

Final Exam

INSTRUCTIONS, SCORING: The points for each question are specified at the beginning of the question. There are 160 points on the exam. Please show your work, so that we can give you partial credit. Good luck!

1. (20 points) Consider the binary relation \succ defined as follows on the reals: For $x, y \in \mathbf{R}$, $x \succ y$ if $x - y$ is an integer.

- (a) Is this binary relation transitive?
(b) Is this relation negatively transitive?

Solution: (a) Yes; if $x - y$ and $y - z$ are integers, then $x - z = (x - y) + (y - z)$ is an integer.

(b) No; for example $1 \not\succeq .5$ and $.5 \not\succeq 0$, but $1 \succ 0$.

2. (20 points) Let Ω be a finite state space, O a set of outcomes, and $u : O \rightarrow \mathbf{R}$ a payoff function. Let D denote a finite set of acts $d : \Omega \rightarrow O$. Define $u_d(\omega) = u(d(\omega))$. Denote by $r_d(\omega)$ the regret from using d in state ω .

- (a) Define $r_d(\omega)$.

Solution: $r_d(\omega)$ is the difference between the utility of the best outcome in state ω and the utility of $d(\omega)$.

- (b) Let p be a probability distribution on Ω . Define $d \succeq_1 e$ iff $E_p[u_d(\omega)] \geq E_p[u_e(\omega)]$. Define $d \succeq_2 e$ if and only if $E_p[r_d(\omega)] \leq E_p[r_e(\omega)]$. Show that $\succeq_1 = \succeq_2$.

Solution: You're being asked to show that expected utility is the same as expected regret. This was done in class. If d_ω is the best act in state ω , then

$$\begin{aligned} E_p[r_d] &= \sum_{\omega} p(\omega)(u(d_\omega(\omega)) - u(d(\omega))) \\ &= \sum_{\omega} p(\omega)(u(d_\omega(\omega)) - \sum_{\omega} p(\omega)u(d(\omega))) \\ &= B - E_p(u_d(\omega)), \end{aligned}$$

where $B = \sum_{\omega} p(\omega)(u(d_\omega(\omega)))$ is a fixed constant. It follows that $E_p[r_d] \leq E_p[r_e]$ iff $E_p(u_d(\omega)) \geq E_p(u_e(\omega))$; i.e. $d \succeq_1 e$ iff $d \succeq_2 e$.

- (c) Let $d \succeq_3 e$ if and only if $\min_{\omega} u_d(\omega) \geq \min_{\omega} u_e(\omega)$. Let $d \succeq_4 e$ if and only if $\max_{\omega} r_d(\omega) \leq \max_{\omega} r_e(\omega)$. Show that if $\max_{\omega} u_d(\omega)$ is a constant independent of ω , then $\succeq_3 = \succeq_4$.

Solution: Suppose that $A = \max_{\omega} u_d(\omega)$. Then $r_d(\omega) = A - u_d(\omega)$, for all states ω . Thus, $\max_{\omega} r_d(\omega) = A - \min_{\omega} u_d(\omega)$. It easily follows that $\max_{\omega} r_d(\omega) \leq \max_{\omega} r_e(\omega)$ iff $\min_{\omega} u_d(\omega) \geq \min_{\omega} u_e(\omega)$; i.e. $d \succeq_3 e$ iff $d \succeq_4 e$.

3. (20 points) This problem computes some minimax regret rules. In this problem the state space is $\Omega = [0, 1]$, the action set is $D = \{0, 1\}$ and the payoff is $u_d(\omega) = -(d - \omega)^2$.

- (a) Compute the regret of each act and identify which decisions are optimal under the minimax regret criterion.

Solution: The best act in state ω is 1 if $\omega \in [1/2, 1]$; the utility of act 1 is $-(1 - \omega)^2$. The best act in state ω is 0 if $\omega \in [0, 1/2]$. It is easy to see that in all states, the utility of both 0 and 1 is in $[-1, 0]$. It follows that the regret of both 0 and 1 in any state can be at most 1. The regret of act 1 is 1, because in state 0, the best act (0) gives utility 0, while 1 gives utility -1; in every other state, the regret of 1 is less than 1. Similarly, the regret of 0 is 1, because in state 1, 0 has utility -1 and the best act (1) has utility 0. Thus, the two acts are equivalent under the minimax regret criterion.

- (b) Now suppose that randomized decision procedures can be used. The set of rules is now $p \in [0, 1]$, where p is the probability of choosing $d = 0$. The payoff to rule d in state ω is its expected payoff, $U(\omega, p) = -p\omega^2 - (1-p)(1-\omega)^2$. Find the randomized rules which are optimal under the minimax regret criterion.

Solution: Now the best act is 1/2. To see this, note that the utility of act p in state ω is $-p\omega^2 - (1-p)(1-\omega)^2$. The best act in states $\omega \in [0, 1/2]$ is still 0, and the best act in states $\omega \in [1/2, 1]$ is still 1. It easy to see that if $p > 1/2$, then the regret of p is maximized is maximized at state 1, where it is p ; if $p \leq 1/2$, then the regret of p is maximized at state 0, where it is $(1-p)$. Thus, if $p \neq 1/2$, then the regret of p is greater than 1/2. Thus, the act $p = 1/2$ minimizes regret, and has regret 1/2.

4. (20 points) Given two acts a and b , say that a *weakly dominates* b (with respect to a utility function u) if for all states $s \in S$, we have $u(a(s)) \geq u(b(s))$ and for some $s' \in S$, we have $u(a(s')) > u(b(s'))$. Suppose that a weakly dominates b . Which of $a \succ b$, $a \sim b$, and $b \succ a$ can/must happen with each of the following decision rules:

- (a) maximin,
 (b) minimax regret,

(c) the principle of insufficient reason.

There was a typo in the next line of the question: It said “If you think one of these can’t happen, given an example . . .”. I meant to say “If you think one of these can happen, give an example that demonstrates it; if you think it must happen, explain why. (Note: 1-2 sentence explanations suffice in each case.)” Most people figured out the intent of the question.


Solution: (a) With maximin, we must have $a \succeq b$; we may have $a \succ b$ and we may have $a \sim b$. (Of course, we can’t have $b \succ a$.) If in fact $u(a(s)) > u(b(s))$ for all states $s \in S$ (which is consistent with weak dominance), then we certainly have $a \succ b$. To see that we might have $a \sim b$, let $S = \{s_1, s_2\}$, and suppose that $u(a(s_1)) = u(b(s_1)) = 0$, $u(a(s_2)) = 1$ and $u(b(s_2)) = 0$. Then the worst-case outcome for both a and b is 0, so $a \sim b$.

(b) With minimax regret, like with maximin, we must have $a \succeq b$; we may have $a \succ b$ and we may have $a \sim b$. (Of course, we can’t have $b \succ a$.) If there are no other acts, then it is easy to see that we must have $a \succ b$ (in fact, the regret of a is 0, and the regret of b is $\max_{s \in S} (u(a(s)) - u(b(s))) > 0$). To see that we might have $a \sim b$, let $S = \{s_1, s_2\}$, and suppose that $u(a(s_1)) = u(b(s_1)) = 0$, $u(a(s_2)) = 1$ and $u(b(s_2)) = 0$. Clearly a weakly dominates c . Now let c be another act such that $u(c(s_1)) = 3$ and $u(c(s_2)) = 0$. Then the regret of both a and b is 3, so $a \sim b$.

(c) With the principle of insufficient reason, we must have $a \succ b$, since the expected utility of a is $\sum_{s \in S} u(a(s))/|S|$, which is greater than $\sum_{s \in S} u(b(s))/|S|$, which is the expected utility of b .

[Grading: 7 for each of parts (a) and (b), 6 for part (c).]

5. (20 points) Consider the following Bayesian network containing 3 Boolean random variables (that is, the random variables have two truth values—*true* and *false*):



Suppose the Bayesian network has the following conditional probability tables (where X and \bar{X} are abbreviations for $X = \text{true}$ and $X = \text{false}$, respectively):

$$\begin{aligned}
 \Pr(A) &= .1 \\
 \Pr(B \mid A) &= .7 \\
 \Pr(B \mid \bar{A}) &= .2 \\
 \Pr(C \mid A) &= .4 \\
 \Pr(C \mid \bar{A}) &= .6
 \end{aligned}$$

- (a) What is $\Pr(\bar{B} \cap C \mid A)$?
- (b) What is $\Pr(A \mid \bar{B} \cap C)$?
- (c) Suppose we add a fourth node labeled D to the network, with edges from both B and C to D . For the new network
- (i) Is A conditionally independent of D given B ?
 - (ii) Is B conditionally independent of C given A ?

In both cases, explain your answer.

Solution: (a) $\Pr(\bar{B} \cap C \mid A) = \Pr(\bar{B} \mid A)$, since, according to the Bayesian network, C is independent of B given A . We are given that $\Pr(B \mid A) = .7$, so $\Pr(\bar{B} \mid A) = .3$.

(b) By Bayes' Rule,

$$\Pr(A \mid \bar{B} \cap C) = \Pr(\bar{B} \cap C \mid A) \Pr(A) / \Pr(\bar{B} \cap C).$$

According to the Bayesian network, B and C are conditionally independent given A , so $\Pr(\bar{B} \cap C \mid A) = \Pr(\bar{B} \mid A) \times \Pr(C \mid A) = .3 \times .4 = .12$. We are also given that $\Pr(A) = .1$. Finally,

$$\begin{aligned} & \Pr(\bar{B} \text{ int } C) \\ &= \Pr(\bar{B} \text{ int } C \mid A) \Pr(A) + \Pr(\bar{B} \text{ int } C \mid \bar{A}) \Pr(\bar{A}) \\ &= \Pr(\bar{B} \mid A) \times \Pr(C \mid A) \Pr(A) + \Pr(\bar{B} \mid \bar{A}) \times \Pr(C \mid \bar{A}) \Pr(\bar{A}) \\ &= (.3 \times .4 \times .1) + (.8 \times .6 \times .9) \\ &= .012 + .432 = .444 \end{aligned}$$

Putting this all together, we get that $\Pr(A \mid \bar{B} \cap C) = .012 / .444 = 1/37$.

(c)(i) No; A is not conditionally independent of D given B , since A and D might be correlated via C . (A is conditionally independent of D given B and C .)

(ii) Yes, B is conditionally independent of C given A . In general, a node is independent of its nondescendants given its parents. A is the only parent of B , and C is a nondescendant.

[Grading: 5 for (a), 9 for (b), and 3 each for (c)(i) and (c)(ii).]

6. (20 points) Consider the following CP (ceteris paribus) network:

A --> B --> C

with the following conditional preference tables:

$a \succ \bar{a}$

a	$b \succ \bar{b}$
\bar{a}	$\bar{b} \succ b$

b	$c \succ \bar{c}$
\bar{b}	$\bar{c} \succ c$

- (a) Does it follow from this CPnet that $abc \succ a\bar{b}\bar{c}$? (Explain why or why not; a simple yes or no will not get any points.)
- (b) Does it follow from this CPnet that $\bar{a}\bar{b}\bar{c} \succ abc$? (Again, explain why or why not.)

Solution: (a) Yes, this follows. Here is the argument:

$$\begin{aligned} abc &\succ ab\bar{c} \quad [\text{since } c \succ \bar{c} \text{ conditional on } b] \\ &\succ a\bar{b}\bar{c} \quad [\text{since } b \succ \bar{b} \text{ conditional on } a] \end{aligned}$$

(b) No, this does not follow. The easiest way to see it is just to observe that, unconditionally, a is preferred to \bar{a} . Nothing can change this unconditional preference. In fact, these two outcomes are incomparable (although you didn't have to show that).

7. (20 points) Consider the lotteries:

- A** If Hillary Clinton is the next President then you win \$1 with probability p and \$0 with probability $1 - p$; if she is not then you win \$1 with probability 0 and \$0 with probability 1.
- B** If Hillary Clinton is the next President then you win \$1 with probability 0 and \$0 with probability 1; if she is not then you win \$1 with probability q and \$0 with probability $1 - q$.
- C** You win \$1 with probability 1/2 and \$0 with probability 1/2.

Suppose that $\$1 \succ \0 .

- (a) Suppose p and q are such that the decision maker is indifferent between lotteries A and B. Assuming that this decision maker is a subjective expected utility maximizer (with a state independent payoff function) determine the decision maker's subjective probability that Hillary Clinton is the next President.

Solution: Let H be the event that Hillary is the next President. If the decision maker is indifferent between A and B , then

$$\Pr(H)(pu(1)+(1-p)u(0))+(1-\Pr(H))(u(0)) = \Pr(H)(u(0))+(1-\Pr(H))(qu(1)+(1-q)u(0)).$$

Straightforward computation shows that it follows that

$$(u(1) - u(0))\Pr(H)p = (u(1) - u(0))q(1 + \Pr(H)).$$

Since $1 \succ 0$, it follows that $u(1) - u(0) > 0$, so we can divide both sides by $u(1) - u(0)$ to get $\Pr(H)p = q(1 + \Pr(H))$. It follows that $\Pr(H) = q/(p - q)$. Note that the exact value of $u(0)$ and $u(1)$. This would have to be the case with an expected utility maximizer. The utility function is unique only up to affine transformations; given any utility function u , and real numbers $a < b$, there is a positive affine transformation of u to a utility function u' such that $u'(0) = a$ and $u'(1) = b$.

- (b) Now suppose that $p = q = 1$, the decision maker may or may not be indifferent between A and B and for these values of p and q we have $C \succ A$ and $C \succ B$. Show that in this case the decision maker cannot be a subjective expected utility maximizer.

Solution: If the decision maker is a subjective expected utility maximizer, then his preference can be represented by a probability \Pr and a utility u . As observed in the solution to part (a), since u is unique only up to a positive affine transformation, we can assume without loss of generality that $u(0) = 0$ and $u(1) = 1$. With these utilities, the expected utility of C is $1/2$, the expected utility of A is $\Pr(H)p$ and the expected utility of B is $(1 - \Pr(H))q$. If $p = q = 1$, the expected utility of A is $\Pr(H)$ and the expected utility of B is $1 - \Pr(H)$. If C is preferred to both A and B , then we must have $1/2 > \Pr(H)$ and $1/2 > 1 - \Pr(H)$, which is impossible.

- (c) For the values of p and q and preferences in part (b) suppose that the decision maker acts as if he has a utility function on money u , a set of probabilities Π on the states, Hillary and not-Hillary, and evaluates each lottery using the element of Π that minimizes the expected utility of the lottery. What do you know about Π ?

Solution: If we take Π to be a set of probability measures that includes probability measures \Pr_1 and \Pr_2 such that $\Pr_1(H) > 1/2$ and $\Pr_2(H) < 1/2$, then A has minimum expected utility at most $\Pr_2(H) < 1/2$ and B has minimum expected utility at most $1 - \Pr_1(H) < 1/2$, while C has minimum expected utility $1/2$, so C is preferred to both A and B . To get this result it is necessary and sufficient that Π has two such probability measures (i.e., one

probability measure \Pr_1 such that $\Pr_1(H) < 1/2$ and one probability measures \Pr_2 such that $\Pr_2(H) > 1/2$).

8. (20 points) A decision maker has to allocate his initial wealth $w_0 > 0$ between two assets. One asset is risk free and pays a certain return of 1 for every dollar invested. The other is risky; its return per dollar invested is given by the random variable r which takes on values \bar{r} or \underline{r} . Let $\alpha \in [0, 1]$ be the fraction of wealth invested in the risky asset. So future wealth is $w = (1 - \alpha)w_0 + \alpha w_0 r$.

- (a) In this part of the question the decision maker is a von-Neumann Morgenstern expected utility maximizer with payoff function for future wealth $u(w)$. Assume that $u'(w) > 0$ and $u''(w) < 0$ for all w and that the expectation of r is greater than one. Show that the optimal choice of α is greater than 0.

Solution: Let p be the probability of getting \bar{r} , so $1 - p$ is the probability of getting \underline{r} . Assume without loss of generality that $0 < p < 1$. (If $p = 1$, then $r = \bar{r}$, and we must have $\bar{r} > 1$ since $E[r]$, the expected utility of r , is greater than 1; in that case, the investor should definitely invest in the risky asset, which isn't so risky. Similarly, if $p = 0$, the investor should definitely invest.) Since the expected utility of r is greater than 1, we must have $E[r] = p\bar{r} + (1 - p)\underline{r} > 1$. Let $U(\alpha)$ be the expected utility of the investment as a function of α . Then

$$U(\alpha) = pu((1 - \alpha)w_0 + \alpha w_0 \bar{r}) + (1 - p)u((1 - \alpha)w_0 + \alpha w_0 \underline{r}).$$

Since $u'(w) > 0$ and $u''(w) < 0$, we have $pu(x) + (1 - p)u(y) > u(px + (1 - p)y)$ for all $x \neq y$. In particular, it follows that if $\alpha > 0$, then

$$\begin{aligned} & pu((1 - \alpha)w_0 + \alpha w_0 \bar{r}) + (1 - p)u((1 - \alpha)w_0 + \alpha w_0 \underline{r}) \\ & > u(p(1 - \alpha)w_0 + \alpha w_0 \bar{r}) + (1 - p)u((1 - \alpha)w_0 + \alpha w_0 \underline{r}) \\ & = u((1 - \alpha)w_0 + \alpha w_0 (p\bar{r} + (1 - p)\underline{r})) \\ & = u((1 - \alpha)w_0 + \alpha w_0 E[r]) \\ & > u(w_0) \end{aligned}$$

The last line follows since $E[r] > 1$, so $(1 - \alpha)w_0 + \alpha w_0 E[r] > (1 - \alpha)w_0 + \alpha w_0 = w_0$, and $u'(w) > 0$. Thus, $U(\alpha) > U(0)$ for all $\alpha > 0$, so the optimal choice of α is greater than 0.

- (b) Now the distribution of r is uncertain. It depends on the state of nature which is either H or L. In state H, $r = \bar{r}$ with probability p and $r = \underline{r}$ with probability $1 - p$. In state L, $r = \bar{r}$ with probability q and $r = \underline{r}$ with probability $1 - q$. Suppose that for each α the decision maker's preferences over lotteries satisfy the Anscombe and Aumann axioms. Let π represent the decision maker's

subjective probability of state H and let $u(w)$ be his payoff function for future wealth. Write the decision maker's expected utility as a function of α .

Solution: Since the DM's preferences satisfy the Anscombe and Aumann axioms, we can represent the preference in terms of the subjective probability of states and the utilities on wealth. In particular, the DM's expected utility is

$$\begin{aligned} & U(\alpha, w_0) \\ &= \pi(pu((1-\alpha)w_0 + \alpha w_0 \bar{r}) + (1-p)u((1-\alpha)w_0 + \alpha w_0 \underline{r})) + \\ & \quad (1-\pi)qu((1-\alpha)w_0 + \alpha w_0 \bar{r}) + (1-q)u((1-\alpha)w_0 + \alpha w_0 \underline{r}). \end{aligned}$$

- (c) Continuing with the setup in part (b) suppose that $1 > \pi > 0$. Show that if we allow the payoff function to depend on the state, i.e. $U(w, H) \neq U(w, L)$, then the subjective probability of state H can be represented by any $\hat{\pi}$ with $1 > \hat{\pi} > 0$.

Solution: Given $\hat{\pi}$, we want to find $U(w, H)$ and $U(w, L)$ such that, for each possible amount of wealth w_0 , we have

$$\begin{aligned} & \pi(pu((1-\alpha)w_0 + \alpha w_0 \bar{r}) + (1-p)u((1-\alpha)w_0 + \alpha w_0 \underline{r})) + \\ & \quad (1-\pi)qu((1-\alpha)w_0 + \alpha w_0 \bar{r}) + (1-q)u((1-\alpha)w_0 + \alpha w_0 \underline{r}) = \\ & \hat{\pi}(pU(((1-\alpha)w_0 + \alpha w_0 \bar{r}, H) + (1-p)U(((1-\alpha)w_0 + \alpha w_0 \underline{r}, H) + \\ & \quad (1-\hat{\pi})qU(((1-\alpha)w_0 + \alpha w_0 \bar{r}, L) + (1-q)U(((1-\alpha)w_0 + \alpha w_0 \underline{r}, L)). \end{aligned}$$

This is easy: take $U(w, L) = 0$ for all w and set $U(w, H) = u(w)\pi/\hat{\pi}$.