rather than being indifferent between $(3, 1)$ and $(5, 3)$ as in part (a), we are indifferent between (x_2, y_2) and (x_3, y_3) . Describe all values of (x_2, y_2, x_3, y_3) such that knowing only the information prior to part (a) and the information that ‘we are indifferent between (x_2, y_2) and (x_3, y_3) ’ does *not* allow us to find the utility for all rewards (x, y) .

Q3 Consider the decision table with the utilities of a person as given below, with a choice between four actions (a_1, \dots, a_4) while the unknown state has four possible values $(\theta_1, \dots, \theta_4)$.

	θ_1	θ_2	θ_3	θ_4
a_1	1	3	6	-2
a_2	8	0	-3	1
a_3	5	-1	3	0
a_4	0	1	1	1

Uncertainty about the state is reflected by the probabilities

$$P(\theta_1) = 0.1, \quad P(\theta_2) = 0.2, \quad P(\theta_3) = 0.3, \quad \text{and} \quad P(\theta_4) = 0.4.$$

- (a) Determine the optimal decisions according to each of the following optimality criteria: (1) Maximum expected utility; (2) Maximisation of minimum utility; (3) Minimisation of maximum regret. Explain briefly why one may want to use criterion (3).
- (b) This person decides to choose the action by maximisation of their optimism-pessimism index, with weight α for the minimum utility corresponding to a decision. It turns out that a_3 is the unique optimal decision. What can you conclude about the value of α used?

- Q4** (a) In late 2024, BBC Radio 2 invited listeners to vote for their favorite Elvis song, in order to create a list reflecting the popularity of his songs. Listeners could vote online on 90 selected songs, each person having 5 votes. These could be assigned to 5 different songs, but a listener could also vote multiple times for the same song. The final list was created by simply counting votes for each song, leading to a ranking in the logical way (with ties between two or more songs leading to equal ranking of those songs).

Consider this voting procedure from the perspective of Arrow's theory: State the axioms in Arrow's Impossibility Theorem and explain in detail for each axiom whether or not it is satisfied by this procedure.

- (b) Arrow's Impossibility Theorem for combining the individual preference orderings of members of a group over rewards to derive an overall group preference ordering, applies when there are at least three rewards. Explain, with detailed justification, why the case with two rewards is excluded.

SECTION B

Q5 We are going to cut down a randomly chosen tree from a given forest (so if a forest has proportion p of a tree type, we have probability p of getting that type). We are considering several forests with different proportions of tree types, and we have varying preferences for each tree type.

The tree types are: Cedar (C), Walnut (W), Oak (O), Larch (L), and Apple (A). On consideration, we have the following information:

- (i.) Our favourite wood is Cedar and our least favourite is Larch.
- (ii.) We are equivocal between an all-Walnut forest and a forest with 60% Cedar and 40% Larch.
- (iii.) We are equivocal between an all-Apple forest and a forest with 50% Walnut and 50% Larch.
- (iv.) We are equivocal between a forest which is 70% Oak, 10% Cedar and 20% Larch and a forest which is 50% Apple and 50% Walnut.
- (v.) Our preferences over various gambles over trees satisfy the assumptions of the Von Neumann-Morgenstern utility theorem.

Solve the following problems:

- (a) Find a formula for our utility of a forest with proportions $(p_C, p_W, p_O, p_L, p_A)$ of C, W, O, L, and A respectively (you may assume these are the only tree types, so $p_C + p_W + p_O + p_L + p_A = 1$). Scale the utility to have maximum value 1 and minimum value 0.
- (b) Suppose that forest 1 has 50% Larch, 30% Walnut, and 20% Oak, and forest 2 has 80% Apple and 20% Cedar. There is one (randomly chosen) tree from each forest available to buy, but we do not know the type of tree until we have bought it. We must buy one of the available trees, but we can choose which forest the tree comes from. We can pay a spy to tell us the tree type from forest 1. Our utility function for money (on the same scale as utility for trees) is $U(\mathcal{L}x) = \frac{1}{10}\sqrt{x}$ and we currently have £100. How much would we be willing to pay the spy?

- Q6** Suppose that we have a statistical decision problem for which the prior distribution for world-state W is $W \sim \text{Beta}(\alpha, \beta)$ for some $\alpha, \beta > 0$; that is, the density of W at w is given by:

$$\frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} w^{\alpha-1} (1-w)^{\beta-1},$$

The loss function for world-state w and decision x is given by $L(w, x)$. The following formula may be useful.

If X has a binomial distribution on n samples with probability p then:

$$E[X] = np \quad \text{and} \quad E[X^2] = np + n(n-1)p^2.$$

If X has a Beta distribution with parameters $\alpha, \beta > 0$ then

$$E[X^n] = \frac{\alpha(\alpha+1)\dots(\alpha+n-1)}{(\alpha+\beta)(\alpha+\beta+1)\dots(\alpha+\beta+n-1)}$$

and $\text{var}[X] = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}.$

- (a) Suppose $L(w, x) = (w^{2k})(x-w)^2$, for some integer k with $k \geq 0$. Show that the Bayes decision without sampling is given by

$$\frac{(\alpha + 2k)}{(\alpha + \beta + 2k)}.$$

- (b) Suppose that we have $L(w, x)$ as in the previous question with $k = 0$ and we are to take a binomial sample S of size n . The Bayes risk of sampling can be written as:

$$E_S[h(\alpha, \beta, n, S)] = E_W[g(\alpha, \beta, n, W)].$$

Find the form of *either* h or g and thus write the Bayes risk of sampling in this form. You do *not* need to calculate the expectations. Your expression for h or g should not contain any expectations or integrals.

- (c) In the previous question the Bayes risk of sampling can be shown to be:

$$\frac{\alpha\beta(\alpha + \beta + n + 1)}{(\alpha + \beta)(\alpha + \beta + 1)(\alpha + \beta + n)(\alpha + \beta + n - 1)}.$$

Suppose that the cost of taking n samples is $f(n)$, where f is increasing. Show that if f does not have an upper bound (that is, if $\lim_{n \rightarrow \infty} f(n) = \infty$) then there must be some N for which sampling is not cost-effective for $n > N$.

- Q7** (a) Consider a bargaining problem with 5 options, for which the utilities (u to first person, v to second person) are given in the following table:

	1	2	3	4	5
u	2	2	1	5	0
v	2	1	3	1	0

The utility pair $(0, 0)$ is the status-quo point for this problem.

Identify the feasible region and Pareto boundary for this bargaining problem. Derive the Nash point by optimising the function in the definition of the Nash point. Also derive the equitable distribution point and specify the bargains corresponding to the Nash point and the equitable distribution point.

- (b) Consider a bargaining problem with 4 options for which the utilities u and v are:

	1	2	3	4
u	2	1	0	0
v	0	1	α	0

The utility pair $(0, 0)$ is the status-quo point for this problem.

For *all* possible values of $\alpha > 0$, derive the Nash point for this bargaining problem using only the Nash axioms and a geometric method (so not by optimising the function in the definition of the Nash point), giving a detailed explanation of how the Nash axioms are used.

- Q8** In a particular game, R chooses row R1 or R2, C chooses column C1, C2, C3 or C4. The payoffs to R are as follows

	C1	C2	C3	C4
R1	4	9	6	8
R2	9	3	4	6

The payoff to C is minus the payoff to R .

- Use a graphical method to identify the minimax strategies for R and for C , and the value of the game.
- Explain briefly why a player might choose to play the minimax strategy.
- Suppose now that each player in this game decides on their strategy by minimising their maximum regret. Determine their strategies and the resulting payoff.
- Explain, using an example, why such games become more complex if the sum of the payoff to C and the payoff to R is not constant.