Networks and Markets

Homework 2

Due: March 7, 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 per problem if pages are not matched). Show work throughout, with legible handwriting. Clearly mark your answer by putting a box around it.

1 The Median Mechanism (12 pts)

Consider the following problem. Several people each have a preference that can be represented on the real number line. For example, three people might have preferences $-1, 4$, and 0.5. Our goal is to aggregate these preferences in a way to choose a single outcome x, such that a person with true preference p receives utility $-|x - p|$. So a person receives higher utility if the outcome x is closer to their true preference. (For example, suppose x is our choice for how many quizzes we should have in class, and each of you has different preferences for this number).

Now consider three agents, Nikhil, Nikola, and Nick, who have true preferences $0, 1$, and 3 respectively.

Part (a) (2 pts) Suppose $x = 1$. What are Nikhil, Nikola, and Nick's three utilities? Answer as an ordered triple (Nikhil's utility, Nikola's utility, and Nick's utility).

Part (b) (3 pts) Define the social welfare as the sum of Nikhil's, Nikola's, and Nick's utilities. Give one outcome x that maximizes the social welfare. Show your work.

Part (c) (3 pts) Suppose that we do not know Nikhil's, Nikola's, and Nick's true utilities, but rather have them report their preferences. We need a way to aggregate these reported preferences into an outcome x. Consider the *mean mechanism*, where we choose x to be the mean (average) of the three reported preferences. Show that reporting truthfully is not a Nash equilibrium. Hint: suppose that two agents report the truth. What does the third agent want to do?

Part (d) (4 pts) Instead consider the *median mechanism.* Show that reporting truthfully is a Nash equilibrium. *Hint:* In fact, one can show that reporting truthfully is a *dominant strategy*, meaning that the median mechanism is strategyproof—no matter what other agents do, each agent cannot benefit by lying.

2 Tie Breaking in School Matching (20 pts)

In this problem, we will consider a significant policy problem that arose in the design of school matching: the use of a single tie breaking (STB) or multiple tie breaking (MTB). Preferences of schools over students are often determined by lottery—the question was if a single lottery number should be used to determine a student's priority at all schools, or if a different lottery number should be used at each school.

Intuitively, many parents and students felt that STB was unfair: a single bad lottery number would mean that a student would have a low priority at *all* schools. As it turns out, however, STB actually matches more students to their top choices, and is generally considered better. Consequently, New York City and many other school districts, use STB.

To analyze the use of STB v. MTB, we introduce a model with a *continuum* of students. By considering an infinite number of students, it is actually easier to analyze deferred acceptance.

There is a continuum of students of total mass 1 and two schools each with capacity 1/4. Intuitively, this means that each school has capacity for a quarter of the total amount of students. Each student has a probability 1/2 of preferring school 1 and a probability 1/2 of preferring school 2.

Each student also has a random pair of lottery numbers (ℓ_1, ℓ_2), where $\ell_1, \ell_2 \in (0, 1)$. (We will specify this randomness later.) Now suppose schools 1 and 2 have **cutoffs** P_1 and P_2. Then we say that a student can **afford** school i if and only if $\ell_i > P_i$, i.e., if and only if their lottery number exceeds the school's cutoff. For example, if $P_1 = 0.5$ and $P_2 = 0.7$, then a student with lottery numbers $(0.3, 0.9)$ can afford school 2 but not school 1.

Given a choice of cutoffs, a student is matched to their most preferred school among the schools they can afford. For example, even if the student above prefers school 1, they are matched to school 2 since that is the only school they can afford.

Cutoffs P_1 and P_2 are **market-clearing** if and only if the probabilities a student is matched to school 1 and school 2 are each 1/4. Intuitively, this means that the right number of students is matched to each school. If cutoffs are market-clearing, then the corresponding matches exactly correspond to the outcome of deferred acceptance. *Side note: this is non-trivial to show.*

We will now walk through how to compute market-clearing cutoffs in one case. Consider a single student. A single lottery number is used at both schools, meaning that $(\ell_1, \ell_2) = (\ell, \ell)$ where ℓ is drawn uniformly at random from $(0, 1)$.

Part (a) (1 points). Consider cutoffs $P_1 = P_2 = 1/4$. What is the probability that the student can afford school 1? *Hint: In other words, what is the probability that the student's lottery number is higher than P_1?*

Part (b) (2 points). Consider cutoffs $P_1 = P_2 = 1/4$. What is the probability that the student matches with school 1? *Hint: What is the difference between affording a school versus matching to it?* Briefly explain.

Part (c) (1 points). Briefly explain why the cutoffs $P_1 = P_2 = 1/4$ are not market-clearing.

Part (d) (3 points). There exists P such that the cutoffs $P_1 = P_2 = P$ are market-clearing. What is P? Verify that the probabilities the student is matched with schools 1 and 2 are each 1/4.

Part (e) (2 points). Given the market-clearing cutoffs above, what is the probability that the student matches with their most preferred school?

Now consider the same setup, except now a different lottery number is used at each school for the student, meaning that ℓ_1 and ℓ_2 are each drawn independently and uniformly at random from $(0, 1)$.

Part (f) (1 points). Consider cutoffs $P_1 = P_2 = 1/4$. What is the probability that the student can afford school 1 but not school 2?

Part (g) (3 points). Consider cutoffs $P_1 = P_2 = 1/4$. What is the probability that the student matches with school 1? Briefly explain.

Part (h) (3 points). There exists P such that the cutoffs $P_1 = P_2 = P$ are market-clearing. What is P? Verify that the probabilities the student is matched with schools 1 and 2 are each $1/4$.

Part (i) (2 points). Given the market-clearing cutoffs above, what is the probability that the student matches with their most preferred school?

Part (j) (2 points). Compare your answers to (e) and (h). What does this comparison reveal about the choice between STB and MTB in school matching?

3 Balance Theory (15 points)

There is an ancient proverb: "The enemy of my enemy is my friend." This proverb conveys a basic principle, that two people who share a mutual enemy are likely to be friends. (In some sense, the opposite of the idea behind *triadic closure*, that I am likely to be friends with my friend's friends).

This principle can be used to analyze positive and negative relationships on networks, which we will do in this problem.

Consider a network G and label each of its edges with a $+$ (indicating a positive relationship) or a $-$ (indicating a negative relationship). A triangle in the network (three nodes that are pairwise connected) is *unbalanced* if and only if it contains 1 or 3 negative relationships. (In the former case, it will thus have 2 positive relationships; in the latter, it will have 0 positive relationship).

Part (a) (2 points). Explain why a triangle with 3 negative relationships violates the "enemy of my enemy is my friend" principle.

Part (b) (2 points). A triangle with 1 negative relationship violates a slightly different property: "the ___ of my ___ is my ___." What goes in the blanks? *Hint: there are multiple correct answers; you only need to give one.*

Roughly speaking, if a network is unbalanced, this means that the set of friends (pairs of nodes with positive relationships) and enemies (pairs of nodes with negative relationships) is not stable. We say that a network is *balanced* if and only if it contains *no* unbalanced triangles. We now use the concept of balance to understand the stability of a complete graph (a network where all pairs of nodes are connected by an edge) with 100 nodes labeled $1, 2, \cdots, 100$.

Part (c) (2 points). Suppose that all nodes in the network are enemies with each other. Is this network balanced? Explain.

Part (d) (3 points). Now suppose that nodes $1 - 50$ are part of political party A and nodes $51 - 100$ are part of political party B. Nodes are friends with everyone in their own political party and enemies with everyone in the opposing political party. Is this network balanced? Explain.

Part (e) (3 points.) Consider the network in part (d). Is it possible to change exactly one relationship between two nodes such that the resulting network is balanced (whether or not the answer to part (d) is that the network is currently balanced)? Explain.

Part (f) (3 points). Now suppose that nodes $1 - 33$ are in party A, $34 - 66$ are in party B, and $67 - 100$ are in party C. Again, nodes are friends with everyone in their own political party and enemies with everyone in the opposing political party. Is this network balanced? Explain.

Part (g) (3 points). Now suppose that there are 10 political issues, labeled $0, 1, 2, \cdots, 9$. Now a node cares about a political issue if and only if they share a digit. So node 23 cares about issues 2 and 3, node 7 cares about only issue 7, and node 100 cares about issues 0 and 1. Two nodes are friends if they care about at least one issue in common, and are enemies if they do not care about any of the same issues. Is this network balanced? Explain.

4 Network Effects (18 pts)

In this problem, we consider a simple model of network effects. The basic idea behind network effects is that for some products, the fact that many people use the product increases its value. For example, if many of your friends use WhatsApp to communicate, this directly increases the value of WhatsApp for you. Somewhat more indirectly, if many people use Uber, this increases the value of Uber since more drivers and riders make the platform more efficient. Other goods like shoes can also experience network effects: if many of your classmates are wearing a particular shoe, that may make the shoe more fashionable, and hence, more valuable.

In our model, there is an underlying graph G of users. There are two products A and B. At time $t = 0$, everybody uses product B except for a set of *initial adopters*, who use product A. Then, at each subsequent time step, if at least a p proportion of a user's neighbors use product A, then that user will switch to product A. This continues until no further changes occur. Intuitively, a higher p means that a user will only want to switch if many of their neighbors adopt product A.

Now consider the following network of eight users, labeled 1 through 8.

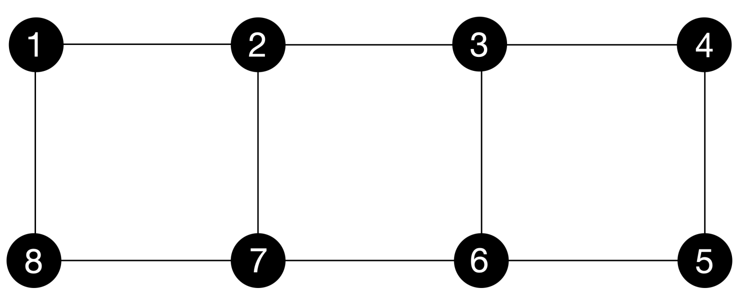

Part (a) (3 points). Suppose that the set of initial adopters is $\{1, 2\}$ and $p = 0.1$. Which users end up adopting product A?

Part (b) (3 points). Suppose that the set of initial adopters is $\{1, 2\}$ and $p = 0.5$. Which users end up adopting product A?

Part (c) (3 points). Suppose that $p = 0.5$. Give a set of two initial adopters such that all users in the network end up adopting product A.

Part (d) (3 points). Suppose that $p = 0.5$. Give a set of four initial adopters such that no additional users adopt product A.

Part (e) (3 points). For which values of p is the following true: any choice of a single initial adopter will result in every user eventually adopting product A?