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 14, 2/26/2020. Scribed by Sagnik Majumder.

In this lecture, we discuss fairness in the context of social-choice and allocation problems; what follows is a study of the epistemic notion of fairness as opposed to participatory ones, insofar as we shall limit ourselves to focusing on the *outcomes* of the democratic process at hand. We begin with the specification of a generic allocation problem and analyze the following axioms of fairness in reference to our chosen problem.

In an economy, there exist n agents with utility functions $U_i : S \to \mathbb{R}_{\geq 0}$ for agents $i = 1, \ldots, n$ and a space S of solutions to the allocation problem. We proceed with an axiomatic approach to describing this space of solutions, so as to endow it with certain 'good' behaviors we believe intuitively should be true. Some of these axioms are very broad in scope, in that they can be meaningfully expressed in any model:

1. **Pareto Efficiency**: An allocation $x \in S$ is Pareto-optimal if there does not exist another allocation $x' \in S$ such that each agent finds x' at least as desirable as x whilst some agent strictly prefers x' to x. In terms of utilities, $x \in S$ is said to be Pareto-optimal if there does not exist $x' \in S$ such that $U_i(x') \geq U_i(x)$ for every $i = 1, \ldots, n$, and for some $i \in \{1, \ldots, n\}$, we have $U_i(x') > U_i(x)$.

Another starting point in classical fair division problems is to bound the physical assignments and/or utilities that each agent can receive. Lower bounds guarantee that an agent's social welfare will be above a certain 'fair' level, as determined by their social endowments and preferences. Upper bounds guarantee that no agent is too much of a burden to the economy. Some of these sensible bounds are actually independent of the agents' preferences:

2. **Preference Agnosticity**: For classical fair division problems, an allocation often satisfies *no-domination of-or-by equal division*. As an illustration consider a solution to the allocation problem, an $n \times k$ matrix A where A_{ij} is the amount of resource j allocated to agent i. Furthermore assume that the total amount of any resource available is 1. This guarantees

$$\sum_{i=1}^{n} A_{ij} \leq 1, \quad \forall 1 \leq j \leq k.$$

   The allocation to agent i, namely A_i is said to be dominated by equal division if $A_{ij} \leq \frac{1}{N}$ for every resource j and for some resource j, we have the strict inequality $A_{ij} < \frac{1}{N}$. Similarly, A_i dominates equal division if $A_{ij} \geq \frac{1}{N}$ for every resource j and for some resource j, we have the strict inequality $A_{ij} > \frac{1}{N}$.

By contrast, the remaining axioms will largely account for the agents' utilities. We begin with some *punctual axioms*; these are requirements that can be enforced upon agents separately, at an 'intra-personal' level. The first of these is a notion almost as pervasive across models as Pareto efficiency

3. **Envy-Freeness**: According to this axiom, no agent should prefer another agent's allocation to their own. Namely, we must have $U_i(A_i) \geq U_i(A_j)$ for every i, j. The no-envy notion makes the most sense when applied to the distribution or redistribution of private resources. It is silent in models of "pure" public choice, where by definition, all agents consume the same thing and a compromise has to be found on what that common assignment should be because agents differ in how they value it.

4. **Proportionality**: This property, in the context of continuous fair division problems states that an agent must obtain at least as much utility from an allocation as they would if they received their proportional share of each resource. Formally, this translates to

$$\forall i = 1, \ldots, n, \quad U_i(A_i) \geq U_i \underbrace{\left(\frac{1}{n}, \ldots, \frac{1}{n}\right)}_{k}.$$

Neither envy-freeness nor proportionality is sufficient to guarantee us preference agnosticity; as an example, each agent may only care about one specific resource and thus never be preference agnostic. Note also that in social choice problems, while Pareto optimality is still a sensible notion, Axioms 2 and 3 are quite meaningless. Proportionality however can sometimes be a meaningful concept, for instance in the case of participatory budgeting. We proceed now to notions of group-fairness; axioms that cannot be verified at an individual level but requires a group of people to know their own utilities to ascertain.

5. **Core**: The concept of an allocation in the core is a generalization of Pareto-optimality and proportionality to groups of individuals; it requires that each subset of agents receive an allocation that is fair relative to its size. An allocation A is in the core if for all sets of agents W and allocations A' that satisfy

   $$\forall\, 1 \leq j \leq k, \quad \sum_{i \in W} A'_{ij} \leq \frac{|W|}{n},$$

   there exists some agent $A_i \in W$ such that $U_i(A'_i) \leq U_i(A_i)$. An allocation that does not satisfy this property is called a *blocking allocation.*

   It's not clear a priori that core-solutions exist. A solution in the core however is necessarily Pareto-optimal and proportional; this follows by considering the absence of blocking allocations of size n and 1 respectively. It is important to note that the core provides a guarantee for every possible subset of agents. Hence, in satisfying the guarantee for a set W, a solution cannot simply make a single member $i \in W$ happy and ignore the rest as this would likely violate the guarantee for the set $W - \{i\}$.

We proceed now to a discussion of notions of fairness that require inter-personal comparisons of utilities.

6. **Maxmin**: An allocation is said to be maxmin if the utility of the poorest agent is as high as possible, that is,

$$\max_{s \in S} \min_{i=1}^{n} U_i(s).$$

Note that a maxmin allocation is not necessarily Pareto optimal since it only focuses on the welfare of the poorest agent; in particular, it may be the case that keeping the poorest agent's utility constant, one can obtain strictly better utility for the rest of the population. However, it is easy to see that at least one maxmin solution satisfies Pareto optimality.

7. **Lexicographic Fairness**: This is an attempt to find Pareto-optimal maxmin allocations, specifically to avoid the situation where the poorest agent drives down the general social welfare. Consider the constrained problem $Cu \leq b$, where C is an $K \times N$ constraint matrix, u a vector of utilities and b the bounds on the constraints. We arrange the utilities $u_1, \ldots, u_n$ in non-decreasing order and maximize the utility of the poorest agent followed by the second poorest agent (holding the poorest agent's utility fixed) and so on. Formally this proceeds as

$$\beta_1 = \max_{s \in S} \min_{i=1}^{n} u_i(s).$$

This is the first step of computing the lexicographically fair solution and is identical to the maxmin solution. This can be formulated as an LP as follows

$$\begin{aligned} &\text{maximize } \beta \\ &\text{subject to } \beta \leq u \\ &\qquad\qquad\quad Cu \leq b. \end{aligned}$$

The next step is to maximize the second smallest utility subject to β_1 being the maxmin objective. To make the lexicographically fair solution envy-free, we can often add envy freeness as a constraint in the optimization problem, especially when utilities are simple functions. For example, if utilities are linear functions of allocation, say $u_i = b_i^\top A_i$, then envy freeness is the set of constraints

$$\forall 1 \leq i \neq i' \leq n, \sum_j b_{ij} A_{ij} \geq \sum_j b_{ij} A_{i'j}.$$

Note that these notions requiring interpersonal comparisons of utility still don't require knowledge of the exact utilities. Applying any monotonic transformation on utilities keeps these properties unchanged. There are however more complex notions of fairness. We end the lecture with a discussion of proportional fairness (via Nash welfare). The solution of the Nash welfare problem guarantees envy-freeness (as the solution to Fisher market equilibrium, if an agent prefers another agent's bundle, they would simply trade). We prove this for a single resource allocation problem with linear utilities. Namely, let x be an allocation vector for a single resource and suppose that

$U_i(x) = a_i x_i$ for each agent A_i. The solution to the constrained Nash welfare problem is obtained by

$$\begin{aligned} \text{maximize} \sum_{i=1}^{n} \log a_i x_i \\ \text{subject to} \sum_{i=1}^{n} c_i x_i \leq 1 \\ x \geq 0. \end{aligned}$$

Using the method of Lagrange multipliers, we have the Lagrangian

$$L(x) = \sum_{i=1}^{n} \log(a_i x_i) - \lambda \left(\sum_{i=1}^{n} c_i x_i - 1 \right).$$

Solving, we have

$$0 = \frac{\delta L(x_i)}{\delta x_i} = \frac{1}{x_i} - \lambda c_i \implies x_i = \frac{1}{\lambda c_i} \quad \forall 1 \leq i \leq n,$$

are optimal. Not only is it clear that the allocations are independent of the utilities, but we also note that they are inversely proportional to externality c_i posed by agent A_i on the economy.

References

[1] F. Brandt, V, Conitzer, U. Endriss, J. Lang and A. D. Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.