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 11, 2/19/2020. Scribed by Grant Vaska.

11.1 Introduction

Nash, Arrow, Schelling, and Maskin are several Nobel Prize winners relevant to this class. We've already discussed some of the research of Nash, Arrow, and Schelling, today we'll look at some of Maskin's work. One of the reasons work in social choice is exciting is it has relevance to some of the most consequential problems we face together, while also being an intellectually vibrant area of study.

11.2 Mixing

Consider, again, sequential deliberation on a line. At a time $t+1$, there is an outside option S_t as well as two agents selected uniformly at random from the set of agents N, u_{t+1} and v_{t+1}. As discussed earlier, if the agents perform Nash bargaining, they will end up choosing the median of their positions and the outside option (which was the result of the previous round). For any finite N, S_t will go to a stationary distribution. In a rough sense, "mixing fast" means that it will go to a stationary distribution quickly.

$$(S_t, u_{t+1}, v_{t+1}) \xrightarrow{\text{Nash Bargaining}} (S_{t+1}, u_{t+2}, v_{t+2})$$

Assume that in two steps of the process, all four of the randomly selected agents are unique. Denote these agents a, b, c, d (though we are not assuming which corresponds to which) and assume without loss of generality that $a <= b <= c <= d$. The probability that a and b arrive together (meaning $a = u_{t+1}$ and $b = v_{t+1}$ or the reverse, or similarly for $t+2$) is $1/3$ since there are $\binom{4}{2}/2 = 3$ ways to divide a, b, c, d into two groups of two. (Alternately, the probability that a is paired with b is $1/3$.) So with probability $1/3$, the Markov chain collapses in two steps, which implies that the mixing time is constant.

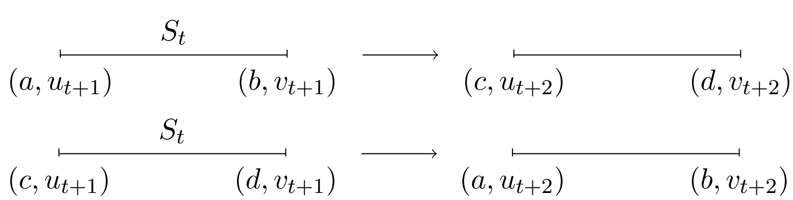

More concretely, let a and b arrive first at time $t+1$, so c and d arrive second at time $t+2$. Then $S_{t+1} = a$ if $S_t <= a$, $S_{t+1} = S_t$ if $a <= S_t <= b$, or $S_{t+1} = b$ if $b <= S_t$. Regardless, since $a <= b <= c <= d$ $S_{t+2} = c$. Similarly, if c and d arrive first then $S_{t+2} = b$. In a sense, with probability 1/3 the process "forgets" the past in two steps to arrive at the stationary distribution, regardless of the actual values of a, b, c, d. We will omit the formal definition of "mixing time" and a formal proof that this implies fast mixing in this class. But the above is enough to claim that the mixing time of this Markov chain is $O(1)$.

Sequential Deliberation Process
$\mathrm{S_1\ S_2\ ...\ S_t\ S_{t+1}?}$
$\mathrm{u_1\ u_2\ ...\ u_t}$
$\mathrm{v_1\ v_2\ ...\ v_t}$

Another way to think about mixing time is "coupling from the past" (CFTP). Suppose the process has been run up to some very large time t. S_1 was given, but all selected agents are random and S_i for $i > 1$ are also random since the agents are randomly selected. The distribution of S_{t+1} isn't known with certainty since it depends on S_1, but the distribution of u_t, v_t, and so on is known. Because of mixing, with probability 1/3 the value of S_{t+1} is known by looking back only two steps (4 agents). Assuming S_1 is on a stationary distribution, then S_{t+1} is also on a stationary distribution. On average, you have to go back 6 steps to know what the value of S_{t+1} is for sure, and this value must be a sample from the stationary distribution. Again, we won't go into the formalism, but this is a stronger result than fast mixing - not only does the Markov chain mix in time $O(1)$ but we can also get a sample from the stationary distribution in time $O(1)$. This mixing time holds for median spaces (the distortion results from before also hold on median spaces).

11.3 Strategic Behavior

Once again, consider sequential deliberation on a line. Now, assume that each participant declares their position and the "social planner" computes the Nash bargaining outcome (the median). If the process is on the final step, it is optimal for the agents to declare their true positions since if they are the median point changing their declaration will only move the result away from their true position, while if they aren't the median point they only change the result by "crossing" the median, which would also move the result further away from their desired position. Therefore, in the final step there is no incentive to report a position other than their true position.

$$...(S_t, u_t, v_t) \xrightarrow{\text{median}=S_{t+1}} (S_{t+1}, u_{t+1}, v_{t+1}) \xrightarrow{\text{median}=S_{t+2}} (S_{t+2}, u_{t+2}, v_{t+2})...$$

Claim 11.1 *Assuming all agents are truthful in future steps, in the current period it's optimal for agents to report their positions truthfully. (This is called a sub-game perfect Nash equilibrium.)*

Proof Outline: By monotonicity. Claim that (1) S_{t+2} is monotone in S_{t+1} and (2) if $S'_{t+1} > S_{t+1}$ then either $\text{median}(S'_{t+1}, u_{t+1}, v_{t+1}) = \text{median}(S_{t+1}, u_{t+1}, v_{t+1})$ or $\text{median}(S'_{t+1}, u_{t+1}, v_{t+1}) >= S_{t+1}$. Assume $u_t < S_{t+1}$. S_{t+1} only changes if u_t falsely reports s.t. $u'_t > S_{t+1}$, so $S'_{t+1} > S_{t+1}$, which is obviously worse for u_t. This reasoning can be applied repeatedly, so misreporting now is never beneficial in the future. ■

It's not feasible to analyze full Nash equilibrium since that would be a strategy over all time periods for all agents, so in a sense this proves the strongest equilibrium concept (sub-game perfect Nash equilibrium) we can hope to achieve without exorbitant effort.

11.4 Median Spaces

Median spaces are metric spaces where for any three points, there is a unique point that lies on three pairwise shortest paths between these. Median spaces always have a Condorcet winner (the Condorcet winner on a line is simply the median). Trees, hypercubes, grids, and lines are examples of median spaces, while triangles and disconnected spaces are not.

Consider the points A, B, C in the simple grid displayed below. The shortest path between points A and B is clearly $A - B$, and similarly the shortest path between B and C is $B - C$. There are two shortest paths between A and C, $A - B - C$ and $A - D - C$. Since it is on a shortest path between all three pairs, point B is the (unique) median point (Condorcet winner).

Consider the points A, B, C in the simple triangle displayed below. The shortest path between any two points is just the direct path between those points, so it's $A - B$ for A and B, $B - C$ for B and C, and $A - C$ for A and C. However, there is no median point since no point appears on three pairwise shortest paths (each appears twice, on the paths where it is an endpoint). This demonstrates that triangles are not median spaces.

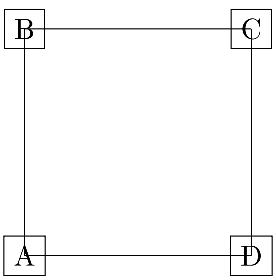

Grid - Median Space

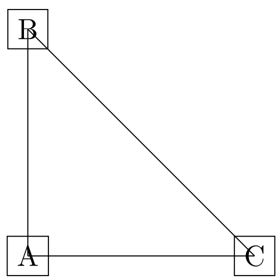

Triangle - Not Median Space

11.5 Trees

Assume there are N agents, where N is odd. Agent opinions are unique nodes in a tree of size N, denoted T. Each edge has a cost, and the cost of an opinion for an agent is the sum of the costs along edges travelled from the agent's position to that opinion. Since trees are median spaces, there is always a Condorcet winner.

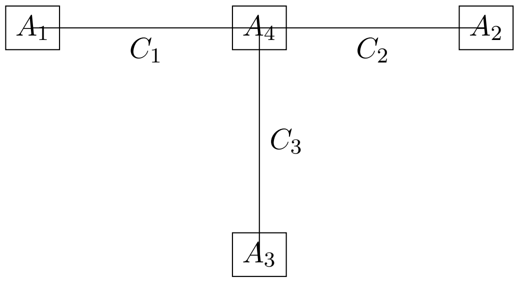

A_4 is the Condorcet winner.

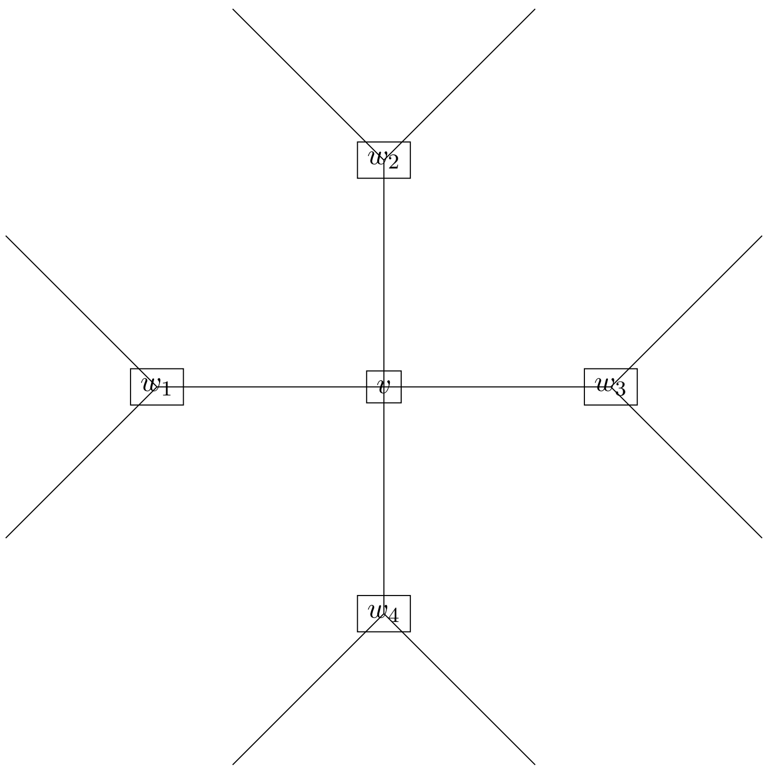

An example tree T.

To find the Condorcet winner (median) of T, start with an arbitrary node v. Denote w_i as neighbors of v and T_i as subtrees of v (we assume v is the root of each T_i). First, we consider the case where $\forall i$, we have $|T_i| < N/2$, so anyone in T_i loses pairwise elections to v, since any node not in T_i will vote for v over any node in T_i. Since v is not in any T_i, it wins pairwise elections against all other nodes in T, meaning it is the Condorcet winner. Next, consider the case where $|T_i| > N/2$ for one i (clearly this could not be true for more than one i), then consider w_i instead of v and note it's still the case that $(|T| - |T_i|) < N/2$ nodes (meaning, the number of nodes outside the subtree T_i is still less than $N/2$). For all neighbors w_i' of w_i and their subtrees T_i', if $|T_i'| < N/2$ then w_i is the Condorcet winner by the same logic as in the $|T_i| < N/2$ case, if not then repeat the process with w_i' and so on until the Condorcet winner is found. In this way, you can "walk" from any node v until you find the Condorcet winner (the process must terminate, since it is also the case that the walk never traverses an edge twice).

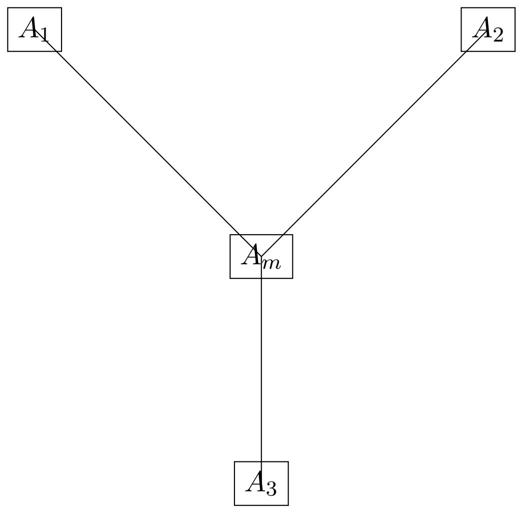

Claim 11.2 *Trees are median spaces.*

PROOF OUTLINE: Consider arbitrary nodes A_1, A_2, and A_3. Due to the tree structure, there is exactly one path (which is trivially also the shortest path) between each pair of nodes. Because there is one path between each pair, the paths between the three pairs must meet at some node, that is the median point. ■

For sequential deliberation, truthfulness on trees can be shown similar to truthfulness on lines.

11.6 Maskin's Framework

We know that truthful social choice rules don't generally exist (due to Gibbard-Satterthwaite), but can we get weaker forms of game theoretic implementation? Maskin's framework is one possibility.

Maskin's framework considers implementation of social choice rules by Nash equilibrium. The social choice rules chooses a set $f(R)$ (R is a profile of strict preferences), while the mechanism (a function g and strategy spaces for each voter, $S_1, S_2, ..., S_N$) implementing it must choose a single outcome $g(s_1, s_2, ..., s_N)$. For voter i, strategy $s_i \in S_i$. $g : S_1 \times S_2 \times ... \times S_N \to A$, where A is the set of candidates. For any alternative in $f(R)$, some pure Nash equilibrium must achieve it. Formally,

$$\begin{aligned} &\forall a \in f(R), \exists s_1, s_2, ..., s_N \ s.t. \\ &\qquad (i)\ g(s_1, ..., s_N) = a \\ &\qquad (ii)\ \forall i, \forall s_i' \neq s_i, a \succcurlyeq_{R_i} g(s_1, ..., s_i', ..., s_N) \end{aligned}$$

Additionally, for any mixed Nash equilibrium, anything in the support must result in something in $f(R)$. Meaning if $P_{(\mu_1,...,\mu_N)}(s_1,...,s_N) > 0$ (where μ_i is a mixed strategy for voter i) then $g(s_1,...,s_N) \in f(R)$.

11.7 Maskin Monotonicity

If a candidate c is in $f(R)$, and one voter changes its ranking without changing the relative position of candidate c with respect to any candidate that c dominated earlier, then c remains in $f(R')$.

Definition 11.1 *A social choice function f satisfies Maskin Monotonicity (MM) if*
$c \in f(R)$
and $\forall i, \forall b \in A, c \succcurlyeq_{R_i} b \implies c \succcurlyeq_{R'_i} b$
$\implies c \in f(R')$.

Maskin monotonicity is a necessary condition for Nash implementability - any social choice rule that is Nash implementable must satisfy MM. (Note that, for example, the Borda and Copeland rules don't satisfy MM.)