

Networks and Markets

Homework 1

Due: February 13, 2024, 11:59pm

Submit solutions as a PDF file to Gradescope. On Gradescope, match pages with the corresponding problem (we will make a one-point deduction if pages are not matched). Show work throughout, with legible handwriting. Clearly mark your answer by putting a box around it.

1 Mixed Nash Equilibria (9 pts)

In this problem, we will consider a game with two players, Alice and Bob. The game involves a cake and a donut, which are both currently in Alice's possession. Alice has two possible actions: to protect the cake (C) or protect the donut (D). Bob has two possible actions: to try to steal the cake (C) or try to steal the donut (D). Bob successfully steals an item if Alice is not protecting it. If Bob steals the cake, Alice gets reward -3 and Bob gets reward 3 . If Bob steals the donut, Alice gets reward -1 and Bob gets reward 1 . Otherwise, if Bob does not successfully steal either item, Alice gets reward 2 and Bob gets reward -2 .

This game is given by the following payoff matrix:

		Bob	
		C	D
Alice	C	$2, -2$	$-1, 1$
	D	$-3, 3$	$2, -2$

Part (a) (2 points). For all four possible pairs of pure strategies, (C, C) , (C, D) , (D, C) , (D, D) , state which player would defect to play a different strategy. (Here, we list Alice's strategy followed by Bob's, so (C, D) corresponds to Alice protecting the cake and Bob trying to steal the donut.)

Part (b) (2 points). Provide a brief intuitive explanation for why there does not exist a pure Nash equilibrium in this game.

Part (c) (5 points). Compute a mixed Nash equilibrium. Give your answer in the form (p, q) , where p is the probability that Alice protects the cake, and q is the probability that Bob tries to steal the cake.

(See section 6.7 and 6.8 in <https://www.cs.cornell.edu/home/kleinber/networks-book/networks-book-ch06.pdf> for reference on calculating mixed equilibria.)

2 Bidding Strategy in Different Auctions (16 pts)

Suppose that Alice's valuation for an item is 2.5. There is one other bidder, Bob, who Alice believes will bid 1, 2, 3, or 4, each with probability $1/4$. Alice is deciding how much to bid, and can bid any real number. Assume that in a tie, Alice always wins the auction.

Part (a) (4 points) If the item is auctioned using a *second-price auction*, what is Alice's expected reward when she bids x ? (Write your answer as a simplified function of x .) What is one value of x that maximizes her expected reward?

Hint: The function you give in this part should be a piecewise function, and so your solution should be in a form similar to the below (note that the solution below is itself incorrect):

$$f(x) = \begin{cases} 5x + 2 & x < 3 \\ 0 & 3 \leq x < 50 \\ -12x & x \geq 50 \end{cases} \quad (1)$$

Part (b) (4 points) If the item is auctioned using a *first-price auction*, what is Alice's expected reward when she bids x ? (Write your answer as a simplified function of x .) What is one value of x that maximizes her expected reward?

Part (c) (4 points) Now suppose the item is auctioned using a *first-price auction* with a *reserve price* of 1.5. This means that the highest bid wins as long as it is at least as big as the reserve price. For example, if Alice bids 4 and Bob bids 3, Alice wins and pays 4 for the item; if Alice bids 1.4 and Bob bids 1, nobody wins the auction since the highest bid did not exceed 1.5.

What is Alice's expected reward when she bids x ? (Write your answer as a simplified function of x .) What is one value of x that maximizes her expected reward?

Part (d) (4 points) An *all-pay auction* is one in which the highest bidder receives the item, but *everyone* pays their own bid. For example, if Alice bid 2 and Bob bid 1, Alice receives the item and pays 2, and Bob pays 1 but receives nothing. If the item is auctioned using an *all-pay auction*, what is Alice's expected reward when she bids x ? (Write your answer as a simplified function of x .) What is one value of x that maximizes her expected reward?

3 Allocating Effort Towards Scientific Discovery (21 pts)

In this problem, we will consider a toy model of scientific discovery, in which two researchers, Alice and Bob, must each decide which of two scientific problems, problem 1 and problem 2, to work on. Each researcher can choose exactly one problem to work on.

Problem i has **reward** r_i and **success probability** p_i . This means that if a researcher chooses to work on problem i , they will find a solution with probability p_i (success will be independent of the other researcher's success on their chosen problem). If they find a solution, they receive payoff r_i , unless the other researcher also finds a solution to the same problem, in which case both researchers receive payoff $r_i/2$. (Therefore, if the researchers find solutions to different problems, then they each receive the full payoff for their chosen problem.)

A Basic Case

We will start with a simple case, in which the rewards and success probabilities for both problems are equal. Take $r_1 = r_2 = 8$ and $p_1 = p_2 = 1/2$. To analyze what happens in this case, we first need to write down a payoff matrix.

Part (a) (2 points). Consider what happens when both Alice and Bob choose problem 1. Then what is Alice's expected payoff?

(Hint: What is the probability that only Alice finds a solution? What is the probability Alice and Bob both find a solution?)

Part (b) (2 points). Now fill in the full payoff matrix with the expected payoffs.

		Bob	
		1	2
Alice	1	,	,
	2	,	,

Part (c) (2 points). Given the payoff matrix above, state all pure Nash equilibria.

Choosing Rewards

We now consider a more involved problem, where the goal is to *choose* the rewards for each problem to maximize societal benefit.

Suppose that the societal utility of solving problem 1 is $u_1 = 100$ and the societal utility of solving problem 2 is $u_2 = 60$. For example, if any researcher solves problem 1 but no researcher solves problem 2, society gets 100 utility. If one researcher solves problem 1 and the other solves problem 2, society gets 160 utility. Society does not get any extra utility if both researchers solve the same problem.

The success probabilities for each researcher is $p_1 = 0.8$ and $p_2 = 0.5$. In this case, problem 1 has more societal benefit, but is actually easier to solve than problem 2. Our goal is to choose the rewards r_1 and r_2 in order to maximize total societal utility.

Part (d) (2 points). First, compute the expected societal utility in the four possible allocations of Alice and Bob's effort:

- (i) Alice and Bob both work on problem 1.
- (ii) Alice works on problem 1 and Bob works on problem 2.
- (iii) Alice works on problem 2 and Bob works on problem 1.
- (iv) Alice and Bob both work on problem 2.

Which allocation(s) maximize expected societal utility?

Part (e) (4 points). One natural approach to choose rewards is to set rewards proportional to societal utility; indeed, it seems intuitive to reward solving problems that are of the greatest societal benefit.

Suppose that we set $r_1 = 100$ and $r_2 = 60$. Fill in the payoff matrix below for Alice and Bob in this case.

		Bob	
		1	2
Alice	1	,	,
	2	,	,

Part (f) (1 point). Given the payoff matrix above, state all pure Nash equilibria.

Part (g) (2 points). Does the choice of rewards result in the societal utility-maximizing allocation? Briefly summarize, intuitively, why it does or does not.

Part (h) (6 points). Keep $r_1 = 100$. Give an example of r_2 such that the resulting pure Nash equilibria maximize societal utility. Give the corresponding payoff matrix in the format below.

		Bob	
		1	2
Alice	1	,	,
	2	,	,

4 Comparing Matching Mechanisms (14 Pts)

Consider three workers A, B, C and three firms $1, 2, 3$ such that the workers have the following preferences:

$$1 \succ_A 2 \succ_A 3 \succ_A \emptyset \tag{2}$$

$$1 \succ_B 3 \succ_B 2 \succ_B \emptyset \tag{3}$$

$$3 \succ_C 1 \succ_C 2 \succ_C \emptyset \tag{4}$$

(In words, worker A prefers firm 1 to 2; 2 to 3; and 3 to being unemployed. Worker B prefers 1 to 3; 3 to 2; and 2 to being unemployed.)

Suppose that the firms all have the same preferences: $A \succ B \succ C \succ \emptyset$.

Part (a) (3 points) Compute the matching that arises from the deferred acceptance algorithm when workers propose. Assume that all workers and firms submit their true preferences. Clearly state which worker is matched to which firm.

Part (b) (2 points) Is the matching stable? If so, why? If it is not stable, provide an example of a blocking pair.

Part (c) (2 points) Would any worker have benefited from submitting a different preference list? If yes, list such a worker and an example of preference list they would have been better off submitting. If not, briefly explain why.

We now consider a different matching algorithm, called the immediate acceptance algorithm (also known as the Boston Mechanism). Like in the deferred acceptance algorithm, each worker begins by proposing to their most-preferred firm. If a firm receives at least one proposal, they *immediately* match with their most-preferred worker among those who proposed to the firm. (This is the key difference in comparison to the deferred acceptance algorithm.) Any worker who is not matched in the first round proposes again in the second round, this time to their second-favorite firm. The process continues until either every student is matched or has proposed to every firm.

Part (d) (3 points) Compute the matching that arises from the immediate acceptance algorithm when workers propose. Assume that all workers and firms submit their true preferences. Clearly state which worker is matched to which firm.

Part (e) (2 points) Is the matching given by the immediate acceptance algorithm stable? If it is not stable, provide an example of a blocking pair.

Part (f) (2 points) Under the Boston mechanism, would any worker have benefited from submitting a different preference list? If yes, list such a worker and an example of a preference list they would have been better off submitting. If not, briefly explain why.