Optimized Democracy (Spring 2023)
Problem Set #1

Due: 2/13/2023 11:59pm ET

Instructions:

- It is fine to look up a complicated sum or inequality, but please do not look up an entire solution. In particular, 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 may discuss the problems with classmates but please write down solutions completely on your own.
- Please *type up* your solution and submit to Gradescope.

Problems:

1. In class we discussed the following notion of monotonicity: If $f(\boldsymbol{\sigma}) = a$ and $\boldsymbol{\sigma'}$ is a profile such that (i) $[a \succ_{\sigma_i} x \;\Rightarrow\; a \succ_{\sigma'_i} x]$ for all $x \in A$ and $i \in N$, and (ii) $[x \succ_{\sigma_i} y \;\Leftrightarrow\; x \succ_{\sigma'_i} y]$ for all $x, y \in A \setminus \{a\}$ and $i \in N$, then $f(\boldsymbol{\sigma'}) = a$. Informally, if you push a upwards and everything else remains the same, a stays the winner. It can be verified that rules like plurality, Borda count and Llull are monotonic; by contrast, we saw that STV is nonmonotnic.

   [**30 points**] Show that Dodgson's Rule is nonmonotonic by giving a counterexample.

2. 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 $f(\boldsymbol{\sigma'}) = a$.

   (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).

(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.

Important 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.*

3. 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.

(b) **[25 points]** Prove that Mechanism 1 is a 3-approximation algorithm for the social cost in the multiple locations setting.

(c) **[20 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.)