

Computational Social Choice: Spring 2007

Ulle Endriss
 Institute for Logic, Language and Computation
 University of Amsterdam

Introduction

The course will cover issues at the interface of *computer science* (including *logic*, *multiagent systems* and *artificial intelligence*) and *mathematical economics* (including *social choice theory*, *game theory* and *decision theory*).

There has been a recent *trend* towards research of this sort. The broad philosophy is generally the same, but people have been using different names to identify various flavours of this kind of work, e.g.

- Algorithmic Game Theory
- Algorithmic Decision Theory
- Social Software
- *and*: Computational Social Choice

No specific *prerequisites* are required to follow the course. Nevertheless, we will frequently touch upon *current research* issues.

Organisational Matters

- **Lecturer:** Ulle Endriss (ulle@illc.uva.nl), Room P.316
- **Timetable:** Tuesdays 11am-1pm in Room P.014
- **Examination:** There will be several *coursework assignments* on the material covered in the course. In the second block, every student will have to study a *recent paper*, write a short essay on the topic, and present their findings in a talk.
- **Website:** Lecture slides, coursework assignments, and other important information will be posted on the course website:
<http://www.illc.uva.nl/~ulle/teaching/comsoc/>
- **Seminars:** There are occasional talks at the ILLC that are directly relevant to the course and that you are welcome to attend (e.g. at the Computational Social Choice Seminar).

Related Courses

- Cooperative Games
Krzysztof Apt
- Game Theory for Information Sciences (taught in autumn)
Peter van Emde Boas
- Logic, Games and Computation (taught in autumn)
Johan van Benthem
- Multiagent Systems and Distributed AI (MSc AI)
Marinus Maris

Plan for Today

- **Part I:** Introduction to the main topics of the course
- **Part II:** Arrow's Theorem (as an example for a classical result in social choice theory)

Part I: Course Topics

Collective Decision Making

The central problem to be addressed in this course is that of *collective decision making*: How can we map the individual preferences of a group of agents into a joint decision? More specifically, we may ask:

- What makes a “good” joint decision? \leadsto *welfare economics*
- What if the number of alternatives is very large or has a combinatorial structure? \leadsto *knowledge representation*
- Can we make sure that the individual agents are going to report their true preferences? \leadsto *game theory* (+ *complexity theory*)
- Once we have settled on a particular mechanism, how can we execute it in a reasonable amount of time? \leadsto *algorithm design*

The classical discipline for the study of collective decision making mechanisms is *social choice theory*. However, research in social choice theory has mostly neglected the *computational aspects* of collective decision making. Hence the name: *computational social choice*.

Ideas and Examples

Next we are going to present a number of examples, problems, ideas, paradoxes, or just *issues* that illustrate the main question addressed in the course:

“How does collective decision making work?”

The remainder of the course will then be devoted to developing these rather vague ideas in a rigorous manner.

Condorcet Paradox

In 1785, the Marquis de Condorcet noticed a problem ...

Suppose there are three agents who have to decide amongst three alternatives. They have the following preferences:

Agent 1: $A \succ B \succ C$

Agent 2: $B \succ C \succ A$

Agent 3: $C \succ A \succ B$

A majority prefers A over B and a majority also prefers B over C , but then again a majority prefers C over A .

So the “social preference ordering” induced by the seemingly natural majority rule fails to be rational (it’s not transitive).

M. le Marquis de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris, 1785

Vote Manipulation

Suppose the *plurality rule* (as in most real-world situations) is used to decide the outcome of an election: the candidate receiving the highest number of votes wins.

Assume the preferences of the people in, say, Florida are as follows:

49%: $\text{Bush} \succ \text{Gore} \succ \text{Nader}$

20%: $\text{Gore} \succ \text{Nader} \succ \text{Bush}$

20%: $\text{Gore} \succ \text{Bush} \succ \text{Nader}$

11%: $\text{Nader} \succ \text{Gore} \succ \text{Bush}$

So even if nobody is cheating, Bush will win in a plurality contest.

Issue: In a pairwise competition, Gore would have defeated anyone.

Issue II: It would have been in the interest of the Nader supporters to *manipulate*, *i.e.* to misrepresent their preferences.

Electing a Committee

Suppose we have to elect a *committee* (not just a single candidate).

If there are k *seats* to be filled from a pool of n *candidates*, then there are $\binom{n}{k}$ possible outcomes.

For $k = 5$ and $n = 12$, for instance, that makes 792 alternatives.

The domain of alternatives has a *combinatorial structure*.

It does not seem reasonable to ask voters to submit their full preferences over all alternatives to the collective decision making mechanism. What would be a reasonable form of balloting?

Multiagent Resource Allocation

As an instance of the general problem of collective decision making, we are going to be interested in the following type of problem:

- Allocate a set of *goods* (or tasks) amongst several *agents*.
- The agents should play an active role in the execution of the *allocation procedure*.
- Their actions may be influenced by their individual *preferences*.

Let’s look at some concrete examples for this kind of problem ...

Earth Observation Satellites

Our agents are representatives of different European countries that have jointly funded a new Earth Observation Satellite (EOS). Now the agents are requesting certain photos to be taken by the EOS, but due to physical constraints not all requests can be honoured ...

Allocations should be both *efficient* and *fair*:

- The satellite should not be underexploited.
- Each agent should get a return on investment that is at least roughly proportional to its financial contribution.

M. Lemaître, G. Verfaillie, and N. Bataille. *Exploiting a Common Property Resource under a Fairness Constraint: A Case Study*. Proc. IJCAI-1999.

Settling Divorce Disputes

Our agents used to be happily married, but have fallen out of love. How should they divide their belongings?

- I value this carpet at €2000. (*quantitative preference*)
- I'd rather have the red than the blue car. (*ordinal preference*)
- No way he's gonna get the piano! (*externalities*)
- I won't be content with less than what she gets! (*envy*)
- I'd rather kill the dog than let him have it! (*pure evil*)

S.J. Brams and A.D. Taylor. *Fair Division: From Cake-cutting to Dispute Resolution*. Cambridge University Press, 1996

First Approach: A Sealed-Bid Auction

Suppose we have just *one item* to be allocated.

Procedure: Each agent sends a letter stating how much they would be prepared to pay for the item to an independent “auctioneer”. The auctioneer awards the item to the agent bidding the highest price and takes the money.

At least if all agents are honest and send in their *true valuations* of the item, then this procedure has some appealing properties:

- The procedure is *simple*, both computationally and in terms of its communication requirements.
- The agent who likes the item the most will obtain it.
- The auctioneer will make maximum profit.
- No other solution would be better for some of the participants without being worse for any of the others (*Pareto optimality*).

Strategic Considerations

But what if the agents are not honest ... ?

What would be the best possible strategy for a *rational* agent? That is, how much should an agent bid for the item on auction?

- ▶ The procedure actually isn't simple at all for the agents, and our nice properties cannot be guaranteed.
- ▶ Try to set up an allocation mechanism giving agents an incentive to bid truthfully: *mechanism design* (\leadsto *game theory*).

Auctioning Multiple Items

Strategic issues aside, our simple auction mechanism is not that bad for allocating a single item. But what if there are *several goods*?

Suppose the auctioneer is first selling a TV and then a DVD player. How should our agents bid?

- Ann wants to watch the news and is not interested in DVDs.
- Bob already owns a TV and only wants the DVD player.
- Chloë has an enormous collection of classic movies on DVD, a pretty low opinion of today's television programming, and no DVD player or TV.

Complements and Substitutes

The value an agent assigns to a *bundle* of goods may relate to the value it assigns to the individual goods in a variety of ways ...

- *Complements*: The value assigned to a set is *greater* than the sum of the values assigns to its elements.

A standard example for complements would be a pair of shoes (a left shoe and a right shoe).

- *Substitutes*: The value assigned to a set is *lower* than the sum of the values assigned to its elements.

A standard example for substitutes would be a ticket to the theatre and another one to a football match for the same night.

In either case a simple auction mechanism that allocates one item at a time is problematic, even if we were to make the (unrealistic) assumption that agents will bid truthfully ...

Combinatorial Auctions

In a *combinatorial auction*, the auctioneer puts several goods on sale and the other agents submit bids for entire bundles of goods.

Given a set of bids, the *winner determination problem* (WDP) is the problem of deciding which of the bids to accept.

- The solution must be *feasible* (no good may be allocated to more than one agent).
- Ideally, it should also be *optimal* (in the sense of maximising revenue for the auctioneer).

Clearly finding a solution to the WDP can be tricky (just how “tricky” we’ll see later on in the course).

So besides the game-theoretical problem of stopping bidders from strategising, in combinatorial auctions we also face a challenging *algorithmic* problem. The same applies to any mechanism for collective decision making in combinatorial domains.

Communicating Bids

Suppose we are running a combinatorial auction with n goods. So there are $2^n - 1$ bundles that agents may want to bid for.

For interesting values of n , it is not possible to communicate your valuations to the auctioneer by simply stating your price for each and every bundle. \rightsquigarrow How do we best communicate/represent preferences in combinatorial domains?

For combinatorial auctions, this is the job of the *bidding language*.

Example: Bid for a small number of concrete bundles with the implicit understanding that you will honour any combination of bids that is feasible and pay the sum of the associated prices.

In general, *preference representation* is a central issue in computational social choice ...

Preference Representation Languages

- *Cognitive relevance*: How close is a given language to the way in which humans would express their preferences?
- *Elicitation*: How difficult is it to elicit the preferences of an agent so as to represent them in the chosen language?
- *Expressive power*: Can the chosen language encode all the preference structures we are interested in?
- *Succinctness*: How compact is the representation of (typical) preferences? Is one language more succinct than another?
- *Complexity*: What is the computational complexity of related decision problems, such as comparing two alternatives?

Centralised vs. Distributed Approaches

So far we have concentrated on auctions as mechanisms for solving resource allocation problems. But what if we cannot find someone who could act as the auctioneer?

- The associated tasks may be too hard *computationally* for the agent supposed to assume the role of auctioneer.
- There may be no agent that enjoys the *trust* of the others.
- The system infrastructure may be truly *distributed* and it may be unnatural to model the problem in a centralised manner.
- The goods may be owned by different agents to begin with (that is, we may have to take an *initial allocation* into account).
- Agents and goods may enter or leave the system *dynamically*.

Auctions are *centralised* mechanisms. All of the above are good reasons to consider *distributed* negotiation schemes as well . . .

Distributed Negotiation Schemes

In truly distributed approaches to resource allocation, allocations emerge as the result of a sequence of local negotiation steps.

However, such systems are more difficult to design and understand than auction-based mechanisms. Issues include:

- What types of deals do we allow for (*bilateral* vs. *multilateral* negotiation)?
- How do we design appropriate communication *protocols*? Note that this problem is relatively easy for auctions.
- To what degree is it possible to *predict* (or even control) the outcome of distributed negotiation processes? Put differently: What is the relationship between the negotiation strategies agents use *locally* and the allocations emerging *globally*?

Efficiency and Fairness

When assessing the quality of an allocation (or any other decision) we can distinguish two types of indicators of *social welfare*.

Aspects of *efficiency* (*not* in the computational sense) include:

- The chosen agreement should be such that there is no alternative agreement that would be better for some and not worse for any of the other agents (*Pareto optimality*).
- If preferences are quantitative, the sum of all payoffs should be as high as possible (*utilitarianism*).

Aspects of *fairness* include:

- The agent that is going to be worst off should be as well off as possible (*egalitarianism*).
- No agent should prefer to take the bundle allocated to one of its peers rather than keeping their own (*envy-freeness*).

Back to Voting

In the context of multiagent resource allocation, computational complexity presents a major challenge. But *high complexity* can also be a *good thing* . . .

We have seen before that common *voting* rules (such as the plurality rule) are susceptible to *manipulation*: if one voter knows how the others are going to vote, they may be able to influence the outcome in their own favour by submitting an insincere ballot.

Idea: If we can design a voting rule that makes it computationally intractable to compute such an insincere ballot, even if if we have all the information on the other voters available, then that may well be considered good enough a protection against manipulation.

Literature

There is no textbook or similar for COMSOC. I will recommend specific papers or book chapters in each lecture. For the general feeling, you may also want to browse through some of these:

- Y. Chevaleyre *et al.* *A Short Introduction to Computational Social Choice*. Proc. SOFSEM-2007, Springer-Verlag, 2007.
This has been a first attempt at giving an overview of the field. Mentions a lot of material, but it is not particularly suited as a tutorial.
- Y. Chevaleyre *et al.* Issues in Multiagent Resource Allocation. *Informatica*, 30:3–31, 2006.
A much more readable and comprehensive survey paper than the above, covering some of our main topics in detail.
- C.H. Papadimitriou. *Algorithms, Games, and the Internet*. Proc. STOC-2001, ACM Press, 2001.
Very nice for inspiration, albeit not exactly what we'll be focussing on.

Part II: Arrow's Theorem

Arrow's Impossibility Theorem

This is probably the most famous theorem in social choice theory. It was first proved by Kenneth J. Arrow in his 1951 PhD thesis. He later received the Nobel Prize in Economic Sciences in 1972.

The theorem shows that there can be no mechanism for aggregating individual preferences into a social preference that would simultaneously satisfy a small number of natural and seemingly innocent axioms.

The exposition of the theorem is taken from Barberà (1980); the proof closely follows Geanakoplos (2005).

K.J. Arrow. *Social Choice and Individual Values*. 2nd edition, Wiley, 1963.

S. Barberà (1980). Pivotal Voters: A New Proof of Arrow's Theorem. *Economics Letters*, 6(1):13–16, 1980.

J. Geanakoplos. Three Brief Proofs of Arrow's Impossibility Theorem. *Economic Theory*, 26(1):211–215, 2005.

Setting

- Finite set of *alternatives* A .
- Finite set of *individuals* $I = \{1, \dots, n\}$.
- A *preference ordering* is a strict linear order on A .
The set of all such preference orderings is denoted \mathcal{P} .
Each individual i has an *individual* preference ordering P_i ,
and we will try to find a *social* preference ordering P .
- A *preference profile* $\langle P_1, \dots, P_n \rangle \in \mathcal{P}^n$ consists of a preference ordering for each individual.
- A *social welfare function* (SWF) is a mapping from preference profiles to social preference orderings: it specifies what preferences society should adopt for any given situation.

Axioms

It seems reasonable to postulate that any SWF should satisfy the following list of axioms:

- **(P)** The SWF should satisfy the *Pareto condition*: if every individual prefers x over y , then so should society.
- $$(\forall \mathbf{P} \in \mathcal{P}^n)(\forall x, y \in A)[[(\forall i \in I)xP_i y] \rightarrow xPy]$$
- **(IIA)** The SWF should satisfy *independence of irrelevant alternatives*: social preference of x over y should not be affected if individuals change their preferences over other alternatives.
- $$(\forall \mathbf{P}, \mathbf{P}' \in \mathcal{P}^n)(\forall x, y \in A)[[(\forall i \in I)(xP_i y \leftrightarrow xP'_i y)] \rightarrow (xPy \leftrightarrow xP'y)]$$
- **(D)** The SWF should be *non-dictatorial*: no single individual should be able to impose a social preference ordering.

$$\neg(\exists i \in I)(\forall x, y \in A)(\forall \mathbf{P} \in \mathcal{P}^n)[xP_i y \rightarrow xPy]$$

The Result

Theorem 1 (Arrow, 1951) *If $|A| > 2$, then there exists no SWF that would simultaneously satisfy all of (P), (IIA) and (D).*

Observe that if there are just two alternatives ($|A| = 2$), then it is easy to find an SWF that satisfies all three axioms (at least for an odd number of individuals): simply let the alternative preferred by the *majority* of individuals also be the socially preferred alternative.

Now for the proof ...

Extremal Lemma

Assume (P) and (IIA) are satisfied. Let b be any alternative.

Claim: For any profile in which b is ranked either top or bottom by every individual, society must do the same.

Proof: Suppose otherwise; that is, suppose b is ranked either top or bottom by every individual, but not by society.

- (1) Then aPb and bPc for distinct alternatives a, b, c and the social preference ordering P .
- (2) By (IIA), this continues to hold if we move every c above a for every individual, as doing so does not affect the extremal b .
- (3) By transitivity of P , we get aPc .
- (4) But by (P), we get cPa . Contradiction. ✓

Existence of an Extremal Pivotal Individual

Fix some alternative b .

We call an individual *extremal pivotal* iff it can move b from the bottom to the top of the social preference ordering.

Claim: There exists an extremal pivotal individual.

Proof: Start with a profile where every individual puts b at the bottom. By (P), so does society.

Then let the individuals change their preferences one by one, moving b from the bottom to the top.

By the Extremal Lemma, there must be a point when the change in preference of a particular individual causes b to rise from the bottom to the top in the social ordering. ✓

Call the profile just before the switch in the social ordering occurred *Profile I*, and the one just after the switch *Profile II*.

Dictatorship: Case 1

Let i be the extremal pivotal individual (for alternative b).

The existence of i is guaranteed by our previous argument.

Claim: Individual i can dictate the social ordering with respect to any alternatives a, c different from b .

Proof: Suppose i wants to place a above c .

Let *Profile III* be like *Profile II*, except that i makes a its top choice: $aP_i bP_i c$. Now let all the others rearrange their relative rankings of a and c as they please.

Observe that in *Profile III* all relative rankings for a, b are as in *Profile I*. So by (IIA), the social rankings must coincide: aPb .

Also observe that in *Profile III* all relative rankings for b, c are as in *Profile II*. So by (IIA), the social rankings must coincide: bPc .

By transitivity, we get aPc . ✓

Dictatorship: Case 2

Let b and i be defined as before.

Claim: Individual i can also dictate the social ordering with respect to b and any other alternative a .

Proof: We can use a similar construction as before to show that for a given alternative c , there must be an individual j that can dictate the relative social ordering of a and b (both different from c).

But at least in *Profiles I* and *II*, i can dictate the relative social ranking of a and b . As there can be at most one dictator in any situation, we get $i = j$. ✓

So individual i will be a *dictator* for *any* two alternatives.

This contradicts (D), and Arrow's Theorem follows.

What next?

The main topics that we are going to cover in this course are:

- (Computational Issues in) Voting Theory
- Preference Representation in Combinatorial Domains
- Multiagent Resource Allocation and Fair Division
- Combinatorial Auctions and Mechanism Design

But as we are going to have to refer to basic concepts from game theory every now and then throughout the course, next week's class is going to be a short *introduction to game theory*.