

# Multiagent Systems

## 9. Social Choice

B. Nebel, C. Becker-Asano, S. Wölfl

Albert-Ludwigs-Universität Freiburg

June 25, 2014

# Motivation

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Making group decisions

- Previously we looked at agents acting strategically
- Outcome in normal-form games follows immediately from agents' choices
- Often a mechanism for deriving group decision is present
- What rules are appropriate to determine the joint decision given individual choices?
- **Social Choice Theory** is concerned with group decision making (basically analysis of mechanisms for voting)
- Basic setting:
  - Agents have preferences over outcomes
  - Agents vote to bring about their most preferred outcome

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Preference Aggregation

# Preference aggregation

Setting:

- $Ag = \{1, \dots, n\}$ : voters (finite, odd number)
- $\Omega = \{\omega_1, \omega_2, \dots\}$ : possible **outcomes** or **candidates**
- $\Pi(\Omega)$ : set of all (strict) preference orderings over  $\Omega$
- $\succsim_i \in \Pi(\Omega)$ : preference ordering of agent  $i$

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Preference aggregation

Setting:

- $Ag = \{1, \dots, n\}$ : voters (finite, odd number)
- $\Omega = \{\omega_1, \omega_2, \dots\}$ : possible **outcomes** or **candidates**
- $\Pi(\Omega)$ : set of all (strict) preference orderings over  $\Omega$
- $\succsim_i \in \Pi(\Omega)$ : preference ordering of agent  $i$

## Preference aggregation

How do we combine a collection of potentially different preference orders in order to derive a group decision?

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Preference aggregation

Task is either to derive a globally acceptable preference ordering, or determine a winner:

## Social welfare/choice functions

- **Social welfare function:** a function that assigns to  $n$  preference relations (one for each agent) a preference relation, i.e.:

$$F: \Pi(\Omega) \times \cdots \times \Pi(\Omega) \rightarrow \Pi(\Omega)$$

- **Social choice function:** a function that assigns to  $n$  preference relations (one for each agent) a candidate, i.e.:

$$f: \Pi(\Omega) \times \cdots \times \Pi(\Omega) \rightarrow \Omega$$

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wöfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Plurality voting

- Voters submit preference orders
- The candidate that appears first in most preference orders wins
- Only submission of the highest ranked candidate is required
- **Simple majority voting** when  $|\Omega| = 2$

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary



# Plurality voting

- Voters submit preference orders
- The candidate that appears first in most preference orders wins
- Only submission of the highest ranked candidate is required
- **Simple majority voting** when  $|\Omega| = 2$

Advantages: simple to implement and easy to understand

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Plurality voting

- Voters submit preference orders
- The candidate that appears first in most preference orders wins
- Only submission of the highest ranked candidate is required
- **Simple majority voting** when  $|\Omega| = 2$

Advantages: simple to implement and easy to understand

Problems:

- Tactical voting
- Strategic manipulation
- Condorcet's paradox

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Plurality voting: An example

- Outcomes:  $\Omega = \{S, G, C\}$
- Assume 1000 voters with the following preference relations:

# Voters	Ranking		
	1	2	3
417	S	G	C
142	G	S	C
441	C	G	S

- Plurality voting:  $C$  wins with 44% of the votes.

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Anomalies with Plurality

- Despite not securing majority,  $C$  wins with 44%
- Even worse:  $C$  is the least preferred option for 56% of voters
- **Tactical voting**: The voters with  $G \succ S \succ C$  may do better by voting for  $S$  instead of their actual preference  $G$
- But is lying bad? Not in principle, but it favours computationally stronger voters, and wastes computational resources
- **Strategic nomination**: manipulate the voting result through the candidate set

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Condorcet's Paradox

- Outcomes:  $\Omega = \{A, B, C\}$
- 3 voters with the following preference orders:

$$A \succ_1 B \succ_1 C,$$

$$C \succ_2 A \succ_2 B; \text{ and}$$

$$B \succ_3 C \succ_3 A$$

- With plurality voting, no decision (a tie)

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Condorcet's Paradox

- Outcomes:  $\Omega = \{A, B, C\}$
- 3 voters with the following preference orders:

$$\begin{aligned}A &\succ_1 B \succ_1 C, \\C &\succ_2 A \succ_2 B; \text{ and} \\B &\succ_3 C \succ_3 A\end{aligned}$$

- With plurality voting, no decision (a tie)
- **Condorcet's Paradox:** The social preference is **not transitive** though all individual preference orderings are transitive

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Condorcet's Paradox

- Outcomes:  $\Omega = \{A, B, C\}$
- 3 voters with the following preference orders:

$$A \succ_1 B \succ_1 C,$$

$$C \succ_2 A \succ_2 B; \text{ and}$$

$$B \succ_3 C \succ_3 A$$

- With plurality voting, no decision (a tie)
- **Condorcet's Paradox:** The social preference is **not transitive** though all individual preference orderings are transitive
- In the example:  $A$  is (more often) preferred to  $B$  and  $B$  is preferred to  $C$ , but  $A$  is not preferred to  $C$

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Condorcet's Paradox

- Outcomes:  $\Omega = \{A, B, C\}$
- 3 voters with the following preference orders:

$$\begin{aligned}A &\succ_1 B \succ_1 C, \\C &\succ_2 A \succ_2 B; \text{ and} \\B &\succ_3 C \succ_3 A\end{aligned}$$

- With plurality voting, no decision (a tie)
- **Condorcet's Paradox:** The social preference is **not transitive** though all individual preference orderings are transitive
- In the example:  $A$  is (more often) preferred to  $B$  and  $B$  is preferred to  $C$ , but  $A$  is not preferred to  $C$
- This means: There are scenarios in which no matter which outcome we choose the majority of voters will be unhappy

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary



# Plurality voting with runoff

- First round: two candidates with the most top highest ranked votes are selected unless one candidate receives absolute majority
- Second round: runoff

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Plurality voting with runoff

- First round: two candidates with the most top highest ranked votes are selected unless one candidate receives absolute majority
- Second round: runoff

In the plurality voting example: candidates  $S$  and  $C$  go to the runoff-round. Given all voters stick to their preferences,  $S$  wins over  $C$ .

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation


Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Sequential majority elections

- Instead of one-step protocol, voting can be done in several steps
- Candidates face each other in **pairwise elections**, the winner progresses to the next election
- **Election agenda** is the ordering of these elections
- Can be organized as a binary voting tree:



- **Key problem:** the final outcome depends on the election agenda

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem


Gibbard-  
Satterthwaite  
Theorem

Summary

# Majority graphs

- Introduce **majority graph** as tool for discussing sequential voting: provides a succinct representation of voter preferences
- Nodes correspond to outcomes, e.g.,  $\omega_1, \omega_2, \dots$
- There is an arc from  $\omega$  to  $\omega'$  whenever a majority of voters rank  $\omega$  above  $\omega'$

Examples:



Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation




Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Majority graphs

- **Tournament:** complete, asymmetric and irreflexive majority graph (produced with odd number of voters)
- **Possible winner:** There is an agenda in which the candidate wins  
E.g.: every candidate in (a) and (b)
- **Condorcet winner:** overall winner for every possible agenda  
E.g.: candidate  $\omega_1$  in (c)
- **Strategic manipulation:** fixing the election agenda



Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# The Borda Count

- In simple mechanisms above, only top-ranked candidate taken into account, rest of orderings disregarded
- **Borda count** looks at entire preference ordering, counts the strength of opinion in favour of a candidate
- For all preference orders and outcomes ( $|\Omega| = m$ ), if  $\omega_i$  has rank  $k$  in a preference ordering,  $\omega_i$  gets  $m - k$  points. Then up sum all points. Candidate with most points wins.

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# The Borda Count

- In simple mechanisms above, only top-ranked candidate taken into account, rest of orderings disregarded
- **Borda count** looks at entire preference ordering, counts the strength of opinion in favour of a candidate
- For all preference orders and outcomes ( $|\Omega| = m$ ), if  $\omega_i$  has rank  $k$  in a preference ordering,  $\omega_i$  gets  $m - k$  points. Then up sum all points. Candidate with most points wins.

Voting example:

- 417 voters with  $S \succ G \succ C$ ; 142 voters with  $G \succ S \succ C$ , and 441 voters with  $C \succ G \succ S$ .
- Borda counts:  
S:  $417 \cdot (3 - 1) + 142 \cdot (3 - 2) + 441 \cdot (3 - 3) = 834 + 142 = 976$   
G:  $417 \cdot (3 - 2) + 142 \cdot (3 - 1) + 441 \cdot (3 - 2) = 417 + 284 + 441 = 942$   
C:  $417 \cdot (3 - 3) + 142 \cdot (3 - 3) + 441 \cdot (3 - 1) = 882$

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem


Gibbard-  
Satterthwaite  
Theorem

Summary

# The Slater Ranking

- Idea: how can we minimize disagreement between the majority graph and the social choice?
- For each possible ordering measure the degree of disagreement with the majority graph
- Degree of disagreement: number of edges that need to be flipped
- NP-hard to compute

Example:



Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem


Summary



# The Slater Ranking

- Idea: how can we minimize disagreement between the majority graph and the social choice?
- For each possible ordering measure the degree of disagreement with the majority graph
- Degree of disagreement: number of edges that need to be flipped
- NP-hard to compute

Example:



Consider:

- $\omega_1 \succ^* \omega_2 \succ^* \omega_4 \succ^* \omega_3$   
Cost is 2, we have to flip the edges  $(\omega_3, \omega_4)$  and  $(\omega_4, \omega_1)$
- $\omega_1 \succ^* \omega_2 \succ^* \omega_3 \succ^* \omega_4$   
Cost is 1, we have to flip the edge  $(\omega_4, \omega_1)$ ; this is the ordering with the lowest disagreement

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Arrow's Theorem

# Desirable properties (I)

- **Pareto condition/Partial unanimity:**

If every voter ranks  $\omega_i$  above  $\omega_j$ , then  $\omega_i \succ^* \omega_j$

... satisfied by plurality and Borda, but not by sequential majority

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Desirable properties (I)

- **Pareto condition/Partial unanimity:**

If every voter ranks  $\omega_i$  above  $\omega_j$ , then  $\omega_i \succ^* \omega_j$

... satisfied by plurality and Borda, but not by sequential majority

- **Condorcet winner condition:**

The outcome would beat every other outcome in a pairwise election

... satisfied only by sequential majority elections

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Desirable properties (II)

- **Independence of irrelevant alternatives (IIA):**

The social ranking of two outcomes  $\omega_i$  and  $\omega_j$  should depend only on their relevant ordering in the voters' preference orders (and not on the ordering of other outcomes)

... Plurality, Borda and sequential majority elections do not satisfy IIA

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Desirable properties (II)

- **Independence of irrelevant alternatives (IIA):**

The social ranking of two outcomes  $\omega_i$  and  $\omega_j$  should depend only on their relevant ordering in the voters' preference orders (and not on the ordering of other outcomes)

... Plurality, Borda and sequential majority elections do not satisfy IIA

- **Non-Dictatorship:** A social welfare function  $F$  is a **dictatorship** if there exists a voter  $i$  such that

$$F(\succ_1, \dots, \succ_n) = \succ_i$$

for all orderings  $\succ_1, \dots, \succ_n$ .

... Dictatorships satisfy Pareto condition and IIA

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Arrow's Theorem

- Overall vision in social choice theory: identify “good” social choice procedures
- Unfortunately, a fundamental theoretical result gets in the way

## Arrow's Impossibility Theorem

In situations with more than two possible outcomes, every social welfare function satisfying unanimity and independence of irrelevant alternatives must be a dictatorship.

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Arrow's Theorem

- Disappointing, basically means we can never achieve combination of good properties without dictatorship
- ... in other words, there exists no social welfare function that satisfies (partial) unanimity, IIA, and non-dictatorship at the same time (in situations with more than two alternatives)
- Most social welfare functions satisfy unanimity and non-dictatorship, i.e., the problem is usually IIA
- This is related to **strategic voting**: add irrelevant candidates ...

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary



# Gibbard-Satterthwaite Theorem

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Strategic Manipulation

- As stated above, while lying could be allowed as part of rational behaviour, it is unfair and wasteful
- Can we design voting procedures that are immune to manipulation?

## Incentive compatibility

A social choice function  $f$  is **manipulable** by voter  $i$  if for some collection of preference profiles  $\succ_1, \dots, \succ_n$  there exists  $\succ'_i$  such that

$$f(\succ_1, \dots, \succ'_i, \dots, \succ_n) \succ_i f(\succ_1, \dots, \succ_i, \dots, \succ_n)$$

$f$  is **incentive-compatible** if  $f$  can never be manipulated by any voter.

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Gibbard-Satterthwaite Theorem

## Dictatorship

$f$  is a **dictatorship** if there exists a voter  $i$  such that for all preference profiles  $\succ_1, \dots, \succ_n$ ,  $f(\succ_1, \dots, \succ_n)$  is the unique candidate that is most preferred w.r.t.  $\succ_i$ .

The Gibbard-Satterthwaite Theorem is an analogous result to Arrow's Impossibility Theorem: social choice functions instead of social welfare functions

**Surjective** social choice function: one that does not exclude **ex ante** any possible outcome.

## Gibbard-Satterthwaite Theorem

In situations with more than two outcomes, the only incentive-compatible and surjective social choice functions are dictatorships.

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Complexity of manipulation

- So we have another negative result: strategic manipulation is possible in principle in all desirable mechanisms
- But how easy is it to manipulate effectively?
- Distinction between being easy to compute and easy to manipulate
- Mechanisms can be designed for which manipulation is very computationally complex (but often only in the worst case)
- Are there non-dictatorial voting procedures that are easy to compute but not easy to manipulate?
- Yes, for example **second-order Copeland**

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wöfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary

# Summary

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary  
Thanks

# Summary

- Discussed procedures for making group decisions
- Plurality, Sequential Majority Elections, Borda Count, Slater Ranking
- Desirable properties
- Dictatorships
- Strategic manipulation and computational complexity
- Next time: Coalition Formation

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary  
Thanks

# Acknowledgments

These lecture slides are based on the following resources:

- Dr. Michael Rovatsos, The University of Edinburgh  
<http://www.inf.ed.ac.uk/teaching/courses/abs/abs-timetable.html>
- Michael Wooldridge: **An Introduction to MultiAgent Systems**, John Wiley & Sons, 2nd edition 2009.

Multiagent  
Systems

B. Nebel,  
C. Becker-  
Asano,  
S. Wölfl

Motivation

Preference  
Aggregation

Arrow's  
Theorem

Gibbard-  
Satterthwaite  
Theorem

Summary  
Thanks