MS&E 336/CS 366: Computational Social Choice. Aut 2021-22
Course URL: `http://www.stanford.edu/~ashishg/msande336/index.html`.
Instructor: Ashish Goel, Stanford University.

Lecture 4, 9/29/2021. Scribed by Max Kanwal.

Warm-up

Q: *Is guessing how many jellybeans are in a jar via group deliberation an example of social choice?*
A: *No.* The utility function of each individual is the same: Everyone is trying to guess the correct answer. Generally, in social choice there is no ground truth (e.g., when deciding tax policy). An essential component of social choice is that, to arrive at a collective decision, one must balance a variety of utility functions that don't entirely align.

Recap

In Lecture 3, we discussed the example of a 2D landscape of *single-peaked* preferences where voter v_i has ideal preference $x_i \in \mathbb{R}^2$.

Theorem 4.1 (Black's Theorem) *If v_i's 'unhappiness' with outcome z is given by $u_i(z) = \|z - x_i\|_1$, then using the coordinate-wise median to aggregate preferences is* ***incentive-compatible*** *(i.e., voters report their ideal preferences accurately).*

(Note: In our context, **strategy-proofness** and incentive compatibility are interchangeable terms signifying that voters do not stand to benefit by misreporting their preferences.)

Remark 4.1 *Black's theorem applies more generally to all L^p-norms in any k-dimensional space of preferences (i.e., $u_i(z) = \|z - x_i\|_p$ where $x_i \in \mathbb{R}^k$). To see why, notice that the p-norm decomposes into a sum of independent components in each dimension. Therefore, the logic behind Black's Theorem for voting on a line applies to each dimension independently.*

Using the arithmetic mean to aggregate preferences does not maintain this property, therefore resulting in strategic behavior. But strategy-proofness is not the only reason why the median is a nice aggregation rule when preferences fall along a line—it's rather icing on the cake. The median is defined such that it minimizes the sum of the absolute distances to everyone's ideal preferences. Furthermore, for single-peaked preferences along a line, the median is also the Condorcet winner (i.e., it pairwise-beats all other candidates). Unfortunately, despite the generality of Black's Theorem, the median loses these more attractive latter two properties when we depart from using the L^1-norm.

Nash Equilibrium

In a single-stage game, a *pure* **Nash equilibrium** (NE) is a deterministic set of actions such that no player has any incentive to deviate. A pure NE may not always exist (e.g., in Rock-Paper-Scissors), in which case we can inject randomness and study *mixed-strategy* Nash equilibria. For exammple, in Rock-Paper-Scissors, the mixed-NE amounts to uniformly picking each option with 1/3 probability.

Theorem 4.2 (Nash) *Let player i choose action A_i from distribution D_i. Suppose that the distributions are public but not the random choices. Then, any single-shot game where strategies are known in this way will have a mixed-strategy Nash equilibrium.*

Note: Often in social choice/welfare functions, pathological NE exist around *tragedy of the commons* scenarios. For example, at stopped railroad crossings in India, drivers often try to cut ahead of others by using the opposite lane; however, by the time the train has crossed, both sides will be occupying all lanes, resulting in gridlock.

Bargaining

Bargaining involves two players A and B who are given a default outcome Z, but if they can agree on X through negotiation, then they will both get X. Analyzing the Nash equilibria doesn't lead to much insight in this context: If both agree on X, then X is a NE; if both cannot agree, then Z is a NE.

Instead, if we want to model how people will behave—or determine how ought they to behave—we need to make further assumptions.

Nash's Axioms

- *Invariance to Affine Transformations*: Effectively, if we change the units of measurement (e.g., \$ → ¢), the outcome should not change.
- *Pareto-optimality*: If both players prefer X to Z, Z should not be the outcome.
- *Symmetry*: If the preferences of each player are swapped, the outcome should not change.
- *Independence of Irrelevant Alternatives* (IIA): Had we removed an alternative that was not selected as the outcome after bargaining, the outcome should not change.

Together these axioms lead to a unique solution known as **Nash bargaining**: The unique solution is the one that maximizes the product of the *additional* utility each player receives.

Example 4.3 *Consider the following utilities for two children A and B determining where to eat:*

	A	B
*Chipotle**	*10*	*0*
Subway	*12*	*10*
Pizza Hut	*20*	*2*
Burger King	*19*	*7*
PF Chang's	*16*	*9*

**Suppose that the default option is Z = Chipotle. Below is a plot of the additional utility (above Z) that each player receives under each alternative:*

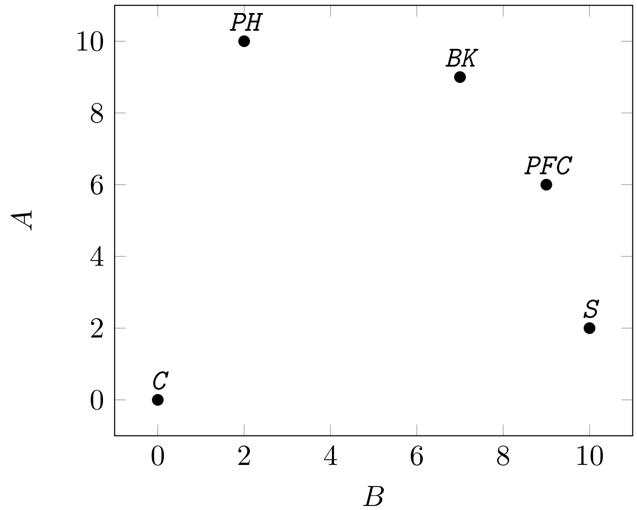

Then according to Nash bargaining, Burger King should be the outcome given that it maximizes the product of additional utilites across A and B.

This example brings to light why IIA might be a shaky assumption to make about real-life bargaining. If Pizza Hut were eliminated as an option, then IIA states that the outcome should still be Burger King. To many, this seems unlikely: BK is A's top choice, S is B's top choice, and PFC looks to be the natural compromise.

The upside of the Nash bargaining formulation is that it gives us a way to analyze bargaining by assuming some behavior. For example, it suggests an important life lesson: Whoever is happier with Z has more bargaining power.