MS&E 336/CS 366: Computational Social Choice. Winter 2019-20
Course URL: `http://www.stanford.edu/~ashishg/msande336.html`.
Instructor: Ashish Goel, Stanford University.

Lecture 10, 2/12/2020. Scribed by Jaymee Sheng.

Sequential Deliberation

In complex decision spaces, standard voting methods and ordinal preferences may be insufficient or impractical (for example, the optimal outcome may lie outside the first choices of all the voters). So we turn to negotiation and deliberation to design a mechanism to reach consensus. In particular, we would like the mechanism (or algorithm) to satisfy the following properties [1]:

1. The designer of the algorithm does not need to understand the decision space.
2. The outcomes under simple analytical models should beat random dictatorship.
3. The mechanism should restrict cognitive load on users and encourage negotiation and deliberation.

Let us assume all users and all possible decisions lie in a common metric space d, with $d(v,x)$ as the cost to user v for decision x. Assume also that each agent has a "bliss point", the opinion that lies at the same point in d as the agent, resulting in 0 cost for that agent. We want to find the optimum decision that minimizes total costs for all users, i.e., find x that minimizes $S(x) = \sum_v d(v,x)$. If x^* is the optimum decision, then the **distortion** of a randomized algorithm that produces decision x is $E[S(x)]/S(x^*)$. We proceed with the following protocol for sequential deliberation.

Protocol *Let N be the set of all agents. Suppose we start with an initial suggestion proposed by a random agent, call it a^1. For rounds from $t = 1$ to $t = T$,*

- *Choose two agents u^t, v^t independently and uniformly at random with replacement.*
- *Let a^t be the consensus from the previous around.*
- *Agents u^t and v^t bargain with a^t as the outside alternative. If they agree, we set a^{t+1} to their consensus; otherwise, $a^{t+1} = a^t$.*

Example Suppose we ask agents A and B the question: how many times in a year do you want to eat pizza? Their reported ideal numbers and the outside alternative are as follows:

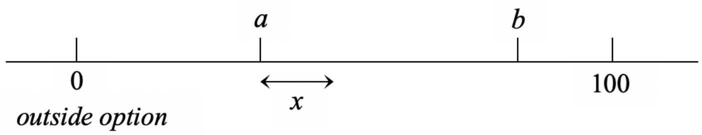

Figure 1: Hypothetical bargaining between two agents

In this case, a is the Nash bargaining outcome. To see this, suppose we increase the number of times from a by x. The utility gain for A is $a - x$ whereas the utility gain for B is $a + x$ relative to the outside alternative 0. Then the solution to $\text{Max}\,[(a - x)(a + x)]$ is $x = 0$, which implies a, the median of the three possible outcomes in this round, is the optimum. In fact, this result can be generalized to a class of decision problems in median spaces.

Definition *A **median space** is a metric space in which for any three points, there exists a unique point that lies on three pairwise shortest paths.*

We know the following results on median spaces:

1. Nash bargaining between agents u and v with ideal points p_u and p_v using disagreement outcome a finds the median of p_u, p_v, a.

2. All agents bargaining by truthfully representing their ideal point is a sub-game perfect Nash equilibrium of the extensive form game defined by sequential bargaining.

3. The chosen alternative converges to a stationary distribution in O(1) steps.

Furthermore, we can show that the **distortion** of sequential deliberation is no more than 1.208 [2]. The remaining lecture focuses on the proving this result for the special case of the line.

Distortion

To illustrate the idea of distortion, let us first consider a simple example. Suppose there are N agents with 1 agent preferring point a and $N - 1$ agents preferring point b. Suppose a and b are 1 unit distance apart. The optimum point, then, is b, which incurs total cost of 1. Formally, let x^* denote the optimum decision. So $S(x^*) = 1$. The random dictator mechanism chooses a with probability $\frac{1}{N}$ and b with probability $\frac{N-1}{N}$, resulting in the following expected distortion:

$$E[S(x_{RD})] = \frac{1}{N} \cdot (N - 1) + \frac{N-1}{N} \cdot 1 = 2 \cdot \frac{N-1}{N}$$

In fact, this is the worst case for random dictatorship, which can achieve a distortion of at most 2. We will now analyze the upper bound of the distortion of sequential deliberation when the metric space is a line.

Setup: Assume N agents' ideal points are $x_1, x_2, ..., x_N$ located on a line. Further assume that every agent does honest Nash bargaining at every step. We already know that the Nash bargaining outcome is the median of the points of the agents chosen and the outside alternative at round t. Hence, only "bliss points" are chosen, and the probability of choosing a point on the line only depends on its order relative to other points.

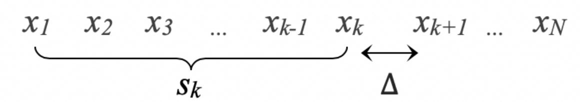

Figure 2: Sequential deliberation on a line

We proceed as follows:

Let p_i be the probability of being at point x_i on the line in stationary distribution (assume one exists), and let $S_k = \sum_{i \le k} p_i$ (see Figure 2).

For the outcome to be in $\{x_1, x_2, ..., x_{k-1}, x_k\}$ at time $t+1$, it must be the case that at time t:

- **either** both chosen agents were in $\{x_1, ..., x_k\}$
- **or** the outside alternative and exactly one of the agents were in $\{x_1, ..., x_k\}$,

which implies

$$S_k = \frac{\binom{k}{2}}{\binom{N}{2}} + \frac{S_k \cdot k(N-k)}{\binom{N}{2}} = \frac{\binom{k}{2}}{\binom{N}{2} - k(N-k)} = \frac{\binom{k}{2}}{\binom{k}{2} + \binom{N-k}{2}}$$

Assume $\frac{k}{N}$ is fixed at z as N goes to infinity. Observe that

$$\lim_{N \to \infty} \frac{\binom{k}{2}}{\binom{N}{2}} = \lim_{N \to \infty} \frac{Nz(Nz-1)}{N(N-1)} = z^2$$

So we can express S_k in terms of S_z as

$$S_z = \frac{z^2}{z^2 + (1-z)^2} \tag{1}$$

For now, let us pretend that we only care about how many agents have to "cross" over this barrier of Δ (Figure 2), which means we can express distortion as:

$$\begin{aligned} D(z) &= \frac{(1-z)S_z + zS_{1-z}}{z} \\ &= \frac{(1-z)S_z + z(1-S_z)}{z} \\ &= \frac{(1-z)z^2 + z(1-z)^2}{z(z^2 + (1-z)^2)} \qquad (\textit{substituting}\,(1)) \\ &= \frac{1-z}{z^2 + (1-z)^2} \end{aligned} \tag{2}$$

Let $\gamma(z) = \frac{1}{D(z)} = \frac{z^2}{1-z} + (1-z)$. Differentiating $\gamma(z)$ with respect to z gives

$$\frac{d\gamma(z)}{dz} = \frac{2z}{1-z} + \frac{z^2}{(1-z)^2} - 1$$

Setting the above to 0 yields $z^* = 1 - \frac{1}{\sqrt{2}} \approx 0.293$ which minimizes $\gamma(z)$. The maximum distortion, then, is $D(z^*) = 1.208$, which concludes the proof.

Intuitively, when $z = 0$, the algorithm always chooses something on the right, which is the same as the optimum, so $D = 1$. Similarly, when $z = 1$, the algorithm always chooses something on the left, which is again the same as the optimum, so $D = 1$. When $z = \frac{1}{2}$, the algorithm chooses left or right with equal probability and the distortion is again 1.

More formally, we can analyze distortion by directly comparing the costs. We have

$$E[\text{Cost of algorithm}] = \sum_{k=1}^{N-1} (x_{k+1} - x_k)[S_k(N-k) + S_{N-k}k]$$

and

$$\begin{aligned} \text{Cost of optimum} &= \sum_{k=1}^{N-1} (x_{k+1} - x_k) \cdot \begin{cases} N-k & \text{if optimum chooses } x_1, \text{...}, x_k \\ k & \text{if optimum chooses } x_{k+1}, \text{...}, x_N \end{cases} \\ &\geq \sum_{k=1}^{N-1} (x_{k+1} - x_k)\, min\{N-k, k\} \end{aligned}$$

Let $A(k) = S_k(N-k) + S_{N-k}k$ and $B(k) = min\{N-k, k\}$. W.L.O.G., assume $k \leq \frac{N}{2}$. We already proved $\frac{A(k)}{B(k)} \leq 1.208$ earlier, which shows formally that the distortion is at most 1.208.

References

[1] B. Fain, A. Goel, and K. Munagala. Sequential Deliberation Slides. https://web.stanford.edu/ ashishg/msande336_handouts/sequential_class_slides.pdf

[2] B. Fain, A. Goel, K. Munagala, and S. Sakshuwong. *Sequential Deliberation for Social Choice.* arXiv:1710.00771 [cs.GT], 2017.