# Optimized Democracy (Spring 2023)
# Problem Set #2

Due: 2/28/2023 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. This problem deals with the Hotelling model with policy-motivated candidates (slides 14–15 of Lecture 5). 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 $\frac{1}{2}(|x_i^\star - x_1| + |x_i^\star - x_2|)$.

   **[15 pt]** 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.

2. In class we discussed the Mallows model, which gives an expression for the probability of a ranking $\sigma$ given the ground truth $\pi$. 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 $\sigma$. 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 $\phi$ 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 $\pi \in \mathcal{L}$,

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

3. A shortcoming of the epistemic approach we discussed in class is that the "optimal" rule depends on the details of the noise model. For example, Kemeny is an MLE with respect to the Mallows model, but wouldn't be an MLE if the noise had a different form. In this problem we will instead explore a worst-case epistemic approach.

   Let $\mathcal{L}$ be the set of rankings over alternatives. Let

   $$d : \mathcal{L} \times \mathcal{L} \to [0, \infty)$$

   be a *metric* over $\mathcal{L}$, which means that, for any rankings $\sigma_1, \sigma_2, \sigma_3 \in \mathcal{L}$,

   - $d(\sigma_1, \sigma_2) = 0 \iff \sigma_1 = \sigma_2$,
   - $d(\sigma_1, \sigma_2) = d(\sigma_2, \sigma_1)$, and
   - $d(\sigma_1, \sigma_3) \le d(\sigma_1, \sigma_2) + d(\sigma_2, \sigma_3)$ (this is called the *triangle inequality*).

   Think of $d$ as an abstract way of measuring the distance between two preference rankings. For example, it is easily verified that $d_{KT}$ satisfies all three axioms.

   Suppose that there is some ground truth ranking $\pi \in \mathcal{L}$, and we are given an input preference profile $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \ldots, \sigma_n) \in \mathcal{L}^n$ with the guarantee that the average distance between $\pi$ and rankings in $\boldsymbol{\sigma}$ is at most some constant $t \ge 0$. That is, we assume

   $$\pi \in B_t(\boldsymbol{\sigma}) := \{\tau \in \mathcal{L} \mid d(\boldsymbol{\sigma}, \tau) \le t\},$$

   where

   $$d(\boldsymbol{\sigma}, \tau) := \frac{1}{n}\sum_{i\in N} d(\sigma_i, \tau).$$

   We then choose $\widehat{\pi} \in \mathcal{L}$ to minimize the worst-case distance from $\widehat{\pi}$ to the unknown ground truth $\pi$. Under this process, the worst-case distance to the ground truth is given by

   $$k := \max_{\boldsymbol{\sigma}\in\mathcal{L}^n} \min_{\widehat{\pi}\in\mathcal{L}} \max_{\pi\in B_t(\boldsymbol{\sigma})} d(\widehat{\pi}, \pi).$$

   We seek to understand the behavior of $k$ as a function of $t$ (as it turns out, $n$ and $d$ don't matter too much).

   Prove the following statements:

(a) **[8 points]** $k \leq 2t$. In words, given $\boldsymbol{\sigma}$ whose average distance from $\pi$ is at most $t$ it is always possible to find a ranking $\widehat{\pi} \in \mathcal{L}$ that is guaranteed to be at distance at most $2t$ from $\pi$.

(b) **[12 points]** Suppose that, instead of allowing for an arbitrary $\widehat{\pi} \in \mathcal{L}$, we require that $\widehat{\pi}$ be one of the rankings $\sigma_i$ of the input profile $\boldsymbol{\sigma}$, i.e., define

$$k' := \max_{\boldsymbol{\sigma} \in \mathcal{L}^n} \min_{i \in N} \max_{\pi \in B_t(\boldsymbol{\sigma})} d(\sigma_i, \pi).$$

Then $k' \leq 3t$.

(c) **[10 points]** Assume that $t$ is in the image of $d$. Then $k \geq \frac{t}{2}$.

(d) **[15 points]** Assume that $t$ is in the image of $d$, and that $d$ is neutral (i.e., the distance between two rankings is invariant to renaming the alternatives). Then $k \geq t$.

**Hint:** Start from a profile $\boldsymbol{\sigma}$ in which all voters have the same ranking $\sigma \in \mathcal{L}$. Also note that neutrality and the assumption that $t$ is in the image of $d$ imply that for any ranking $\pi \in \mathcal{L}$ there is a ranking $\pi' \in \mathcal{L}$ such that $d(\pi, \pi') = t$.

4. Recall the epistemic liquid democracy model, which was introduced and analyzed in Lecture 7 slides 5–12. We saw that local delegation mechanisms cannot satisfy do no harm (DNH) and positive gain (PG).

Now consider the following non-local delegation mechanism. For each voter $i = 1, \ldots, n$, if $A_G(i) \neq \emptyset$ ($i$ approves other voters), the mechanism determines the lowest $j$ such that $i$ approves $j$, that is, $j = \min A_G(i)$. Then $i$ delegates to $j$ if and only if $j$ has not already delegated their vote and there is no other voter who has already delegated to $i$ or $j$. Intuitively, under this mechanism, some delegations could take place but no voter would ever have a weight of more than 2.

For simplicity, let us fix $\alpha = 0.1$ for this problem, that is, $i$ approves $j$ if and only if $(i, j) \in E$ and $p_j > p_i + 0.1$.

(a) **[5 pt]** Show that the above mechanism satisfies PG.

(b) **[15 pt]** Show that the above mechanism does *not* satisfy DNH.

**Note and hint:** This is surprising because the mechanism prevents the problem of voters amassing large weight. It turns out that the mechanism does satisfy DNH with an additional assumption: for all $i \in N$, $p_i \in [\beta, 1 - \beta]$ for $\beta > 0$. Here you are asked to give a family of counterexamples that would necessarily have to violate this assumption; in particular, some voters can have $p_i = 0$ or $p_i = 1$.