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 5, 1/23/2024. Scribed by Tomer Zaidman.

A Quick Recap of Participatory Budgeting: We write $x^{(i)} = (x_1^{(i)}, x_2^{(i)}, \ldots, x_N^{(i)})$ for the ideal budget of the i-th user, $z = (z_1, z_2, \ldots, z_n)$ for the aggregated budget, and B for the cap on the sum of both the user budget and any chosen budget. There are several reasonable ways of modelling an individual's utility from an aggregated budget z, but we looked in particular at the overlap utility, defined by $u_i(z) = \sum_{j=1}^{N} \min\{z_j, x_j^{(i)}\}$, which is equivalent to using the L1 norm between $x^{(i)}$ and z, $d_i(z) = ||x^{(i)}, z||_1$. The problem of maximizing the sum of individual utilities can be set up as a linear program, or solved by way of the water-filling algorithm we saw on the participatory budgeting site.

From Ordinal to Cardinal Social Choice

We return to the original problem posed in the first lecture – voters provide a ranking of alternatives, and a social choice function selects a societal winner. Up until now, the rankings provided were purely ordinal – it only matters in which order voters rank preferences. But now, we would like to measure the quality of our aggregation, in which case it would be nice to be able to refer to some (hidden) notion of utility or cost that induce each voter's ordinal preferences, *even if those utilities and/or costs are not known to the social planner.*

We already saw an example of this – when costs are distances on a line, we select the median vote because it is the Condorcet winner, hence the Copeland winner, and also the choice that minimizes total cost to voters. The reasonable assumption was that voters prefer candidates in decreasing order of distance. Even when we do not know the actual distances between candidates, the rank list suffices to find a Condorcet winner, who will sit at the median.

How, then, will we evaluate a social choice function? We will do it on the basis of **distortion**, a measure which quantifies the worst-case ratio between the total cost/utility to society of the alternative selected by the SCF, and the total cost/utility to society of the optimal choice.

5.1 Distortion and its Bounds

Suppose there are M candidates labelled $c_1, c_2, \ldots, c_M$, and N voters, labelled $v_1, v_2, \ldots, v_N$. Voter i provides a strict ranking of candidates $\succ_i$. Assume further that there is a hidden[1] cost $d(v_i, c_j) = d_i(c_j)$ that denotes the unhappiness (or cost) v_i experiences with the outcome c_j. Alternatively, there is a hidden utility $u(v_i, c_j) = u_i(c_j)$ that v_i derives from outcome c_j.

[1]This means the mechanism designer only knows M, N, and $\succ_i$ for each i.

Assumption 5.1 *Preferences are consistent with hidden costs/utilities. That is, voter i ranks c_j over $c_{j'}$ only if $d_i(c_j) \leq d_i(c_{j'})$, or equivalently $u_i(c_j) \geq u_i(c_{j'})$.*

The total cost associated with candidate c_j with hidden cost functions d_i is $TC_d(c_j) = \sum_{i=1}^N d(v_i, c_j)$, and the total utility with hidden utility functions u_i is $TU_u(c_j) = \sum_{i=1}^N u(v_i, c_j)$.

Definition 5.1 *The **distortion** of c_j given hidden cost functions d_i satisfying assumption 5.1 is defined as*

$$D(c_j) = \max_{c \in \{c_1, \ldots, c_M\}} \frac{TC_d(c_j)}{TC_d(c)} = \max_{c \in \{c_1, \ldots, c_M\}} \frac{\sum_{i=1}^N d_i(c_j)}{\sum_{i=1}^N d_i(c)} = \frac{\sum_{i=1}^N d_i(c_j)}{\min_c \sum_{i=1}^N d_i(c)}$$

Equivalently, we can define distortion for hidden utility functions u_i by

$$D(c_j) = \max_{c \in \{c_1, \ldots, c_M\}} \frac{TU_u(c)}{TU_u(c_j)} = \max_{c \in \{c_1, \ldots, c_M\}} \frac{\sum_{i=1}^N u_i(c)}{\sum_{i=1}^N u_i(c_j)} = \frac{\max_c \sum_{i=1}^N u_i(c)}{\sum_{i=1}^N u_i(c_j)}$$

The distortion of a social choice function is thus the expected distortion as computed over candidates it selects with positive probability.

The question we will ask is, what are the maximum and minimum (expected) distortion that a social choice function might achieve, where the optimization is taken over all possible hidden cost/utility functions (and thus over all preference orderings)?

Proposition 5.2 (Impossibility Theorem) *Suppose voters have either hidden cost functions c_i or hidden utility functions u_i which induce their preferences on outcomes. Then:*

- *In a world of utility, no deterministic SCF has a finite lower bound on distortion, and no randomized SCF has a lower bound below M.*
- *In a world of cost, no deterministic or randomized algorithm has a finite lower bound on distortion.*

For example, consider the following voter profile, and family of utility profiles representing ordinal preferences for $\epsilon > 0$:

v_1	v_2
c_1	c_2
c_2	c_1

	v_1	v_2
c_1	ϵ	0
c_2	0	$1/\epsilon$

A deterministic SCF, without loss of generality, selects c_1 with certainty. Then the distortion of this SCF is the distortion of c_1 which is $\frac{\epsilon}{1/\epsilon} = \epsilon^2 \to \infty$ as ϵ goes to infinity. As another example, consider instead the following voter profile:

v_1	v_2	$\cdots$	v_M
c_1	c_2	$\cdots$	c_M
?	?	$\cdots$	?
$\vdots$	$\vdots$	$\cdots$	$\vdots$
?	?	$\cdots$	?

Without loss of generality, suppose a randomized algorithm chooses c_1 with probability at most $1/M$. Then we can consider the following utility function: $u(v_1, c_1) = 1$, and $u(v_i, c_j) = 0$ for all other (i, j). Hence the expected total utility of this randomized algorithm is at most $1/M$, while the optimal total utility is 1, achieved when c_1 is selected, leading to a distortion of at least M. Finally, consider the following voter and cost profiles for $\epsilon > 0$:

v_1	v_2
c_1	c_2
c_2	c_1

	v_1	v_2
c_1	0	1
c_2	0	0

The deterministic algorithm which picks c_1 with certainty has infinite distortion, and we can find another representative utility profile so that the algorithm picking c_2 with certainty also has infinite distortion.

5.2 So What Can We Do?

So we have two sides of the spectrum. If we restrict preferences to lie on a line, then we can select the median voter and do reasonably well. In general, however, the level of distortion cannot be limited. Can we achieve an intermediate result with a reasonable restriction on utilities and costs? It turns out we can, to an extent.

Assumption 5.3 *The sum of utilities must be 1 for any voter. Alternatively, assume voters and candidates lie in a metric space, and the cost of a voter for a candidate is exactly the distance between them in this metric space.*[2]

Proposition 5.4 *Suppose assumption 5.3 holds. Then:*

- *In a world of utilities, a deterministic SCF still has a lower bound of M on distortion, while a randomized algorithm has both lower and upper bounds on the order of $\sqrt{M}$.*
- *In a world of costs, deterministic SCFs have a distortion of at most 3, and randomized SCFs have both an upper bound of 3 and a lower bound of 2.*

The result pertaining to deterministic SCFs in a world of metric costs is especially striking, but it is among several results with fairly accessible proofs. For example:

Theorem 5.5 *Suppose there are two voters and two candidates, and that costs for both voters are represented by distance on a line. Then the distortion of a deterministic algorithm is at most 3.*

Proof: Consider the usual voting profile below, and a deterministic algorithm that, without loss of generality, chooses c_2.

[2]For those who have not come across metric spaces before, the important characteristic is a "distance function" on a set which satisfies the triangle inequality:

$$d(v_i, c_j) \leq d(v_i, c_{j'}) + d(v_{i'}, c_{j'}) + d(v_{i'}, c_j)$$

for all $v_i, v_{i'}, c_j, c_{j'}$.

v_1	v_2
c_1	c_2
c_2	c_1

These preferences are induced by a some unknown metric. Without loss of generality, suppose that metric has $d(v_1, c_1) = 0$, and $d(v_1, c_2) = 2$, so the space under consideration is a line with endpoints 0 and 2 at which are positioned v_1, c_1, and c_2 respectively:

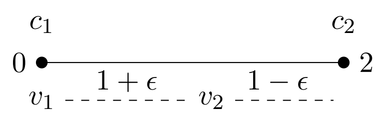

The voter v_2 lies in the middle at distance $1+\epsilon$ from c_1 and $1-\epsilon$ from c_2. Then the total cost of our deterministic social choice function is $2+(1-\epsilon) = 3-\epsilon$ while the optimal total cost is obtained from selecting c_1 and is $0+(1+\epsilon) = 1+\epsilon$. The distortion is therefore $\frac{3-\epsilon}{1+\epsilon}$ which approaches 3 in the limit as ϵ goes to zero. We could repeat this analysis for a deterministic function choosing c_1, producing representative utility functions which give a distortion of at most 3. ■

We can prove another remarkable result pertaining to one of our tried and true Condorcet consistent voting rules:

Theorem 5.6 *The Copeland rule has a distortion of at most 9.*[3]

We prove this by way of two lemmas:

Lemma 5.7 *If C is the Copeland winner, then for any other candidate C', either C beats C' in a pairwise election, or there exists C'' such that C beats C'' and C'' beats C', each in a pairwise election.*

Proof: Suppose for the sake of contradiction that neither is true. That is, C' beats C in a pairwise election, and no C'' as described exists. Hence for every C'' which is beaten by C, it is also beaten by C'. Then C' beats more alternatives in pairwise elections than does C, so C cannot be a Copeland winner. ■

Lemma 5.8 *If C beats C' in a pairwise election, then the total cost of C under any metric is at most 3 times the total cost of C' under that same metric.*

Proof: Let C beat C' in a pairwise election. Define the set $S_1 = \{v_i : C \succ_i C'\}$, and $S_2 = V \setminus S_1 = \{v_i : C \prec_i C'\}$. For every voter $v \in S_2$, assign a unique match $m(v) \in S_1$. This is possible since C beating C' in a pairwise election implies $|S_1| > |S_2|$. Fix any metric d. The total cost of alternative C is thus

$$TC(C) = \sum_{v \in S_1} d(v, C) + \sum_{v \in S_2} d(v, C)$$

[3] The literature actually proves a bound of 5, see [1], but this bound is much more accessible.

Voters in S_1 prefer C to C', so their distance to C cannot be any less than their distance to C'. Thus

$$TC(C) \leq \sum_{v \in S_1} d(v, C') + \sum_{v \in S_2} d(v, C)$$

Furthermore, the triangle inequality applied to $v \in S_2$, $m(v)$, C, and C' gives

$$TC(C) \leq \sum_{v \in S_1} d(v, C') + \sum_{v \in S_2} \big(d(v, C') + d(m(v), C') + d(m(v), C)\big)$$

Again, $m(v) \in S_1$ implies their distance from C is no move than their distance from C', so

$$TC(C) \leq \sum_{v \in S_1} d(v, C') + \sum_{v \in S_2} \big(d(v, C') + d(m(v), C') + d(m(v), C')\big)$$

Restating the sum over S_2 of $m(v)$ as a sum over S_1 of v (noting that the second set is bigger so we again increase the size of the sum), we have

$$TC(C) \leq 3 \sum_{v \in S_1} d(v, C') + \sum_{v \in S_2} d(v, C') \leq 3 \sum_{v \in S_1 \cup S_2} d(v, C') = 3TC(C')$$

■

Proof of Theorem 5.6: Suppose C is the Copeland winner, and C' is the optimal candidate as measured by total cost. By Lemma 5.8, either C beats C' in a pairwise election, or there exists C'' such that C beats C'' and C'' beats C'. In the first case, Lemma 5.9 gives us that $TC(C) \leq 3TC(C')$. In the second case, applying Lemma 5.9 twice gives $TC(C) \leq 3TC(C'') \leq 9TC(C')$. Hence, in the worst case scenario, the Copeland rule picks a candidate with total cost 9 times that of the optimal candidate, and thus the distortion of the Copeland rule is at most 9. ■

Remaining results will be covered next lecture.

References

[1] Elliot Anschelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. "Apprx-oimating optimal social choice under metric preferences." *Artifical Intelligence* 264: 27-51, 2018