THE UNIVERSITY OF WARWICK

FIRST YEAR EXAMINATION: Mock Examination Paper

GAMES, DECISIONS AND BEHAVIOUR

Time Allowed: **2 hours**

Read carefully the instructions on the answer book and make sure that the particulars required are entered on each answer book.

Calculators are not needed and are not permitted in this examination.

ANSWER 3 QUESTIONS.

---

1. a) Consider a game which involves rolling two normal, unbiased 6-faced dice (one of which is red and the other blue, so we can distinguish between them). You may assume that the dice are "fair" in the sense that each number from 1 to 6 is equally likely to be obtained if one of them is rolled.

(i) What is the sample space (the set of possible outcomes), Ω, for this experiment?

We may write the result of rolling the two dice as an ordered pair with the first element corresponding to the result of the red die and the second to the result of the blue die, so a roll of "red 3, blue 4" would be represented as $(3,4)$. Using this representation:

$$\Omega = \{(1,1),(1,2),\ldots,(1,6),(2,1),\ldots,(6,6)\}$$

Or, more compactly:

$$\Omega = \{(x,y) : x \in \{1,\ldots,6\}; y \in \{1,\ldots,6\}\}$$

i.e. Ω is the set containing all pairs of integers from 1 to 6. **[2]**

(ii) If we are only interested in three possible classes of outcome:

- Both numbers are odd.
- Both numbers are even.
- One number is odd and the other even (in either order).

then what is the smallest algebra which we could employ?

We are interested in the three sets:

$$\begin{aligned} A_o = &\quad \{(x,y) : x \text{ odd}; y \text{ odd}\}, \\ A_e = &\quad \{(x,y) : x \text{ even}; y \text{ even}\} \text{ and} \\ A_m = &\quad \Omega \setminus (A_o \cup A_e) \end{aligned}$$

As with any algebra we must also include Ω and $\Omega^c = \emptyset$. Considering complements and unions of the sets we're interested in, it's clear we must also include $A_o^c = A_e \cup A_m$, $A_e^c = A_o \cup A_m$ and $A_m^c = A_o \cup A_e$. Check that these elements are sufficient to provide an algebra. It may be clearer if you think about what these sets correspond to (using words rather than symbols). **[3]**

(iii) What are the atoms of this algebra?

It's easy to see that A_o, A_e and A_m are all atoms and that all of the other sets are either empty or are supersets of these atoms (and thus they can't themselves be atoms). **[2]**

(iv) What is the associated probability mass function?

A direct solution would involve assuming that all points in Ω have the same probability and then finding the size of A_o, A_e and A_m. There are 3 odd and 3 even numbers available for each die, and each of the six numbers was assumed equally likely to be obtained. Thus each die is equally likely to produce an odd or even number (and there must be a probability of half associated with these, although a probability measure with respect to our minimal algebra does not encode that information). Both odd or both even must occur for $\frac{1}{4}$ of the possible configurations, ergo:

$$\mathbb{P}(A_o) = \frac{1}{4} = \mathbb{P}(A_e), \text{ and } \mathbb{P}(A_m) = 1 - \frac{1}{4} - \frac{1}{4} = \frac{1}{2}$$

[3]

(v) How big is the algebra used in this experiment – how many elements does it contain? How big is the largest algebra over Ω... which is it easier to work with?

The simple algebra we used had 8 elements. Ω has $6 \times 6 = 36$ elements and so the largest algebra over Ω (the set of all subsets or *power set* of Ω) contains 2^{36} elements, a rather large number. To convince yourself of this note that one way of representing a subset of Ω is a vector of length 36 which takes a value 1 in each position if the corresponding element of Ω is contained in the set and 0 otherwise. Each such vector describes a unique, valid subset of Ω and there are clearly 2^{36} such vectors. Although the algebra we used leads to a much less powerful description of the underlying probabilistic event it will be rather less cumbersome than the full probabilistic description and is sufficient if we only care about these particular events. **[2]**

b) Joe owes £$100,000$ to the Mafia. They want the money immediately. He has £$60,000$. He has the opportunity to place a bet in a casino in which with probability $\frac{1}{2}$ he will make a profit of £$40,000$ and with probability $\frac{1}{2}$ he will lose £$60,000$.

(i) What is his EMV decision?

The two possible decisions are to play or not to play. Not playing has an expected reward of zero; playing has an expected reward of -£10,000. The EMV decision is not to play. **[2]**

(ii) Does that coincide with common sense?

Most people in this position would probably accept that having £60,000 is unlikely to be of any value whilst having £100,000 might keep them alive and would be prepared to indulge in the bet, even if it's an unfair bet. **[2]**

(iii) What utility function would you advise Joe to use in this circumstance?

Joe needs to win at least £40,000 in order to be able to pay his debts. This is his sole concern. If he assigns an outcome in which he wins at least $\pounds 40,000$ a utility of 1 and other outcomes a utility of 0 then this represents these beliefs.

$$U(x) = \begin{cases} 0 & x < 40,000 \\ 1 & x \geq 40,000 \end{cases}$$

[2]

(iv) What is the expected utility of the two possible decisions? And what is the expected utility decision?

The expected utility of not playing is zero: there is no possible way of achieving a utility other than zero.

The expected utility of playing is $0.5[1+0] = 0.5$. It's not a coincidence that this is equal to the probability of achieving an outcome with a utility of one under the specified utility function.

Consequently, maximising the expected utility in this case tells us that Joe should play. **[2]**

---

2. a) The diagnosis of a certain disease D is not straight forward. However, there is a certain task T at which people without the disease outperform people with the disease. People without D succeed at T with a probability of π, while people with the disease end up guessing corresponding to a success probability of 0.5. One way of diagnosing D is to have a person perform T a number of times and count the number r of successes. Let p bet the prevalence of D (i.e. the probability someone drawn at random from the reference population has D).

Immediate treatment (d_1) for D can be given at cost C, regardless of whether or not a person has D. It is a harmless intervention. Deferred treatment (d_2) is also harmless, but comes at cost $\beta\, C$ and is only given to a person if the person has.

You are working on mathematical models to advice doctors in this decision. From a financial point of view, the decision reduces to a relationship between the expected losses $\overline{L}_i = \mathbb{E}[L(d_i)]$ for d_i $(i = 1, 2)$. Immediate treatment will be performed if $\overline{L}_1 \leq \overline{L}_2$, otherwise it will be deferred. For simplicity, the loss function L consists of just the costs described above.

(i) Write down the formula for the loss function and calculate the expected losses $\overline{L}_i$ $(i = 1, 2)$ as functions of the parameters p, π, β, n and the number r of correctly performed tasks.

$L_1(d_1)$ is always C, so independently of any of the other parameters, $\overline{L}_1 = C$.

$$\overline{L}_2 = \beta C P(D|r) + 0 \cdot P(D^C|r)$$

$P(D|r)$ can be calculated using Bayes theorem:

$$P(D|r) = \frac{\binom{n}{r} 0.5^n\, p}{\binom{n}{r}\left(0.5^n\, p + \pi^r (1-\pi)^{n-r}(1-p)\right)}$$

so

$$\overline{L}_2 = \frac{\beta C \binom{n}{r} 0.5^n\, p}{\binom{n}{r}\left(0.5^n\, p + \pi^r (1-\pi)^{n-r}(1-p)\right)}$$

[4]

(ii) Derive a formula that expresses the condition $\overline{L}_1 \leq \overline{L}_2$ as a condition on r.

Using (i), $\overline{L}_1 \leq \overline{L}_2$ the condition is equivalent to

$$0.5^n\, p + \pi^r (1-\pi)^{n-r}(1-p) \leq \beta\, 0.5^n\, p$$

which can further be transformed into

$$\pi^r (1-\pi)^{n-r} \leq (\beta - 1)\, 0.5^n \frac{p}{1-p},$$

in other words,

$$\left(\frac{\pi}{1-\pi}\right)^r \leq (\beta - 1)\left(0.5/(1-pi)\right)^n \frac{p}{1-p}$$

Since log is a strictly increasing functions, this is equivalent to

$$r\log\left(\frac{\pi}{1-\pi}\right) \leq \log\left((\beta-1)\left(0.5/(1-\pi)\right)^n \frac{p}{1-p}\right)$$

and considering that $\frac{\pi}{1-\pi} \in (0,1)$ where log is negative, this is equivalent to

$$r \geq \frac{\log\left((\beta-1)\left(0.5/(1-\pi)\right)^n \frac{p}{1-p}\right)}{\log\left(\frac{\pi}{1-\pi}\right)}$$

[5]

(iii) Calculate the threshold for r more explicitly for $p = 0.2, \pi = 0.9, \beta = 10$ and $n = 4$.

Using R plugging in the parameters above yields that the threshold for r is about 3.299017. This means, for $r = 0, 1, 2, 3$ the EMV approach recommends using strategy L_2 and for $r = 4$ it recommends using L_1.
Note: This bit would not have been included if this was a real exam, because it needs a calculator. **[1]**

b) Consider a zero sum game with the following payoff matrix (for player 1; remember player 2 has payoffs corresponding to the negative of those of player 1 in a zero sum game):

	δ_1	δ_2	δ_3
d_1	0	5/6	1/2
d_2	1	1/2	3/4

(i) What is player 1's maximin mixed strategy?

The maximin strategy has associated probabilities $[x, 1-x]$ with x chosen to maximise the expected return obtained if player 2 makes the worst possible move. Plot expected return against x for each of player 2's possible moves:

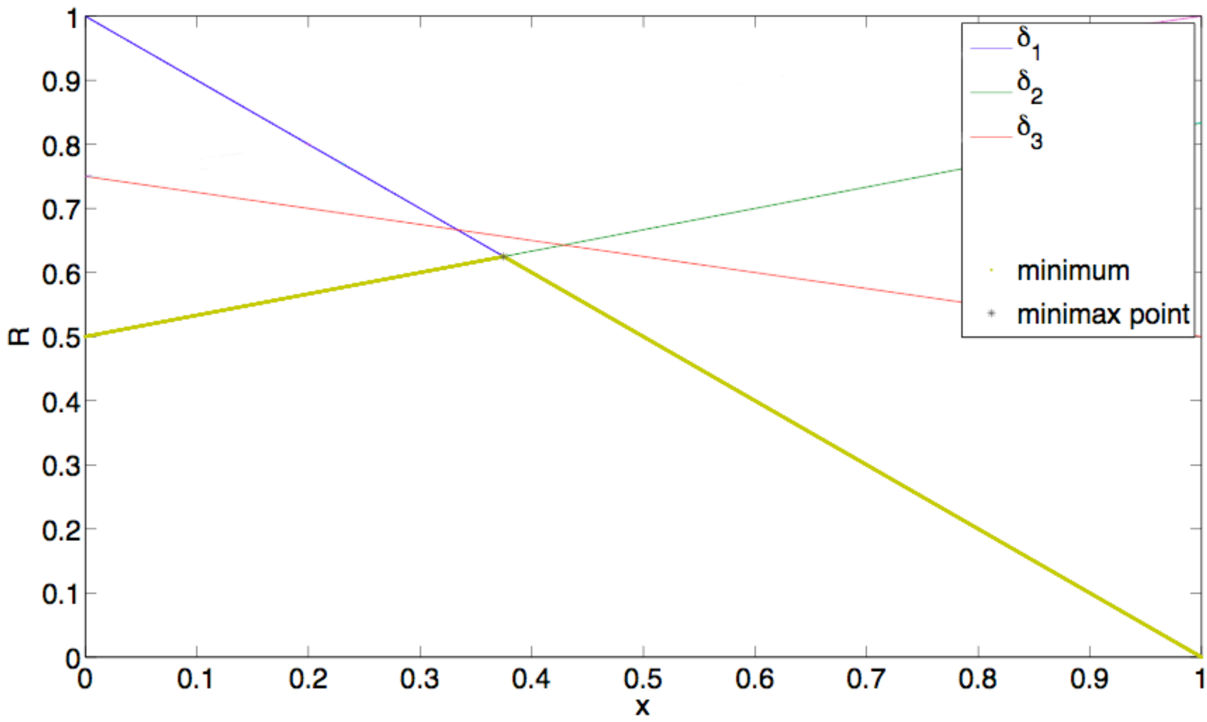

[3]

From this figure it's clear that δ_1 and δ_2 have associated lines which intersect at the maximin point. Checking by calculation, it's clear that this intersection occurs at $x = 0.375$ at which point player 1 has an expected reward of 0.625.

[2]

(ii) What is player 2's maximin mixed strategy?

Player 2 may use the fact that the value of the game is the expected reward of player 1 at their maximin strategy (i.e. 0.625) and that player 1 only has to consider moves δ_1 and δ_2 (knowing that they will do better if player 2 plays any other strategy). Thus their maximin mixed strategy is $(y^\star, 1 - y^\star, 0, 0, 0)$ with $y^\star$ chosen to achieve a reward for player 1 of at most 0.625. Thus, $\frac{5}{6}(1 - y^\star) \leq 0.625$ and $y^\star + \frac{1}{2}(1 - y^\star) \leq 0.625$. And we must have $y^\star = 0.25$.

[3]

(iii) What is the value of this game?

The value of the game is the expected reward of the first player with both adopting their maximin mixed strategy: 0.625.

[2]

3. a) Consider a game with the following payoff matrix:

	α	β	γ	δ
a	(4,5)	(1,4)	(1,5)	(3,3)
b	(5,6)	(3,4)	(2,5)	(4,7)
c	(3,7)	(6,10)	(8,11)	(2,8)
d	(4,8)	(8,5)	(10,4)	(3,4)

Using iterated elimination of dominated strategies show that there is a single strategy which the players should, under the assumption that *rationality is common knowledge*, adopt deterministically.

Player 1 should always prefer d to c. c is eliminated. Player 1 should always prefer b to a. a is eliminated. Knowing that player 1 will play b or d, player 2 optimizes their strategy for the reduced payoff matrix:

	α	β	γ	δ
b	(5,6)	(3,4)	(2,5)	(4,7)
d	(4,8)	(8,5)	(10,4)	(3,4)

Player 2 should prefer α to both β and γ as it always provides a higher payoff in this reduced game. Player 1 does addresses the further reduced payoff matrix:

	α	δ
b	(5,6)	(4,7)
d	(4,8)	(3,4)

Now, b now dominates d, so d can be eliminated. From the remaining options, player 2 will choose move δ.

Thus we end up with the rationally-justified solution that player 1 plays b and player 2, δ. **[7]**

b) Consider the elicitation of preferences below.

I Which of the following situations would you prefer?
A: Certainty of receiving 1 million.
B: 10% chance of 5 million; 89% chance of 1 million; 1% chance of nothing.

II Which of the following situations would you prefer?
C: 11% chance of 1 million; 89% chance of nothing.
D: 10% chance of 5 million; 90% chance of nothing.

Empirical studies have shown that most people prefer A to B and D to C.

(i) Why do these preferences contradict expected utility theory (EUT)?

The choices above can be rewritten as follows (using m for million):

I Which of the following situations would you prefer?
A: 89% chance of 1m; 11% chance of 1m.
B: 89% chance of 1m; 1% chance of nothing; 10% chance of 5m.

II Which of the following situations would you prefer?
C: 89% chance of nothing; 11% chance of 1m.
D: 89% chance of nothing; 1% chance of nothing; 10% chance of 5m.

Let u be the utility function of a person making these choices. Calculating all the expected utilities $\overline{U}$ (in millions) yields

$$\begin{aligned}\overline{U}(A) &= 0.89\,u(1) + 0.11\,u(1)\\ \overline{U}(B) &= 0.89\,u(1) + 0.01\,u(0) + 0.10\,u(5)\\ \overline{U}(C) &= 0.89\,u(0) + 0.11\,u(1)\\ \overline{U}(D) &= 0.89\,u(0) + 0.01\,u(0) + 0.10\,u(5)\end{aligned}$$

In comparisons some of the addends are redundant leading to

$$\begin{aligned}\overline{U}(A) > \overline{U}(B) &\Leftrightarrow 0.11\,u(1) > u(0) + 0.10\,u(5)\\ \overline{U}(C) > \overline{U}(D) &\Leftrightarrow 0.11\,u(1) > u(0) + 0.10\,u(5)\end{aligned}$$

The conditions on the right hand side are identical. Hence, a subject making decisions according to EUT should either prefer A and C or should prefer B and D. (*Note that these preferences are independent of the utility function. All we are using in this argument is that the same utility function is used consistently.*) This is a contradiction to the empirical observation. [5]

(ii) Model and explain this behaviour using prospect theory. For the value function v assume $v(0) = 0, v(1) = 1$ and $v(5) = 2$ (defined on millions of the currency unit). For the probability weighting function use the Prelec function, $w(p) = \exp(-\delta(-log(p))^\gamma)$ ($p \in [0,1]$), with parameters $\delta = 1.2$ and $\gamma = 0.25$.

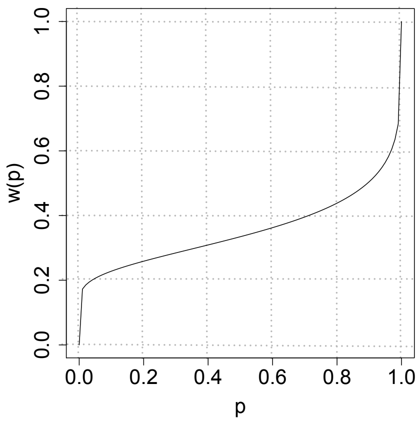

p	w(p)
0.00	0.0
0.01	0.2
0.10	0.2
0.11	0.2
0.89	0.5
0.90	0.5
1.0	1.0

(rounded)

Using the original formulation of the pay-offs, prospect theory models the expected value of the different deals as follows (in millions):

$$V(A) = w(1)\,v(1)$$
$$V(B) = w(0.10)\,v(5) + w(0.89)\,v(1) + w(0.01)\,v(0)$$
$$V(C) = w(0.11)\,v(1) + w(0.89)\,v(0)$$
$$V(D) = w(0.10)\,v(5) + w(0.90)\,v(0)$$

Plugging in the (approximative) values for w as given in the table, and then the values for v

$$V(A) = 1\,v(1) = 1 \cdot 1 = 1$$
$$V(B) = 0.2\,v(5) + 0.5\,v(1) + 0.2\,v(0) = 0.2 \cdot 2 + 0.5 \cdot 1 + 0.2 \cdot 0 = 0.9$$
$$V(C) = 0.2\,v(1) + 0.5\,v(0) = 0.2 \cdot 1 + 0.5 \cdot 0 = 0.2$$
$$V(D) = 0.2\,v(5) + 0.5\,v(0) = 0.2 \cdot 2 + 0.5 \cdot 0 = 0.4$$

Hence, $V(A) > V(B)$, but $V(C) < V(D)$, which reflects the empirical findings. **[5]**

(iii) Explain which aspect of the probability distortion expressed by w leads to a modification of preferences when using prospect theory rather than EUT. Include a crucial quantitative observation about w to make your point.

People perceive the certain option A as superior to the bet B, because they overrate the 1% chance of getting nothing in B. Being certain of something has an additional value not accounted for in EUT. Prospect theory accounts for this by including a probability weighting function w with steep increase close to $p = 1$. The probability of winning at least one million in B is 99%, but after probability weighting it only amounts to 0.7. **[3]**

[1]

---

[1] *Note: Historical background. The two choice situations were proposed, hypothetically, in a paper by Allais titled "Le Comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine" (English translation: Rational man's behavior in the presence of risk: critique of the postulates and axioms of the American school) in Econometrica 21, 1953. To make the monetary values from the original paper appear more realistic, they have been converted into nouveau franc, anticipating its introduction in 1958. (In 1963, franc became anciens franc and nouveau franc became franc.)*

4. a) You are a 4th year student doing a Master's dissertation project on cognitive heuristics and biases of the type studied extensively by Tversky and Kahneman in their research program in the 1970s to 1980s. Part of your work consists in running your own empirical studies on this. You are given the opportunity to study this empirically by conducting a short questionnaire based survey in an ST222 lecture.

Among other things, you want to study how prior exposure to numerical quantities affects the estimate of an unrelated numerical quantity. You plan to give all of the students the same estimation task, but expose half of them to a different quantity then the other half. Design a suitable experiment by working through the following steps.

(i) Formulate the hypothesis you want to test.

The respondents will answer differently to the questions, depending on which quantity they were exposed to prior to filling out the questionnaire. **[1]**

(ii) Select a numerical quantity you will ask the students to estimate.

How many students typically graduate with a First in Warwick Maths (or MORSE, or MathStat)? **[1]**

(iii) How would you split them into two groups (in practical terms)?

It depends a bit on the lecture theatre. Could be split it into a left hand and a right hand side. Or, if there are three blocks of seating rows, it could be split into an inner block and a combination of the two outer blocks. **[2]**

(iv) Which prior exposures would you give them, and how?

Put some introductory text on the questionnaire that contains a statistical summary about graduate destinations. For one group, make this number low (e.g. 10%) and for the other one make it high (e.g. 90%). **[2]**

(v) Name two difficulties that may arise during the conduction of the survey. *Use no more than 10 words each.*

1. Unequal block sizes, 2. subjects may look at others' questionnaires. **[2]**

(vi) Briefly describe two essential limitations of your study and how they would effect the results. *Use no more than 20 words each.*

1. Students in mathematical degrees are not representative of the population and may process numbers differently. Hence the results may not be generally valid. 2. Students may assume some connection between the introductory text and the questions. So the study is not testing the hypothesis as intended. **[2]**

b) Let Ω be a space of outcomes and $\mathcal{F}$ be an algebra of subsets of Ω. Let $A, B \in \mathcal{F}$ be two sets of events with $A \subset B$. One of your flat mates holds the belief that $P(A) > P(B)$. Can you construct a Dutch book against him? Detail one way how you can do that.

The bet $b(n,k)$ pays $£1$ if a ball of colour c is drawn from a bag that contains n balls k of which are of colour c. The bet $b(A)$ pays $£1$ if A happens. Let $k(A)$ and $k(B)$ be the smallest integers such that the bets $b(n,k(A))$, $b(n,k(B-A))$ and $b(n,k(B))$ are equivalent to the bets $b(A)$, $b(B-A)$ and $b(B)$. $b(n,k(A))$ refers to drawing red balls, and $b(n,k(B-A))$ refers to drawing green balls. Holding them together is equivalent to holding $b(n,k(B))$. Since $A \subset B$, for n large enough, there is an integer $k \geq 1$, such that $k(A) = k(B) + k$. Assume the flat mate owns $b(B)$. Given he believes $P(A) > P(B)$, he would be prepared to pay an amount $M > 0$ to exchange this for $b(A)$. However, this is equivalent to $b(n,k(A))$, but since $k(A) = k(B) + k < k(B)$, he actually made a loss. **[5]**

c) Let w be a probability weighting function and v be a value function in a prospect theory model. Two prospects (x_i, p_i) $(i = 1,2)$ are equally preferable if their values $v(x_i)\,w(p_i)$ $(i = 1,2)$ are equal.

Let $\lambda > 0$. A probability weighting function w satisfies the λ-*reduction invariance* property if the following is true:

The prospects $((x,p),q)$ and (x,r) are equally preferable, if and only if the prospects $((x,p^\lambda),q^\lambda)$ and (x,r^λ) are equally preferable.

Prove that any Prelec function $w(p) = \exp(-\delta(-\log p)^\gamma)$ $(p \in [0,1])$ with parameters $\delta, \gamma > 0$ is λ-reduction invariant for all $\lambda > 0$.

It is sufficient to show that the condition

$$\big((x,p^\lambda),q^\lambda\big) \sim (x,r^\lambda) \qquad (\star)$$

yields independently of λ. Using strict monotonicity of $\log$,

$$\begin{aligned}
(\star) &\Leftrightarrow \big(v(x)w(p^\lambda)\big)w(q^\lambda) = v(x)w(r^\lambda) \\
&\Leftrightarrow \log(w(p^\lambda)w(q^\lambda)) = \log w(r^\lambda) \\
&\Leftrightarrow -\delta(-\log p^\lambda)^\gamma - \delta(-\log q^\lambda)^\gamma = -\delta(-\log r^\lambda)^\gamma \\
&\Leftrightarrow (-\lambda \log p)^\gamma + (-\lambda \log q)^\gamma = (-\lambda \log r)^\gamma \\
&\Leftrightarrow (-\log p)^\gamma + (-\log q)^\gamma = (-\log r)^\gamma
\end{aligned}$$

Hence, whether or not the condition $(\star)$ holds is independent of λ. (The last line in the above calculation is equivalent to the condition $\big((x,p),q\big) \sim (x,r)$.) **[5]**
