MS&E 336/CS 366: Computational Social Choice. Aut 2010-21
Course URL: `http://www.stanford.edu/~ashishg/msande336/index.html`.
Instructor: Ashish Goel, Stanford University.

Lecture 10, 10/20/2021. Scribed by Juan Langlois.

10.1 Lecture plan for the rest of the quarter

- Week 6: Market-based allocation
- Week 7: Opinion Dynamics
- Week 8: Fairness
- Week 9: Loose Threads (additional social choice rules, K-S bargaining, Maskin and other monoticity)
- Week 10: Deliberation

10.2 Sequential Deliberation

The goal of this and the following lecture is to analyze the expected distortion of the stationary distribution of the result of the sequential deliberation when we assume that the agents do Nash Bargaining. This analysis is based on [1].

In a typical social choice setting, the mechanism designer could elicit solutions from the users and have them vote. However, voting on these options does not necessarily find the social optimum because it is not clear that the social optimum is even on the ballot. In such a setting, deliberation between individuals could find entirely new alternatives. This mechanism leads to finding a social optimum over a wider space of semi-structured outcomes that the system/mechanism designer was not originally aware of and the participants had not initially articulated.

With this idea in mind, the authors of [1] start with the following premise: The mechanism designer may not be able to enumerate the outcomes in the decision space or know their structure, and this decision space may be too big for most ordinal voting schemes. However, we can assume that agents can still reason about their preferences and, small groups of agents can negotiate over this space and collaboratively propose outcomes that appeal to all of them. The authors' goal is to design protocols based on such a primitive by which small group negotiation can lead to an aggregation of societal preferences without the need to formally articulate the entire decision space and without every agent having to report ordinal rankings over this space.

10.2.1 Bargaining Theory Review [1]

Two-person bargaining, as framed in [2], is a game wherein there is a disagreement outcome and two agents must cooperate to reach a decision; failure to cooperate results in the adoption of the disagreement outcome. Nash postulated four axioms that the bargaining solution ought to satisfy assuming a convex space of alternatives: Pareto optimally (agents and an outcome that cannot be simultaneously improved for both of them), symmetry between agents, invariance with respect to affine transformations of utility (scalar multiplication or additive translation of any agent's utility should not change the outcome), and independence of irrelevant alternatives (informally that the presence of a feasible outcome that agents do not select does not influence their decision). Nash proved that the solution maximizing the Nash product is the unique solution satisfying these axioms.

10.2.2 Sequential Pairwise Deliberation

The reference [1] uses 2-person bargaining as a basic primitive, and view deliberation as a sequence of pairwise interactions that refine good alternatives into better ones as time goes by. More formally, there is a decision space $\mathcal{S}$ of feasible alternatives (these may be projects, sets of projects, or continuous allocations) and a set $\mathcal{N}$ of agents. We assume each agent has a hidden cardinal utility for each alternative. We encapsulate deliberation as a sequential process. In each round $t = 1, 2, \ldots, T$:

(a) A pair of agents u^t and v^t are chose independently and uniformly at random with replacement.

(b) These agents are presented with a disagreement alternative a^t, and perform bargaining.

(c) Agents u^t and v^t are asked to output a consensus alternative; if they fail to reach a consensus then the alternative a^t is output.

(d) Let o^t denote the alternative that is output in round t. We set $a^{t+1} = o^t$, where we assume a^1 is the bliss point of an arbitrary agent.

The final social choice is a^T. Note that as $T \to \infty$, this is equivalent to drawing a sample from the stationary distribution (assuming one exists). Also, we can draw many samples if we want by repeating the process with a different set of participants.

10.2.3 Analytical Model

We are going to assume that the set $\mathcal{S}$ of alternatives lie on a line. That is, we assume each agent u has a bliss point $p_y \in \mathcal{S}$, and his cost for an alternative $a \in \mathcal{S}$ is simply $d(p_u, a)$, where $d(\dot{)}$ is the distance function on a line:

$$d(p_u, a) = |p_u - a| \tag{1}$$

The assumption in the paper [1] is much weaker. In their case, the set $\mathcal{S}$ of alternatives are vertices of a median graph. This definition is more general than a line but less general than metric spaces. A median graph has the property that for each triplet of vertices u, v, w, there is a unique point that is common to the three sets of shortest paths (since there may be multiple pairwise shortest paths), those between u, v, between v, w, and between u, w. This point is the unique median of u,v, w. For a median graph, the cost function is the shortest path distance function on the median graph. Several natural graphs are median graphs, including trees, points on the line, hypercubes, and grid graphs in arbitrary dimensions.

Digression for Project 2

Project 2 considers a setting where voters are reluctant to be on the losing side, so they choose a region in which they will tolerate any result (such that there is more space in which they can "be seen choosing the winner"). This strategy is still truthful because the region contains the true point. In other words, voter i provides an approval set A_i such that their ideal point $x_i^* \in A_i$. Under these circumstances, the aggregation mechanism chooses the point z that belongs to the most amount of A_i. The penalty for voter i is then

$$c_i(z) = d(x_i^*, z) + \lambda \min_{x \in A_i} d(x, z)$$

.

With $\lambda > 0$. The first term is the cost of the social choice outcome and the second term is the perception cost. The latter assumes that the approval sets A_i are public. If the ideal points x_i^* lie on a line, and we know their median z, then the optimal strategy is to define a set x_i that covers the range between x_i^* and z. Now, if instead, we claim that we know which side the median of the ideal points are,

- can we extend to multiple rounds?
- what is the strategy if the median's direction is unknown?
- can we extend results from a line to median spaces?

Social Cost and Distortion. The objective of the mechanism/system designer is to find a solution that minimizes total cost for all voters. The distortion of the solution s is given by:

$$D(a) = \frac{\sum_{v \in \mathcal{N}} d(p_v, a)}{\min\limits_{a' \in \mathcal{S}} \sum_{v \in V} d(p_v, a')} \tag{2}$$

where we use the expected social cost if a is the outcome of a randomized algorithm. The denominator of the distortion is the minimizer of social cost, i.e., the generalized median. The expected distortion under the stationary distribution $s \sim \pi$ is

$$\mathbb{E}_{a \sim \pi}\left[D(a)\right] = \sum_{a \in \mathcal{S}} \mathbb{P}_\pi(a) D(a) \tag{3}$$

Randomized dictator gives an expected distortion of 2 in the worst case.

Nash Bargaining. The model for two-person bargaining is simply the classical Nash bargaining solution described before. Given a disagreement alternative a, agents u and v choose that alternative $o \in \mathcal{S}$ that maximizes:

$$\text{Nash product} = (d(p_u, a) - d(p_u, o)) \times (d(p_v, a) - d(p_v, o)) \tag{4}$$

subject to individual rationality, that is, $d(p_v, o) \leq d(p_v, a)$ and $d(p_u, o) \leq d(p_u, a)$. The Nash product maximizer need not be unique; in the case of ties the reference [1] postulates that agents select the outcome that is closest to the disagreement outcome. The Nash product is a widely studied axiomatic notion of pairwise interactions, and is therefore a natural solution concept in our framework.

10.2.4 Results

In this section we are going to work toward the proof that under sequential deliberation, the expected Distortion of outcome a^T has an upper bound approaching 1.208. In particular, we are going to analyze only the special case of the line.

When the disagreement alternative a^t lies between p_u and p_v, there is no consensus that can reduce their cost of one agent without reducing the cost of the other. So the Nash Bargining output is going to be the median $o^t = a^t$, i.e. disagreement. On the other hand, when the disagreement alternative a^t lies outside the line p_u and p_v, the agents are going to agree on an alternative o^t that lies between p_u and p_v. In particular, the agents are going to agree on the bliss point of the agent whose preferred decision is closest to the disagreement alternative. For example, and without loss of generality, let us assume that $a^t \leq p_u \leq p_v$ and $d(p_u, a^t) = d$ and $d(p_v, a^t) = d + p_v - p_u$ with $d \geq 0$. The output decision is going to be a point o^t such that $d(p_u, o^t) = x$ and $d(p_v, o^t) = p_v - p_u - x$ with $x \geq 0$. The Nash product under this circumstances is going to be

$$\begin{aligned}
(d(p_u, a) - d(p_u, o)) \times (d(p_v, a) - d(p_v, o)) &= ((d) - (x)) \times ((d + p_v - p_u) - (p_v - p_u - x)) \\
&= (d - x) \times (d + x) \\
&= d^2 - x^2
\end{aligned}$$

This product is clearly maximized when $x = 0$, that is $o^t = p_u$.

References

[1] B. Fain, A. Goel and K. Munagala, S. Sakshuwong. *Sequential Deliberation for Social Choice.* http://arxiv.org/abs/1710.00771, 2017.

[2] J.F. Nash. *The bargaining problem.* Econometrica 19(2), 1950.