MS&E 336/CS 366: Computational Social Choice. Win 2023-24
Course URL: `http://www.stanford.edu/~ashishg/msande336/index.html`.
Instructor: Ashish Goel, Stanford University.

Lecture 6, 1/25/2024. Scribed by Orrie B. Page.

7 Random Dictator Distortion Bounds

7.1 Upper Bound

In this section we prove the upper bound on metric distortion of a random dictator social choice function is 2 when the bliss point assumption is in place.

First we will need to define variables. Let there be N voters (v_i for $i \in \{1, \ldots, N\}$) who submit strict rankings over M candidates (c_i for $i \in \{1, \ldots, M\}$). Call this submitted set of voter rankings S and let the distortion of a chosen candidate c be $D(c, S)$, let the cost to any voter v_i of candidate c_j be $d(v_i, c_j)$. For notational convenience we will say that $d \in S$ if $d(\cdot, \cdot)$ is consistent with voter submission S. We define the distortion $D(c, S)$ as follows:

$$D(c, S) = \max_{d \in S} \frac{\sum_{i=1}^{N} d(v_i, c)}{\min_{j \in \{1,\ldots,M\}} \sum_{i=1}^{N} d(v_i, c_j)}$$

Next we need to define the concept of a bliss point. Let $J(i)$ be the favorite candidate of voter v_i, meaning the candidate they chose first. The bliss point assumption states that $d(v_i, J(i)) = 0 \ \forall \ i \in \{1, \ldots, N\}$, or in words, that every voter is perfectly "in bliss" when their top choice is picked. We will be using this assumption for the rest of this proof.

Now consider the random dictator social choice function. Assume without loss that the true optimal candidate under d is J^*. The expected total cost can be written as follows:

$$\begin{aligned}
\text{Cost}_{RD} &= \frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{N} d(v_j, J(i)), \\
&\leq \frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{N} d(v_j, J^*) + d(v_i, J^*) + d(v_i, J(i)), \\
&\leq \frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{N} d(v_j, J^*) + \frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{N} d(v_i, J^*), \\
&\leq \sum_{j=1}^{N} d(v_j, J^*) + \sum_{i=1}^{N} d(v_i, J^*) = 2 \times \text{ the optimal total cost.}
\end{aligned}$$

This concludes our proof that the upper bound of the metric distortion from the random dictator social choice function is 2 when the bliss point assumption is used.

7.2 Lower Bound

Next we will demonstrate that a random dictator social choice function cannot achieve a distortion lower than 2, even with the bliss point assumption. To see that imagine $n+1$ voters arranged in a star with the $(n+1)th$ voter at the center and the other n voters all shooting off from that center. We then say that each of the candidates is exactly on top of each of our voters in our metric space and define our distance function as follows:

$$d(v_i, c_j) = \begin{cases} 0, & \text{if } i = j \\ 1, & \text{if } i \neq n+1, j = n+1 \\ 2, & \text{if } j \neq n+1, i \neq j \end{cases}$$

In this setting the expected total cost can be written as follows:

$$\begin{aligned} \text{Cost}_{RD} &= \frac{1}{n+1} \sum_{i=1}^{n+1} \sum_{j=1}^{n+1} d(v_j, J(i)), \\ &= \frac{1}{n+1}(n + n(1 + 2(n-1))), \\ &= \frac{1}{n+1}(n + n + 2n^2 - 2n)), \\ &= \frac{2n^2}{n+1} \end{aligned}$$

Note that in the optimal case we choose the center candidate (c_{n+1}) corresponding to a total cost of m. Since the distortion is the realized cost divided by the optimal cost the distortion is $2n/(n+1)$ which in the limit as $n \to \infty$ approaches 2.

This concludes our proof that the random dictator social choice function cannot achieve a distortion lower than 2, even with the bliss point assumption.

8 Existence is all we Need

In this section we show that if all we care about is minimizing the metric distortion, we don't need to know the social choice rule that does so. Here we will define the distortion of a candidate c_1 with respect to another candidate c_2 as if c_2 had been the optimal candidate. We will call this $D(c_1, c_2)$ and define it as follows:

$$D(c_1, c_2) = \frac{\sum_{i=1}^{N} d(v_i, c_1)}{\sum_{i=1}^{N} d(v_i, c_2)}$$

Where all variables other than $D(\cdot,\cdot)$ are defined as in section 7.1. Our aim is to choose some c^* such that the following is true:

$$c^* = \arg\min_{c_1} \max_{c_2} D(c_1, c_2)$$

Next we will set this up as a linear program with a constraint that $\forall\ (i, i', j, j')\ d_{ij} \leq d_{ij'} + d_{i'j'} + d_{i'j}$, where d_ij represents $d(v_i, c_j)$. We will also add constraints to ensure that $d_{ij} \geq 0\ \forall\ (i, j)$. For our optimization function we can simply divide out the numerator of our $D(c_1, c_2)$ function to linearize it. Solve this linear program for every pair of candidates will then allow us to easily identify a c^* such that $c^* = \arg\min_{c_1} \max_{c_2} D(c_1, c_2)$.

Hence, we did not need any social choice function to identify the candidate that minimizes metric distortion, we can simply use the above linear program.

9 Sequential Deliberation

9.1 The General Problem

In this section we cover an infinite horizon game that is set up as follows. There are N agents, each with some preference over all possible solutions s in some solution space S. There is also some predetermined initial solution s^0. The game then proceeds as follows in every round t:

1. Two agents are drawn uniformly–with replacement–from the pool of all agents. Denote these two agents u^t and v^t.

2. The two agents attempt to agree on a new solution s. If they can agree $s^t = s$, if not the outside option is the former solution, so $s^t = s^{t-1}$.

An important property of these games is that as the game continues and $t \to \infty$ it is guaranteed to converge to a distribution δ over the solution space S.

9.2 An Explicit Example

In this section we look at a specific variant of the general game format described in section 9.1. We will again assume there are n agents but we will now assume that each agent i is placed on a line at a position such that $x_i = i$, so agent 1 is at 1, agent 2 is at 2, and so on. We will also assume that agents preferences are defined strictly by the proximity of the solution to their position, where the solution must be a point in $\mathcal{R}^1$.

The first thing to notice in this game is that thanks to Nash Bargaining (which we discussed briefly in class) for any u^t, v^t and s^{t-1}, the chosen solution will be $s^t = \text{median}\{u^t, v^t, s^{t-1}\}$. Note that because of this, in any long run we will have that all solutions are one of the digits 1,...,n.

We can then solve for the convergent distribution as follows. Define p_i as $p_i = \lim_{t\to\infty} P(s^t = i)$ and define C_i as $C_i = \sum_{j=1}^{i} p_j$. Note that there are two ways in which s^t can be less than or equal to i:

- EITHER: $s^{t-1} \le i$ and $\min\{u^t, v^t\} \le i$ (case 1)
- OR: $s^{t-1} > i$ and $\max\{u^t, v^t\} \le i$ (case 2)

Since C_i represents the probability that $s^t \le i$, we can solve for it by summing over the probability of the above two cases occurring. We know already that in the long run the probability that $s^t \le i$ is going to be whatever C_i is so we can use that fact to solve as follows:

$$C_i = C_i\left(1 - \left(\frac{n-i}{n}\right)^2\right) + (1 - C_i)\left(\frac{i}{n}\right)^2,$$

$$C_i = C_i - \frac{C_i(n-i)^2}{n^2} + \frac{i^2}{n^2} - \frac{i^2 C_i}{n^2},$$

$$C_i = C_i\left[1 - \frac{(n-i)^2}{n^2} - \frac{i^2}{n^2}\right] + \frac{i^2}{n^2},$$

$$C_i n^2 = C_i[n^2 - (n-i)^2 - i^2] + i^2,$$

$$C_i[(n-i)^2 + i^2] = i^2.$$

Which simplifies to our final identity of:

$$C_i = \frac{i^2}{i^2 + (n-i)^2}.$$