CHAPTER 12: MAKING GROUP DECISIONS

An Introduction to Multiagent Systems

`http://www.csc.liv.ac.uk/~mjw/pubs/imas/`

Social Choice

- *Social choice theory* is concerned with *group decision making*.
- Classic example of social choice theory: *voting*.
- Formally, the issue is *combining preferences* to derive a *social outcome*.

Components of a Social Choice Model

- Assume a set $Ag = \{1, \ldots, n\}$ of *voters*.
  These are the entities who will be expressing preferences.
- Voters make group decisions wrt a set $\Omega = \{\omega_1, \omega_2, \ldots\}$ of *outcomes* .
  Think of these as the *candidates*.
- If $|\Omega| = 2$, we have a *pairwise election*.

Preferences

- Each voter has preferences over Ω: an *ordering* over the set of possible outcomes Ω.
- Example. Suppose

$$\Omega = \{gin, rum, brandy, whisky\}$$

then we might have agent mjw with preference order:

$$\varpi_{mjw} = (brandy, rum, gin, whisky)$$

meaning

$$brandy \succ_{mjw} rum \succ_{mjw} gin \succ_{mjw} whisky$$

Preference Aggregation

The fundamental problem of social choice theory:

> *given a collection of preference orders, one for each voter, how do we combine these to derive a group decision, that reflects as closely as possible the preferences of voters*?

Two variants of preference aggregation:

- *social welfare functions*;
- *social choice functions*.

Social Welfare Functions

- Let $\Pi(\Omega)$ be the set of preference orderings over Ω.
- A *social welfare function* takes the voter preferences and produces a *social preference order*:

$$f : \underbrace{\Pi(\Omega) \times \cdots \times \Pi(\Omega)}_{n \text{ times}} \to \Pi(\Omega).$$

- We let $\succ^*$ denote to the outcome of a social welfare function
- Example: beauty contest.

Social Choice Functions

- Sometimes, we just one to select *one* of the possible candidates, rather than a social order.
- This gives *social choice functions*:

$$f : \underbrace{\Pi(\Omega) \times \cdots \times \Pi(\Omega)}_{n \text{ times}} \to \Omega.$$

- Example: presidential election.

Voting Procedures: Plurality

- Social choice function: selects a single outcome.
- Each voter submits preferences.
- Each candidate gets one point for every preference order that ranks them first.
- Winner is the one with largest number of points.
- Example: Political elections in UK.
- If we have only two candidates, then plurality is a *simple majority election*.

Anomalies with Plurality

- Suppose $|Ag| = 100$ and $\Omega = \{\omega_1, \omega_2, \omega_2\}$ with:

  40% voters voting for ω_1

  30% of voters voting for ω_2

  30% of voters voting for ω_3

- With plurality, ω_1 gets elected even though a *clear majority* (60%) prefer another candidate!

Strategic Manipulation by Tactical Voting

- Suppose your preferences are

$$\omega_1 \succ_i \omega_2 \succ_i \omega_3$$

while you believe 49% of voters have preferences

$$\omega_2 \succ_i \omega_1 \succ_i \omega_3$$

and you believe 49% have preferences

$$\omega_3 \succ_i \omega_2 \succ_i \omega_1$$

- You may do better voting for ω_2, *even though this is not your true preference profile*.
- This is *tactical voting*: an example of *strategic manipulation* of the vote.

Condorcet's Paradox

- Suppose $Ag = \{1, 2, 3\}$ and $\Omega = \{\omega_1, \omega_2, \omega_3\}$ with:

$$\omega_1 \succ_1 \omega_2 \succ_1 \omega_3$$
$$\omega_3 \succ_2 \omega_1 \succ_2 \omega_2$$
$$\omega_2 \succ_3 \omega_3 \succ_3 \omega_1$$

- For every possible candidate, there is another candidate that is preferred by a majority of voters!
- This is *Condorcet's paradox*: there are situations in which, *no matter which outcome we choose, a majority of voters will be unhappy with the outcome chosen*.

Sequential Majority Elections

A variant of plurality, in which players play in a series of rounds: either a *linear* sequence or a *tree* (knockout tournament).

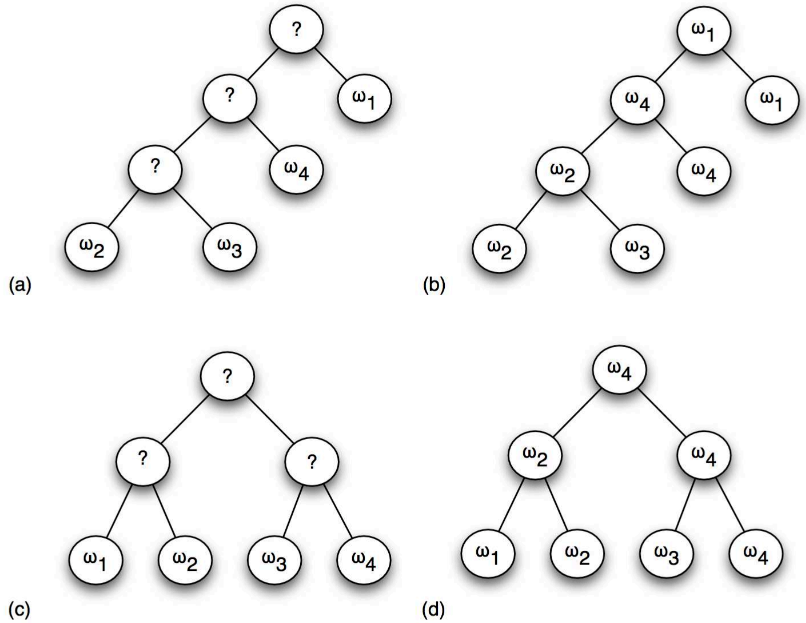

Linear Sequential Pairwise Elections

- Here, we pick an ordering of the outcomes – the *agenda* – which determines who plays against who.
- For example, if the agenda is:

$$\omega_2, \ \omega_3, \ \omega_4, \ \omega_1.$$

then the first election is between ω_2 and ω_3, and the winner goes on to an election with ω_4, and the winner of this election goes in an election with ω_1.

Anomalies with Sequential Pairwise Elections

Suppose:

- 33 voters have preferences

$$\omega_1 \succ_i \omega_2 \succ_i \omega_3$$

- 33 voters have preferences

$$\omega_3 \succ_i \omega_1 \succ_i \omega_2$$

- 33 voters have preferences

$$\omega_2 \succ_i \omega_3 \succ_i \omega_1$$

Then *for every candidate, we can fix an agenda for that candidate to win in a sequential pairwise election!*

Majority Graphs

- This idea is easiest to illustrate by using a *majority graph*.
- A directed graph with:

  vertices = candidates

  an edge (i, j) if i would beat j is a simple majority election.
- A *compact representation of voter preferences*.

Majority Graph for the Previous Example

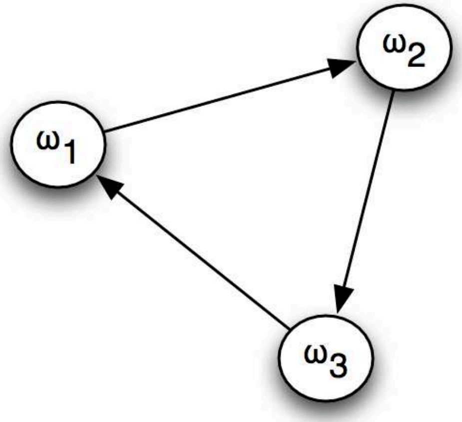

with agenda $(\omega_3, \omega_2, \omega_1)$, ω_1 wins
with agenda $(\omega_1, \omega_3, \omega_2)$, ω_2 wins
with agenda $(\omega_1, \omega_2, \omega_3)$, ω_3 wins

Another Majority Graph

Give agendas for each candidate to win with the following majority graph.

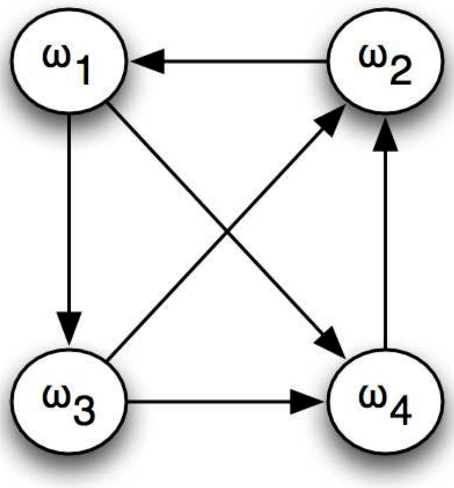

Condorcet Winners

A *Condorcet winner* is a candidate that would beat *every other candidate* in a pairwise election.
Here, ω_1 is a Condorcet winner.

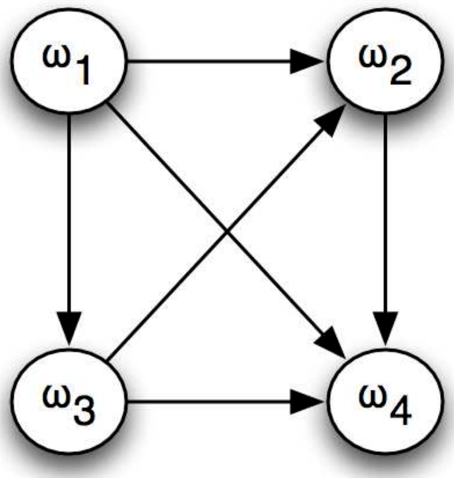

Voting Procedures: Borda Count

- One reason plurality has so many anomalies is that it *ignores* most of a voter's preference orders: it only looks at the *top ranked candidate*.
- The *Borda* count takes *whole* preference order into account.

- For each candidate, we have a variable, counting the strength of opinion in favour of this candidate.
- If ω_i appears first in a preference order, then we increment the count for ω_i by $k-1$; we then increment the count for the next outcome in the preference order by $k-2$, . . . , until the final candidate in the preference order has its total incremented by 0.
- After we have done this for all voters, then the totals give the ranking.

Desirable Properties of Voting Procedures

Can we classify the properties we want of a "good" voting procedure?
Two key properties:

- *The Pareto property*;
- *Independence of Irrelevant Alternatives (IIA)*.

The Pareto Property

If everybody prefers ω_i over ω_j, then ω_i should be ranked over ω_j in the social outcome.

Independence of Irrelevant Alternatives (IIA)

Whether ω_i is ranked above ω_j in the social outcome should depend only on the relative orderings of ω_i and ω_j in voters profiles.

Arrow's Theorem

For elections with more than 2 candidates, the only voting procedure satisfying the Pareto condition and IIA is a dictatorship, *in which the social outcome is in fact simply selected by one of the voters.*

This is a *negative* result: there are fundamental limits to democratic decision making!

Strategic Manipulation

- We already saw that sometimes, voters can benefit by *strategically misrepresenting their preferences*, i.e., lying – tactical voting.
- Are there any voting methods which are *non-manipulable*, in the sense that voters can *never* benefit from misrepresenting preferences?

The Gibbard-Satterthwaite Theorem

The answer is given by the Gibbard-Satterthwaite theorem:

> *The only non-manipulable voting method satisfying the Pareto property for elections with more than 2 candidates is a dictatorship.*

In other words, every “realistic” voting method is prey to strategic manipulation . . .

Computationally Complexity to the Rescue!

- Gibbard-Satterthwaite only tells us that manipulation is possible *in principle*.

  It does not give any indication of *how* to misrepresent preferences.
- Bartholdi, Tovey, and Trick showed that there are elections that are prone to manipulation in principle, but where manipulation was *computationally complex*.
- “Single Transferable Vote” is NP-hard to manipulate!