

Optimized Democracy (Fall 2026)

Problem Set #2

Due: 10/9/2026 11:59pm ET

Instructions:

- You must not use AI in any way.
- You must not use AI *in any way*.
- *You must not use AI in any way!!!*
- You may discuss the problems with classmates but please write down solutions completely on your own.
- The solutions to many of the problems that we give can be found in papers, but, needless to say, you should avoid reading the proof if you come across the relevant paper. If for some reason you did see the solution before working it out yourself, please say so in your solution.
- Please *type up* your solution and submit to Gradescope.

Problems:

1. **Note:** The following problem is identical to one from the spring 2026 edition of CS 1360. If you've taken that course, you may submit the same solution.

We saw in class a proof sketch of the Gibbard-Satterthwaite Theorem for the special case of strategyproof and neutral voting rules with $m \geq 3$ and $m \geq n$. That proof relied on two key lemmas. In this problem, you will prove the two lemmas and formalize the theorem's proof for this special case.

Prove the following statements.

- (a) **[5 points]** Let f be a strategyproof voting rule, $\sigma = (\sigma_1, \dots, \sigma_n)$ be a preference profile, and $f(\sigma) = a$. If σ' is a profile such that $[a \succ_{\sigma_i} x \Rightarrow a \succ_{\sigma'_i} x]$ for all $x \in A$ and $i \in N$, then $f(\sigma') = a$.
- (b) **[5 points]** Let f be a strategyproof and onto voting rule. Furthermore, let $\sigma = (\sigma_1, \dots, \sigma_n)$ be a preference profile and $a, b \in A$ such that $a \succ_{\sigma_i} b$ for all $i \in N$. Then $f(\sigma) \neq b$.

Hint: use part (a).

- (c) **[10 points]** Let m be the number of alternatives and n be the number of voters, and assume that $m \geq 3$ and $m \geq n$. Furthermore, let f be a strategyproof and neutral voting rule. Then f is dictatorial.

Note: There are many proofs of the full version of the Gibbard-Satterthwaite Theorem; here the task is specifically to formalize the proof sketch we did in class.

2. Consider a facility location game with n agents in which each agent controls k locations, and denote the set of locations that agent i controls by $\mathbf{x}_i = (x_{i1}, \dots, x_{ik})$. Therefore, the entire location profile is $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$.

Let a deterministic mechanism in the multiple locations setting be defined as a function $f : \mathbb{R}^k \times \dots \times \mathbb{R}^k \rightarrow \mathbb{R}$; that is, it takes in a location profile and returns a single location based on all the locations reported by each agent.

The cost of facility location y to an agent i is the sum of distances from y to each of the locations that i controls, or $\text{cost}_i(y, \mathbf{x}_i) = \sum_{j \in [k]} |y - x_{ij}|$. The social cost of a location y is the sum of costs of each agent for location y :

$$\text{cost}(y, \mathbf{x}) = \sum_{i \in [n]} \sum_{j \in [k]} |y - x_{ij}|.$$

Consider the following mechanism for the facility location game in the multiple locations setting.

MECHANISM 1

- For each agent i with reported locations $\mathbf{x}_i = (x_{i1}, \dots, x_{ik})$, let $\text{med}(\mathbf{x}_i)$ be the median of these locations.
- Return the median of $(\text{med}(\mathbf{x}_1), \dots, \text{med}(\mathbf{x}_n))$.

Intuitively, Mechanism 1 creates a new bid for each agent at the median of the locations under its control, and then returns the median of these new bids. When n is even the median refers to the $n/2$ order statistic, but below you may assume that both n and k are odd when it simplifies the proof.

- (a) **[5 points]** Prove that Mechanism 1 is strategyproof.
- (b) **[20 points]** Prove that Mechanism 1 is a 3-approximation algorithm for the social cost in the multiple locations setting.
- (c) **[15 points]** Consider the case of two agents. Prove that for any $\varepsilon > 0$, there exists a k such that any strategyproof deterministic mechanism $f : \mathbb{R}^k \times \mathbb{R}^k \rightarrow \mathbb{R}$ cannot have an approximation ratio better than $3 - \varepsilon$ for the social cost in the multiple locations setting.
- Hint:** After choosing k , first prove that for any strategyproof mechanism f , there must exist distinct locations $a, b \in \mathbb{R}$ such that

$$f(\underbrace{(a, \dots, a)}_k, \underbrace{(b, \dots, b)}_k) \in \{a, b\}.$$

(Note that you must prove this; you cannot assume it is true.)

3. This problem deals with the Hotelling model with policy-motivated candidates (slides 8–9 of the “electoral competition” lecture). We showed (informally) that if $x_1^* < m < x_2^*$ then (m, m) is the unique Nash equilibrium; this is more generally true when $x_1^* \leq m \leq x_2^*$. Our goal is to examine the (almost) complement case of $x_1^* < x_2^* < m$, where (m, m) is no longer the unique Nash equilibrium.

To avoid any ambiguity, let us make the following simplifying assumptions. As before, there are two candidates. The distribution of voters is the uniform distribution over $[0, 1]$, so $m = 1/2$. For a winning position x_j , the cost of candidate i is $|x_i^* - x_j|$, and if there is a tie between the two candidate positions x_1 and x_2 then the cost of candidate i is $\frac{1}{2}(|x_i^* - x_1| + |x_i^* - x_2|)$.

[20 points] Assuming that $x_1^* < x_2^* < 1/2$, prove that (x_1, x_2) is a Nash equilibrium if and only if $x_2^* < x_2 = x_1 \leq 1/2$ or $(x_2 = x_2^*$ and $x_1 \leq x_2)$ or $(x_2 = x_2^*$ and $x_1 > 1 - x_2^*)$.

Note: Please prove both directions.

4. In class we discussed the Mallows model, which gives an expression for the probability of a ranking σ given the ground truth π . So computing the probability of a given ranking is easy, but how can we sample from this distribution?

Assume that $a_1 \succ_\pi a_2 \succ_\pi \dots \succ_\pi a_m$, and consider the following generative model, defined by probabilities p_{ij} for all $i = 1, \dots, m$ and $j = 1, \dots, i$, which iteratively constructs the ranking σ . In round 1, a_1 is inserted into the first (and only) position of the constructed ranking with probability $p_{11} = 1$. In round 2, a_2 is inserted into position 1 (above a_1) with probability p_{21} and into position 2 (below a_1) with probability p_{22} . More generally, in round i , for each $j = 1, \dots, i$, a_i is inserted into position j with probability p_{ij} .

[20 points] Prove that the Mallows Model with parameter ϕ is equivalent to this generative model with $p_{ij} = \phi^{i-j} \frac{1-\phi}{1-\phi^i}$. (This means that sampling rankings from the Mallows model is indeed easy.)

Hint: You may use the fact that for all $\pi \in \mathcal{L}$,

$$(1 + \phi)(1 + \phi + \phi^2) \dots (1 + \phi + \dots + \phi^{m-1}) = \sum_{\tau \in \mathcal{L}} \phi^{d_{KT}(\tau, \pi)}.$$