Optimized Democracy (Fall 2025)
Problem Set #2
— Solutions —

Due: 10/8/2025 11:59pm ET

Instructions:

- 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.
- You must not use AI in any way.
- Please *type up* your solution and submit to Gradescope.

Problems:

1. **Note:** The following problem is identical to one from the spring 2025 edition of CS 1360. If you've taken that course, simply say so (without solving the problem) and you'll get full credit.

   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, $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ be a preference profile, and $f(\boldsymbol{\sigma}) = a$. If $\boldsymbol{\sigma'}$ 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 $f(\boldsymbol{\sigma'}) = a$.

   Solution: The proofs of all three parts are copied from Svensson [1]. Suppose first that $\sigma_i = \sigma'_i$ for $i > 1$. Let $f(\succ_{\sigma'_1}, \succ_{\sigma'_{-1}}) = b$. From strategyproofness it follows that $a \succeq_{\sigma_1} b$, and hence from the assumption of the lemma, $a \succeq_{\sigma'_1} b$. Strategyproofness also implies that $b \succeq_{\sigma'_1} a$, and because preferences are strict it follows that $a = b$. The lemma now follows after repeating this argument while changing the preferences for only $i = 2$, then $i = 3$, etc.

(b) **[5 points]** Let f be a strategyproof and onto voting rule. Furthermore, let $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ be a preference profile and $a, b \in A$ such that $a \succ_{\sigma_i} b$ for all $i \in N$. Then $f(\boldsymbol{\sigma}) \neq b$.

Hint: use part (a).

Solution: Suppose that $f(\boldsymbol{\sigma}) = b$. Since f is onto there is a profile $\boldsymbol{\sigma}'$ such that $f(\boldsymbol{\sigma}') = a$. Let $\boldsymbol{\sigma}''$ be such that for all $i \in N$, $a \succ_{\sigma_i''} b \succ_{\sigma_i''} x$ for all $x \in A \setminus \{a, b\}$, and the rest of the alternatives are ranked identically to σ_i. By strong monotonicity (part (a)), $b = f(\boldsymbol{\sigma}) = f(\boldsymbol{\sigma}'')$ and $a = f(\boldsymbol{\sigma}') = f(\boldsymbol{\sigma}'')$, which is a contradiction. Hence $f(\boldsymbol{\sigma}) \neq b$.

(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.*

Solution: For this part of the proof it is convenient to define the preferences of each $i \in N$ via a utility function u_i such that for $x, y \in A$, $x \succ_{\sigma_i} y$ if and only if $u_i(x) > u_i(y)$. Therefore, $f(\boldsymbol{u})$ is well defined. We will also denote $A = \{a_1, \ldots, a_m\}$.

For each $i \in N$, let

$$u_i(a_j) = \begin{cases} n + i - j & i \leq j \leq n \\ i - j & j < i \\ n - j & j > n \end{cases}$$

That is, the ranking of $a_1, \ldots, a_n$ is shifted, and all other alternatives are ranked below them. By Pareto optimality (part b), $f(\boldsymbol{u}) = a_j$ for some $j \leq n$. Assume w.l.o.g. that $f(\boldsymbol{u}) = a_1$. Let $\boldsymbol{u}'$ be defined as follows:

$$\begin{aligned} &u_1'(a_1) = n + 2 \text{ and } u_1'(a_n) = n + 1, \\ &u_i'(a_n) = n + 2 \text{ and } u_i'(a_1) = n + 1 \text{ for } i > 1, \\ &u_i'(a_j) = u_i(a_j) \text{ otherwise} \end{aligned}$$

Hence all voters consider the alternatives a_1 and a_n to be better than the other alternatives. Also note that the ranking of a_1 and a_n is the same in the profiles $\boldsymbol{u}$ and $\boldsymbol{u}'$; and in $\boldsymbol{u}'$, a_1 and a_n are both ranked above other alternatives. Hence by strong monotonicity (part (a)), $f(\boldsymbol{u}') = f(\boldsymbol{u}) = a_1$.

Finally, define profiles $\boldsymbol{u}^k$ for $k = 1, \ldots, n$, where $\boldsymbol{u}^1 = \boldsymbol{u}'$, and

$$u_i^{k+1}(x) = \begin{cases} u_i^k(x) & i \neq k + 1 \\ u_{k+1}^k(x) & i = k + 1 \text{ and } x \in A \setminus \{a_1\} \\ -m & i = k + 1 \text{ and } x = a_1 \end{cases}$$

By Pareto optimality (part (b)), $f(\boldsymbol{u}^k) \in \{a_1, a_n\}$. But strategyproofness implies that $f(\boldsymbol{u}^k) = a_1$, and hence $f(\boldsymbol{u}^n) = a_1$. In $\boldsymbol{u}^n$, a_1 is ranked at the top by voter 1, and at the bottom by every other voter. Monotonicity (part (a)) implies that a_1 is the winner whenever voter 1 puts a_1 at the top. Neutrality then implies that voter 1 is a dictator.

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 $\boldsymbol{x}_i = (x_{i1}, \ldots, x_{ik})$. Therefore, the entire location profile is $\boldsymbol{x} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$.

   Let a deterministic mechanism in the multiple locations setting be defined as a function $f : \mathbb{R}^k \times \cdots \times \mathbb{R}^k \to \mathbb{R}$; 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 $\text{cost}_i(y, \boldsymbol{x}_i) = \sum_{j \in [k]} |y - x_{ij}|$. The social cost of a location y is the sum of costs of each agent for location y:

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

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

   > **MECHANISM 1**
   >
   > - For each agent i with reported locations $\boldsymbol{x}_i = (x_{i1}, \ldots, x_{ik})$, let $\text{med}(\boldsymbol{x}_i)$ be the median of these locations.
   > - Return the median of $(\text{med}(\boldsymbol{x}_1), \ldots, \text{med}(\boldsymbol{x}_n))$.

   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.

   Solution: First, we show that each agent i has single peaked preferences with peak located at their median $\text{med}(\boldsymbol{x}_i)$. Indeed, consider two points y_1, y_2 with $\text{med}(\boldsymbol{x}_i) \leq y_1 < y_2$ ($y_2 < y_1 \leq \text{med}(\boldsymbol{x}_i)$ is symmetric). Let $d = y_2 - y_1$ be the distance moved to the right. For simplicity, we suppose that k is odd. When moving the facility from y_1 to y_2, the distance from all $(k+1)/2$ points located at or to the left of the median has increased by d, while the remaining points $(k-1)/2$ could have distance decreased by at most d. Hence, $\text{cost}_i(y_2, \boldsymbol{x}_i) > \text{cost}_i(y_1, \boldsymbol{x}_i)$, implying single peakedness.

   Hence, Mechanism 1 is simply taking the median of agent peaks, which, as argued in the slides, implies it is strategyproof.

   (b) **[20 points]** Prove that Mechanism 1 is a 3-approximation algorithm for the social cost in the multiple locations setting.

   Solution: Fix preferences $\boldsymbol{x} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$. Let the true optimum location be a^* and the location returned by the mechanism be a.

   First, we show that the number of total locations at or below a is at least $nk/4$ (and, symmetrically, that the number of total locations at or above a is at least $nk/4$). The proofs are symmetric; we only provide the former below.

Let $\{x_{ij} : x_{ij} \le a\}$ be the set of all locations that are to the left of a. Furthermore, let $\tilde{x}_i$ denote the median of agent i. Because Mechanism 1 returns the median of the $\tilde{x}_i$ values, we know that at least half of the individual medians are below a. Further, we know that at least half of each of these agents' locations are below their median, so we have $|\{x_{ij} : x_{ij} \le a\}| \ge nk/4$. Symmetrically, we know that $|\{x_{ij} : x_{ij} \ge a\}| \ge nk/4$.

Now, we look at the social cost of the facility location that the mechanism returns. If $a = a^*$, then $\mathrm{sc}(a, \boldsymbol{x})/\mathrm{sc}(a^*, \boldsymbol{x}) = 1$ and we are done. Next, suppose $a \ne a^*$. Let $d = |a - a^*|$ and WLOG assume $a < a^*$. Then,

$$\begin{aligned}
\mathrm{sc}(a, \boldsymbol{x}) &= \sum_{i,j} |x_{ij} - a| \\
&= \left(\sum_{i,j:\, x_{ij} \le a} (a - x_{ij}) + \sum_{i,j:\, a < x_{ij} \le a^*} (a - x_{ij}) + \sum_{i,j:\, x_{ij} > a^*} (x_{ij} - a) \right) \\
&\le \left(\sum_{i,j:\, x_{ij} \le a} (a - x_{ij}) + \sum_{i,j:\, a < x_{ij} \le a^*} d + \sum_{i,j:\, x_{ij} > a^*} (d + (x_{ij} - a^*)) \right) \\
&= \left(\sum_{i,j:\, x_{ij} \le a} (a - x_{ij}) + |\{i,j : x_{ij} > a\}| d + \sum_{i,j:\, x_{ij} > a^*} (x_{ij} - a^*) \right) \\
&\le \left(\sum_{i,j:\, x_{ij} \le a} (a - x_{ij}) + \frac{3}{4} nkd + \sum_{i,j:\, x_{ij} > a^*} (x_{ij} - a^*) \right) \\
&\le 3 \left(\sum_{i,j:\, x_{ij} \le a} (a - x_{ij}) + \sum_{i,j:\, x_{ij} > a^*} (x_{ij} - a^*) + \frac{1}{4} nkd \right) \\
&\le 3 \left(\sum_{i,j:\, x_{ij} \le a} (a - x_{ij} + d) + \sum_{i,j:\, x_{ij} > a^*} (x_{ij} - a^*) \right) \\
&\le 3 \left(\sum_{i,j:\, x_{ij} \le a} (a - x_{ij} + d) + \sum_{i,j:\, a < x_{ij} \le a^*} (a^* - x_{ij}) + \sum_{i,j:\, x_{ij} > a^*} (x_{ij} - a^*) \right) \\
&= 3 \sum_{i,j} |x_{ij} - a^*| \\
&= 3 \cdot \mathrm{sc}(a^*, \boldsymbol{x}).
\end{aligned}$$

Thus, Mechanism 1 is a 3-approximation algorithm for social cost.

(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 $f : \mathbb{R}^k \times \mathbb{R}^k \to \mathbb{R}$ 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

$$f((\underbrace{a, \ldots, a}_{k}), (\underbrace{b, \ldots, b}_{k})) \in \{a, b\}.$$

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

Solution: Fix an arbitrary $\varepsilon > 0$. Let $\ell \in \mathbb{N}$ be such that $3 - \varepsilon < \frac{3\ell+1}{\ell+1}$ ($\ell = \lceil 2/\varepsilon \rceil$ will do). Choose $k = 2\ell + 1$. Fix an arbitrary strategyproof deterministic mechanism $f : \mathbb{R}^k \times \mathbb{R}^k \to \mathbb{R}$. We'll show that f cannot have an approximation ratio better than $\frac{3\ell+1}{\ell+1}$, which implies it is not better than $3 - \varepsilon$, as needed.

To that end, we'll begin by showing that there exist distinct points $x_1 \neq x_2 \in \mathbb{R}$ such that

$$f((x_1, \ldots, x_1), (x_2, \ldots, x_2)) \in \{x_1, x_2\}. \tag{1}$$

To show this, first, fix two arbitrary distinct points $x \neq y$ and suppose

$$f((x, \ldots, x), (y, \ldots, y)) = z.$$

If $z = x$ or $z = y$, then we have found such points and we are done. On the other hand, suppose $z \neq x$ and $z \neq y$. Consider $z' := f((x, \ldots, x), (z, \ldots, z))$. We claim $z' = z$, showing that x and z are our desired points. Indeed, if this was not the case, then f would not be strategyproof, as agent 2 could benefit by deviating from $(z, \ldots, z)$ to $(y, \ldots, y)$, improving from a location not at z (with positive cost) to one exactly at z (with zero cost).

Fix such points x_1, x_2 satisfying Equation (1) and, without loss of generality, suppose $f((x_1, \ldots, x_1), (x_2, \ldots, x_2)) = x_1$. Now we claim that:

$$f((\underbrace{x_1, \ldots, x_1}_{\ell+1}, \underbrace{x_2, \ldots, x_2}_{\ell}), (x_2, \ldots, x_2)) = x_1. \tag{2}$$

Indeed, as argued in part a), each agent's most preferred location is the median of their locations, so under the reported preferences for agent 1 in Equation (2), this remains x_1. If f did not choose x_1, then it would not be strategyproof as agent 1 would benefit by deviating to $(x_1, \ldots, x_1)$ (achieving their most desired location).

Finally, we claim that choosing x_1 in (2) only provides an at least $\frac{3\ell+1}{\ell+1}$ approximation to the optimal cost. Let $d = |x_2 - x_1|$ be the distance between x_1 and x_2. Note that the cost of x_1 is $(\ell+1) \cdot 0 + (\ell+k) \cdot d = (3\ell+1)d$. On the other hand, the cost for x_2 is only $(\ell+1) \cdot d + (\ell+k) \cdot d = (\ell+1)d$. This shows that the approximation to the social cost of f is at most $\frac{3\ell+1}{\ell+1}$, as needed.

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 $x_1^\star < m < x_2^\star$ then (m, m) is the unique Nash equilibrium; this is more generally true when $x_1^\star \leq m \leq x_2^\star$. Our goal is to examine the (almost) complement case of $x_1^\star < x_2^\star < m$, 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 $|x_i^\star - x_j|$, and if there is a tie between the two candidate positions x_1 and x_2 then the cost of candidate i is $\frac{1}{2}(|x_i^\star - x_1| + |x_i^\star - x_2|)$.

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

Note: Please prove both directions.

Solution: Fix values $x_1^\star, x_2^\star$, such that $x_1^\star < x_2^\star < 1/2$ and a pair (x_1, x_2).

We begin with the forward direction. Suppose (x_1, x_2) is a Nash equilibrium. We will first prove

$$x_2 \in [x_2^\star, 1/2]. \tag{3}$$

To that end, we handle each of several cases separately that cover all possible ways this could be false, and show they are not in equilibrium.

- Suppose $x_2 < x_2^\star$ and $x_1 > x_2^\star$ and candidate 1 wins. Then, candidate 1 can increase her payoff by moving to $x_1^\star$.
- Suppose $x_2 < x_2^\star$ and $x_1 > x_2^\star$ and candidate 1 loses or ties. Then, candidate 2 can increase her payoff by moving to $x_2^\star$.
- Suppose $x_2 < x_2^\star$ and $x_1 = x_2^\star$. Then, candidate 1 can increase her payoff by moving to $x_1^\star$.
- Suppose $x_2 < x_2^\star$ and $x_1 < x_2^\star$. Then, candidate 2 can increase her payoff by moving to $x_2^\star$.
- Suppose $x_2 > 1/2$ and $x_1 < x_2^\star$. Then, candidate 2 can increase her payoff by moving to $x_2^\star$.
- Suppose $x_2 > 1/2$ and $x_1 > 1/2$. Then, candidate 2 can increase her payoff by moving to $1/2$.
- Suppose $x_2 > 1/2$ and $x_2^\star \le x_1 \le 1/2$ and candidate 1 wins. Then, candidate 1 can increase her payoff by moving to a point in $(\max(1 - x_2, x_1^\star), x_1)$, which must be a nonempty interval because candidate 1 is currently winning, so $x_1 > 1 - x_2$, and $x_1 \ge x_2^\star > x_1^\star$.
- Suppose $x_2 > 1/2$ and $x_2^\star \le x_1 \le 1/2$ and candidate 2 wins or ties. Then, candidate 2 can increase her payoff by switching to x_1.

This proves Equation (3). Now we split into two more cases, either $x_2 \in (x_2^\star, 1/2]$ or $x_2 = x_2^\star$.

- If $x_2 \in (x_2^\star, 1/2]$, then we must have $x_1 = x_2$. To show this, we show that all other cases are not equilibria.
  - Suppose $x_1 > x_2$ and candidate 1 is winning or tied. Then, candidate 1 can improve her payoff by moving to x_2.
  - Suppose $x_1 > x_2$ and candidate 1 is losing. Then, candidate 2 can improve her payoff by moving to a point in $(1 - x_1, x_2)$, which must be nonempty because candidate 2 is currently winning.
  - Suppose $x_1 < x_2$, then candidate 2 can improve her payoff by moving to $x_2^\star$.

  This means it is true that $x_2^\star < x_2 = x_1 \le 1/2$.
- If $x_2 = x_2^\star$, we claim that $x_1 \notin (x_2^\star, 1 - x_2^\star]$. For any such value of x_1, candidate 1 can improve by moving to $x_2^\star$. This implies that either $x_1 \le x_2^\star$ or $x_1 > 1 - x_2^\star$, as needed.

Next we handle the reverse direction. First, suppose $x_2^\star < x_2 = x_1 \leq 1/2$. This is an equilibrium because even though both candidates would prefer to move left, they cannot do so without losing the election leaving the outcome the same. Next, suppose $x_2 = x_2^\star$ and either $x_1 \leq x_2^\star$ or $x_1 > 1 - x_2^\star$. This is an equilibrium because even though candidate 1 wants to move the outcome to the left, they cannot do so no matter what they report, and candidate 2 is already receiving their optimal payoff.

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 $a_1 \succ_\pi a_2 \succ_\pi \cdots \succ_\pi a_m$, and consider the following generative model, defined by probabilities p_{ij} for all $i = 1, \ldots, m$ and $j = 1, \ldots, i$, 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 $j = 1, \ldots, i$, 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 $\pi \in \mathcal{L}$,

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

   Solution: First, note that the insertion model is equivalent to generating a vector v of m elements, where the i^{th} element is an integer in $[1, i]$ corresponding to the location at which it is inserted in the current subranking. It is clear that this vector uniquely defines a complete ranking over all alternatives. Furthermore, the vector v that corresponds to the correct final ranking is $\langle 1, 2, \ldots, m\rangle$.

   Now, given a vector v that generates a final ranking σ, the Kendall-Tau distance between σ and π is the L1 norm between v and $\langle 1, 2, \ldots, m\rangle$. That is, $d_{KT}(\sigma, \pi) = \sum_{i=1}^m (i - v_i)$. This is because when you put the i^{th} item in spot $j < i$, this necessarily flips $i - j$ comparisons between item i and other items that are supposed to be ranked before it. The argument then proceeds by induction on the number of elements.

   The probability of any ranking $\rho \in \Pi$ can be decomposed as follows:

$$\begin{aligned}
\Pr[\pi] &= p_{1,v(1)} p_{2,v(2)} \cdots p_{m,v(m)} \\
&= \phi^{1-v(1)}\frac{1-\phi}{1-\phi^1} \cdot \phi^{2-v(2)}\frac{1-\phi}{1-\phi^2} \cdot \cdots \cdot \phi^{m-v(m)}\frac{1-\phi}{1-\phi^m} \\
&= \left(\phi^{1-v(1)}\phi^{2-v(2)}\cdots\phi^{m-v(m)}\right)\frac{(1-\phi)^m}{\prod_{k=1}^m(1-\phi^k)} \\
&= \phi^{d_{KT}(\rho,\sigma)}\frac{(1-\phi)^m}{\prod_{k=1}^m(1-\phi^k)} && \text{(because of the above fact)} \\
&= \phi^{d_{KT}(\rho,\sigma)}\left(\frac{1-\phi}{1-\phi}\cdot\frac{1-\phi}{1-\phi^2}\cdot\cdots\cdot\frac{1-\phi}{1-\phi^m}\right)
\end{aligned}$$

$$= \phi^{d_{KT}(\rho,\sigma)} \left(1 \cdot \frac{1}{1+\phi} \cdot \dots \cdot \frac{1}{1+\phi+\dots+\phi^{m-1}}\right)$$

$$= \phi^{d_{KT}(\rho,\sigma)} \frac{1}{\sum_{\tau\in\mathcal{L}} \phi^{d_{KT}(\tau,\sigma)}}. \qquad \text{(by the hint)}$$

This is the same as the Mallows model, so by setting $p_{ij} = \phi^{i-j}\frac{1-\phi}{1-\phi^i}$ in the insertion model, we get the Mallows ϕ-model.

References

[1] L.-G. Svensson. The proof of the Gibbard-Satterthwaite theorem revisited. Working Paper No. 1999:1, Department of Economics, Lund University, 1999.