Networks and Markets

Homework 5

Due: 05/05/2026, 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.

As always, you are allowed to use computational tools to perform algebra/calculus, as long as you clearly show the algebra setup and explain how you calculated it.

1 Strategic voting, and Voting for a committee (18 points)

This problem explores two aspects of voting: (a) that voters may benefit from being strategic; (b) electing a set of multiple winners can be far more complicated than standard majority voting. In many ways, this homework problem is similar to the problem on diversity in recommendations.

Suppose we have 7 voters (labeled 1 - 7) and 4 candidates (labeled A - D), and we want to elect a committee of either one or two winners (in different subproblems, we want to declare one or two of the four candidates as winners). Suppose the voters have the following preferences over candidates:

$$1 : A \succ B \succ \emptyset$$
$$2 : A \succ B \succ \emptyset$$
$$3 : B \succ A \succ \emptyset$$
$$4 : B \succ A \succ \emptyset$$

$$5 : C \succ D \succ \emptyset$$
$$6 : C \succ D \succ \emptyset$$
$$7 : C \succ D \succ \emptyset$$

Intuitively, the first 4 voters agree on the top two candidates, and the final 3 voters agree on a different set of top 2 candidates, and no voter likes any candidate besides their top two (they do not feel **represented** by them).

In voting, there are typically two components to define a mechanism: (1) **elicitation**, how do we *ask* voters what they want? For example, do we ask them for a full ranking, for their favorite 2 candidates, numeric scores, etc? (2) **aggregation**, how do we *count* votes? For example, do we just count the total number of (first place) votes each candidate receives, give points for first place versus second place, etc? In this problem, we'll explore how different choices of elicitation and aggregation mechanisms might have different implications for who is selected as a winner.

All parts in this problem are 2 points each.

First, we'll analyze the standard election setup, where each voter votes for one candidate, and there is only 1 winner. For the following problems, we'll have: **elicitation:** Each voter votes for 1 candidate. **aggregation:** we simply count how many votes each candidate receives. The winner is the candidate with the most votes.

Part (a) With the above elicitation and aggregation mechanisms, who is the winner, if every voter honestly gives the system their favorite candidate?

Part (b) Let's suppose that a voter feels *represented* by a committee if it has *at least one* member who represents them. How many voters feel *represented* by the elected winner in part (a)? (Recall that a voter feels represented by a winner if that winner appears in their top two choices.)

Part (c) We say that a voter is *strategic* if they lie about who their favorite candidate is. Is there a group of voters (a "coalition") who can all be better represented (i.e., there are more candidates in the winning set that represent them) if they are strategic together? How can they vote such that everyone in the coalition is better off? *Note: the answer may not be unique.*

Second, now suppose we are electing 2 winners. For the following problems, we'll have: **elicitation:** Each voter votes for their favorite 2 candidates (without telling us who is number one or two; a voter also cannot vote for the same candidate twice). **aggregation:** we simply count how many votes each candidate receives. The two winner are those candidates with the most votes.

Part (d) With the above elicitation and aggregation mechanisms, who are the two winners, if every voter honestly gives the system their favorite 2 candidates? *Note: in event of a tie, give all possible sets of winners.*

Part (e) Let's suppose that a voter feels *represented* by a committee if it has *at least one* member who represents them. How many voters feel *represented* by the elected winners in part (d)?

Part (f) We say that a voter is *strategic* if they lie about who their favorite two candidates are. Is there a group of voters (a "coalition") who can all be better represented if they are strategic together? How can they vote such that everyone in the coalition is better off? *Note: the answer may not be unique*

Third, we're still electing 2 winners, but we use a fancier mechanism. For the following problems, we'll have: **elicitation:** Each voter votes for their favorite 2 candidates (without telling us who is number one or two). **aggregation:** We use what is called a Thiele rule or *proportional approval voting.*

Consider a set of winners W. For each voter v, suppose they are represented by $r_v(W)$ candidates in set W. Let $f(r_v(W))$ be how many "points" set W gets from the voter, where $f(0) = 0$ and otherwise,

$$f(r) = \sum_{k=1}^{r} \frac{1}{k} = 1 + \frac{1}{2} + \cdots + \frac{1}{r}. \tag{1}$$

We select the winning set W with the most points, i.e., we select W, where $|W| = 2$, that maximizes

$$\sum_{v \in \text{voters}} f(r_v(W)). \tag{2}$$

Intuitively, there are diminishing returns for a voter: they contribute 1 point if they like 1 candidate in a set, but only 1.5 points if they like 2, and 1.83333 if they like 3, etc. Many multiwinner elections work in a qualitatively similar fashion of diminishing returns, like single transferable voting.

Part (g) With the above elicitation and aggregation mechanisms, who are the two winners, if every voter honestly gives the system their favorite 2 candidates? *Note: in event of a tie, give all possible sets of winners.*

Part (h) Let's suppose that a voter feels *represented* by a committee if it has *at least one* member who represents them. How many voters feel *represented* by the elected winners in part (g) (in event of a tie in part (g), give answer for each answer)?

Part (i) We say that a voter is *strategic* if they lie about who their favorite two candidates are. Is there a group of voters (a "coalition") who can all be better represented if they are strategic together? How can they vote such that everyone in the coalition is better off? *Note: the answer may not be unique*

2 Guessing the number game (14 points)

Let's revisit the game that we've been playing in class all semester. Recall that in this game, you have a group (e.g., of students) who each have to guess an integer in $\{1, 2, \ldots 100\}$. The winner(s) of the game are those who guess closest to $\frac{2}{3}$ of the mean of the guessed numbers. For example, if the mean is 99, then the people who guess closest to 66 would be the winners. In this problem, we'll assume that there are a large number of players, of mass 1 (in other words, an agent doesn't have to consider how their own guess changes the class average; they just have to think about the class average not including their own guess). This is an iterated reasoning game.

Part (a) (2 points). What is the Nash Equilibrium of the game? In other words, in equilibrium, what is the number that everyone will play if everyone is being strategic?

Part (b) (2 points). As we've seen in class, we don't actually reach the Nash Equilibrium because many students deviate from it, for some reason. This reflects real life, when the Nash Equilibria assuming everyone is fully rational is rarely the outcome that is reached. In the rest of this quiz problem, we'll explore different ways that we can model this situation. In particular, we're going to study a model of a K-level strategic behavior. A basic K-level model is as follows: an agent plays like a K-level agent if they play an action that is optimal assuming all their opponents play like $(K-1)$-level agents. A 0-level agent plays uniformly at random (an integer in $\{1, 2, \ldots 100\}$).

As a warmup, we ask: what number does a 1-level agent play? *In other words, what action would be optimal if all opponents played like a* 0*-level agent?*

Part (c) (2 points). What number does a 3-level agent play? Show your work.

Part (d) (2 points). What number does a K-level agent play, when $K \to \infty$?

Part (e) (2 points). Suppose you don't think all your opponents play as a K-level agent, for any fixed K. Rather, you think that half your opponents play as 1-level, a quarter as 2-level, and the remaining as 4-level. What number should you play? Show your work.

Part (f) (2 points). It seems as if the class average has settled to be around 20. Suppose I told you that a fraction a of the class was playing like a 0-level agent, and a fraction $1 - a$ was playing like an ∞-level agent. What is a? Show your work.

Part (g) (2 points). In the previous question, suppose I gave you a class average m instead of $m = 20$ specificaly. For what m would the question in (f) be impossible to answer (i.e., there wouldn't exist such an a). Show your work.

3 Crowd-Sourced Fact Checking (13 points)

In this problem, we will consider a problem motivated by the Community Notes program on X (Twitter) that allows users to add context to misleading posts. Two examples are shown below:

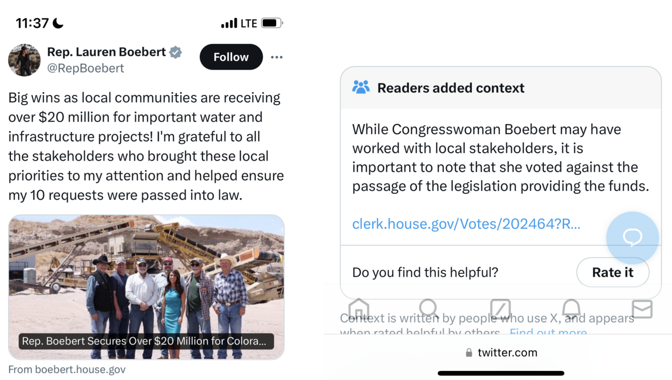

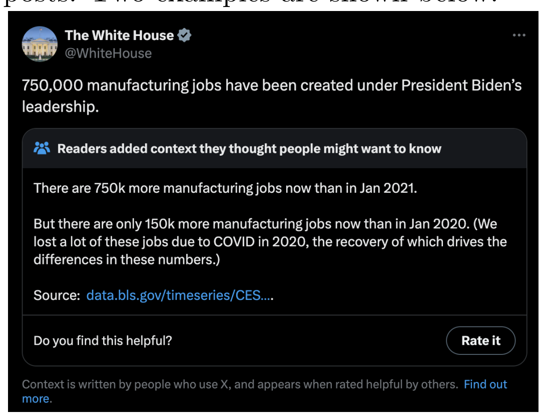

The key idea of Community Notes is to identify pieces of content ("notes" that provide context) that satisfy a "bridging-based" objective—i.e., content that is viewed as helpful across, say, the political spectrum. This avoids showing content that is "too biased." For example, a piece of content that is liked by 80 percent of people from group A, but only 20 percent of people from group B is (intuitively) not very bridging, but a piece of content that is liked by 70 percent of people from both groups is (intuitively) more bridging. Here, we will consider the task of identifying bridging-based content from a voting perspective.

Imagine that there are only two groups of people, A and B. Suppose that n_A and n_B people from groups A and B vote on a piece of content c. We'll assume that of the n_A people from group A, a $p_A(c)$ proportion vote to show content c. Similarly, of the n_B people from group B, a $p_B(c)$ proportion vote to show content c. For simplicity, we allow there to be a fractional number of people who like a piece of content.

We'll consider a simple voting method where content is shown if and only if at least a t proportion of all votes are in favor of showing the content. Specifically, when

$$p_A(c) \cdot n_A + p_B(c) \cdot n_B \geq t \cdot (n_A + n_B) \tag{3}$$

for some $t \in [0, 1]$. Intuitively, we would like to choose a low threshold t, as to be able to show more content, but also must choose a high enough threshold to ensure that content shown satisfies a bridging property.

Part (a) (2 points). Suppose that $n_A = n_B$. Suppose that we would like to show content such that $p_A(c) \geq 0.3$ and $p_B(c) \geq 0.3$. What is the smallest choice of t such that *all* shown content satisfies this property? Explain your answer.

Part (b) (2 points). Suppose that $n_A = n_B$. Suppose that we would like to show content such that $p_A(c) \geq p$ and $p_B(c) \geq p$. What is the smallest choice of t such that *all* shown content satisfies this property? Write your answer as a function of p, considering all values $p = 0$ through $p = 1$. Explain your answer.

Part (c) (2 points). Use your answer in part (b) to explain how the choice of p (the amount of bridging), affects the threshold required.

Part (d) (2 points). Suppose now that n_A and n_B can take on any values. Suppose that we would like to show content such that $p_A(c) \geq 0.3$ and $p_B(c) \geq 0.3$. What is the smallest choice of t such that *all* shown content satisfies this property? Explain your answer.

Part (e) (3 points). Suppose now that n_A and n_B can take on any values, as long as $\frac{1}{c} \cdot n_B < n_A < c \cdot n_B$. Suppose that we would like to show content such that $p_A(c) \geq p$ and $p_B(c) \geq p$. For $c \geq 1$, what is the smallest choice of t that guarantees that *all* shown content satisfies this property? Write your answer as a function of c and p, and explain your answer.

Part (f) (2 points). Use your answer in part (e) to explain how differences in the number of voters in each group changes the threshold required to maintain bridging.

4 Correlated Judges (13 points)

You are evaluating outputs from a language model and want to use "LLM-as-judge" to label them. To improve reliability, you aggregate the verdicts of 3 judges on a binary question (correct / incorrect) via majority vote. Each individual judge is correct with probability $p \in [0, 1]$. Let $V_i \in \{0, 1\}$ denote the event that judge i votes correctly, and let M denote the event that the majority vote is correct.

Part (a) (3 points) Suppose first that the three judges opinions' are independent. What is the probability that $M = 1$?

Part (b) (3 points) Now, suppose that the first judge has some influence on the two others. Independently, with probability $\alpha \in [0, 1]$, each one of them aligns its judgment with the first judge instead of formulating an independent one. What is the probability that $M = 1$?

Part (c) (5 points) Now, suppose instead that there are $2k + 1$ independent judges for $k \in \mathbb{N}$. Denote as m_k the probability that the majority vote is correct. Suppose that $p > \frac{1}{2}$ (hint: this is equivalent to $p > 1 - p$). Show that $m_{k+1} > m_k$.

Part (d) (2 points) Suppose that k is very large, such that with independent judges there is only a negligible risk of error. We consider a general version of the scenario from part (b). Now, all judges labeled from 2 to $2k+1$ align independently with probability α their opinion with that of the first judge. Explain why if $2p > (1-\alpha)^{-1}$ and k is large enough, the risk of error remains negligible (hint: consider the two vote options for the first judge and reason about the expected number of correct votes in each case).