

Voting: Preference Aggregating & Social Choice

(Based on Shoham and Leyton-Brown (2008). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge.)

Leen-Kiat Soh

How do we *know* as a society what we want?

How to find out what the choice (or preference order) of a group of agents?



What's the point? If we (or the agents) know the goals of the group, agents can self-organize and plan accordingly.



Tasks: setting prices of resources, designing fees and taxes, allocating resources; Applications: traffic management, cybersecurity, etc.

So far, our discussions have been focused on 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 look at the “designer perspective”

- **We ask what rules should be put in place by the authority (the “designer”) orchestrating a set of agents to achieve the designer’s objectives**
- **Voting, mechanism design, etc.**

Introduction

- How should a central authority pool the preferences of different agents so as to best reflect the wishes of the population as a whole?
 - Voting is 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
 - Or it would defeat the purpose of trying to find out what the social choice is

Hmm ... think about today's voting mechanisms



**Social choice is *NOT* a
straightforward matter**

Example | Plurality Voting

- Suppose you are babysitting three kids—Will, Liam, Vic—and need to decide on an activity for them
 - Choices are 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
 - $a \succ b$ denotes the proposition that outcome a is preferred to outcome b

Example | Plurality Voting 2

Will: $a \succ b \succ c$

Liam: $b \succ c \succ a$

Vic: $c \succ b \succ a$

Their individual
preference ordering

- **What should you do to identify the social choice?**
- 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**
 - the **plurality** method

Any potential problems with this voting method?



Example | Plurality Voting 3

Will: $a > b > c$

Liam: $b > c > a$

Vic: $c > b > a$

- In case of a tie
 - A **tie-breaking rule** (e.g., we could select the candidate ranked first alphabetically)
 - A more disciplined way: **hold a runoff election among the candidates tied at the top**
- Even absent a tie, the method might not meet the **Condorcet condition**
- 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**
- Assume that each kid votes for his or her top choice, and the **plurality method** would declare a tie between all three candidates
 - Alphabetical order tie-breaker: a wins
 - Condorcet condition tie-breaker: b wins, since two of the three children prefer b to a , and likewise prefer b to c .

Example | Plurality Voting 4

- The Condorcet rule might seem unproblematic (and actually useful!), but now consider a similar example in which the preferences are as follows:

Will: $a \succ b \succ c$

Liam: $b \succ c \succ a$

Vic: $c \succ a \succ b$

- In this case the Condorcet condition does **not** tell us what to do

Think why we do not vote for a government leader in this manner!



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

Formal Model 2

- 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 $u_i(o_1) \geq u_i(o_2)$.**

Formal Model 3

- **Definition 9.2.1 Social choice function.** A social choice function (over N and function O) is a function $C: L_-^n \rightarrow O$.
 - A function that maps preferences to outcomes
- **Definition 9.2.2 Social choice correspondence.** A social choice correspondence (over N and function O) is a function $C: L_-^n \rightarrow 2^O$.
 - 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).

Formal Model 4

- 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 $\forall o' \in O, \#(o \succ o') \geq \#(o' \succ o)$.*
 - A social choice function satisfies the *Condorcet condition* if it always picks a Condorcet winner **when** one exists
 - 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

And preference ordering is needed from every voter/agent



Formal Model 5

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

How to make use
of the Smith set?



Formal Model 6

- The other **important** flavor of social function is the ***social welfare function***
- 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 $W: L_-^n \rightarrow L_-$.
 - A function that maps multiple orderings into one ordering



Examples of this?

- Think BCS polls
- Think Top N lists

Connection to MAS?

How to solicit and aggregate preferences from a group of agents to obtain its social choice, or social welfare?

Do we ask for voters for their top pick? Or do we ask voters for their preference ordering?



Environment is dynamic, uncertain, incomplete, etc., how do we know what the agents, as a whole, want?

How to motivate agents to be truthful in revealing their preferences? (Hint: Mechanism Design)

Silly Question: Suppose that a TV station is collecting votes to come up with a list of Top-10 songs. To encourage audience participation, it will have a draw of prizes for the voters who vote for the Top-3 songs.

What would happen?

