

Handout 13: Basics of Aggregating Preferences: Social Choice

October 3, 2017

(Based on Shoham and Leyton-Brown 2011)

Introduction

- In the preceding chapters we adopted what might be called the “agent perspective”: we asked what an agent believes or wants, and how an agent should or would act in a given situation. We now adopt a complementary, “designer perspective”: we ask what rules should be put in place by the authority (the “designer”) orchestrating a set of agents.
- A simple example of the designer perspective is voting.
- How should a central authority pool the preferences of different agents so as to best reflect the wishes of the population as a whole? It turns out that voting, the kind familiar from our political and other institutions, is only a special case of the general class of *social choice problems*.
- **Social choice is a motivational but nonstrategic theory**—agents have preferences, but do not try to camouflage them in order to manipulate the outcome (of voting, for example) to their personal advantage. (*Note*: Hmm ... think about voting in today’s world ...)

Example: Plurality Voting

To get a feel for social choice theory, consider an example in which you are babysitting three children—Will, Liam, Vic—and need to decide on an activity for them.

You can choose among going to the video arcade (*a*), playing basketball (*b*), and driving around in a car (*c*). Each kid has a different preference over these activities, which is represented as a strict total ordering over the activities and which he or she reveals to you truthfully. By $a > b$ denote the proposition that outcome *a* is preferred to outcome *b*.

Will: $a > b > c$

Liam: $b > c > a$

Vic: $c > b > a$

What should you do? One straightforward approach would be to ask each kid to vote for his or her favorite activity and then to pick the activity that received the largest number of votes. This amounts to what is called the *plurality* method. While quite standard, this method is not without problems.

For one thing, we need to select a **tie-breaking rule** (e.g., we could select the candidate ranked first alphabetically).

A more disciplined way is to hold a runoff election among the candidates tied at the top.

Even absent a tie, however, the method is vulnerable to the criticism that it does not meet the *Condorcet condition*. (*Note*: This condition states that if there exists a candidate *x* such that for all other candidates *y* at least half the voters prefer *x* to *y*, then *x* must be chosen.)

If each child votes for his or her top choice, the plurality method would declare a tie between all three candidates and, in our example, would choose a (i.e., using the alphabetical order as a tie-breaking rule).

However, the Condorcet condition would choose b , since two of the three children prefer b to a , and likewise prefer b to c .

Based on this example the Condorcet rule might seem unproblematic (and actually useful!), but now consider a similar example in which the preferences are as follows.

$Will: a > b > c$
 $Liam: b > c > a$
 $Vic: c > a > b$

(**Note:** Think why we do not vote for a government leader in this manner!)

In this case the Condorcet condition does not tell us what to do, illustrating the fact that it does not tell us how to aggregate arbitrary sets of preferences.

- **Social choice is not a straightforward matter.**

Formal Model

Let $N = \{1, 2, \dots, n\}$ denote a set of agents, and let O denote a finite set of outcomes (or alternatives, or candidates).

Denote the proposition that agent i weakly prefers outcome o_1 to outcome o_2 by $o_1 \succsim_i o_2$.

We use the notation $o_1 \succ_i o_2$ to capture strict preference (shorthand for $o_1 \succsim_i o_2$ **and not** $o_2 \succ_i o_1$) and $o_1 \sim_i o_2$ to capture indifference (shorthand for $o_1 \succsim_i o_2$ and $o_2 \succsim_i o_1$).

Because preferences are transitive, an agent's preference relation induces a *preference ordering*, a (nonstrict) total ordering on O .

Let L_- be the set of nonstrict total orders; we will understand each agent's preference ordering as an element of L_- .

Overloading notation, we also denote an element of L_- using the same symbol we used for the relational operator: $\succsim_i \in L_-$.

Likewise, we define a *preference profile* $[\succsim] \in L_-^n$ as a tuple giving a preference ordering for each agent.

We can define an ordering $\succsim_i \in L_-$ in terms of a given utility function $u_i: O \rightarrow R$ for an agent i by requiring that o_1 is weakly preferred to o_2 if and only if $u_i(o_1) \geq u_i(o_2)$.

Definition 9.2.1 (Social choice function) A social choice function (over N and function O) is a function $C: L_-^n \rightarrow O$.

(**Note:** A function that maps preferences to outcomes.)

A *social choice correspondence* differs from a social choice function **only in that it can return a set of candidates, instead of just a single one.**

In the babysitting example, the *social choice correspondence* defined by plurality voting picks the subset of candidates with the most votes (all). Plurality is turned into a *social choice function* by any deterministic tie-breaking rule (e.g., alphabetical).

Definition 9.2.2 (Social choice correspondence) A social choice correspondence (over N and function O) is a function $C: L_-^n \rightarrow 2^O$.

Let $\#(o_i \succ o_j)$ denote the number of agents who prefer outcome o_i to outcome o_j under preference profile $[\succsim] \in L_-^n$.

Definition 9.2.3 (Condorcet winner) An outcome $o \in O$ is a Condorcet winner if $\forall o' \in O, \#(o \succ o') \geq \#(o' \succ o)$.

(*Note:* A social choice function satisfies the *Condorcet condition* if it always picks a Condorcet winner when one exists. We saw earlier that for some sets of preferences there does *not* exist a Condorcet winner. Thus, the Condorcet condition does not always tell us anything about which outcome to choose.)

Definition 9.2.4 (Smith set) The Smith set is the smallest set $S \subseteq O$ having the property that $\forall o \in S, \forall o' \notin S, \#(o \succ o') \geq \#(o' \succ o)$.

(*Note:* That is, every outcome *in* the Smith set is preferred by at least half of the agents to every outcome *outside* the set. This set *always* exists. When there is a Condorcet winner then that candidate is also the only member of the Smith set; otherwise, the Smith set is the set of candidates who participate in a “stalemate” (or “top cycle”).)

The other important flavor of social function is the *social welfare function*. These are similar to **social choice functions, but produce richer objects, total orderings on the set of alternatives.**

Definition 9.2.5 (Social welfare function) A social welfare function (over N and function O) is a function $W: L_-^n \rightarrow L_-$.

(*Note:* A function that maps multiple orderings into one.)