Logic and Social Choice Theory

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

(Computational) Social Choice

Social Choice Theory studies group decision making: how should we aggregate individual preferences to obtain a "social preference"?

$$\triangle \succ_{\mathbf{1}} \bigcirc \succ_{\mathbf{1}} \square$$

$$\square \succ_{\mathbf{2}} \triangle \succ_{\mathbf{2}} \bigcirc$$

$$\bigcirc \succ_{\mathbf{3}} \square \succ_{\mathbf{3}} \triangle$$

?

Computational Social Choice is a fairly new research area that adds a computational component to such questions.

Talk Outline

- As an example for work in Social Choice Theory:
  - Arrow's Theorem + Proof
- Brief flavour of Computational Social Choice (COMSOC):
  - Possible COMSOC-style approaches to Arrow's Theorem
  - Some other research directions (focus on ILLC + on logic)
- Conclusion / Practical matters:
  - COMSOC as a research area, international activities
  - COMSOC at the ILLC

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.

Our 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* $I = \{1, \ldots, n\}$.
- A *preference ordering* is a strict linear order on A.
  The set of all such preference orderings is denoted by $\mathcal{P}$.
  Each individual i has an *individual* preference ordering P_i, and we will try to find a *social* preference ordering P.
- A *preference profile* $\boldsymbol{P} = \langle P_1, \ldots, P_n \rangle \in \mathcal{P}^n$ 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.
- Remark: We implicitly assume that *any* individual preference orderings are possible (*universal domain* assumption).

Axioms

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

- **(WP)** The SWF should satisfy the *weak Pareto condition* (aka. *unanimity*): if everyone prefers x over y, then so should society.

$$(\forall \boldsymbol{P} \in \mathcal{P}^n)(\forall x, y \in A)[[(\forall i \in I)xP_iy] \rightarrow xPy]$$

- **(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.

$$(\forall \boldsymbol{P}, \boldsymbol{P}' \in \mathcal{P}^n)(\forall x, y \in A)[[(\forall i \in I)(xP_iy \leftrightarrow xP'_iy)] \rightarrow (xPy \leftrightarrow xP'y)]$$

- **(ND)** The SWF should be *non-dictatorial:* no single individual should be able to impose a social preference ordering.

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

The Result

Theorem 1 (Arrow, 1951) *For three or more alternatives, there exists no SWF that satisfies all of* (WP), (IIA) *and* (ND)*.*

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 (WP) 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, applied to (1), we get aPc.

(4) But by (WP), applied to (2), we get cPa. Contradiction. ✓

Existence of an Extremal Pivotal Individual

Fix some alternative b. We call an individual *extremal pivotal* if it can move b from the bottom to the top of the social preference ordering (for some particular profile).

Claim: There exists an extremal pivotal individual i.

Proof: Start with a profile where every individual puts b at the bottom. By (WP), 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 and (WP), 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 (1) i makes a its top choice (that is, aP_ibP_ic), and (2) all the others have rearranged 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. By (IIA), this continues to hold if others change their relative ranking of alternatives other than a, c. ✓

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 (ND), and Arrow's Theorem follows.

Arrow's Theorem and COMSOC

From a COMSOC perspective, Arrow's model of preference aggregation and Arrow's Theorem can raise questions such as:

- What is the right *logic* for formalising this?
- Can we prove the theorem automatically, or can we at least automatically check a known proof? (*automated reasoning*)
- What is the *computational complexity* of relevant questions, e.g., deciding whether a given profile satisfies a given domain condition that is known to avoid the impossibility?
- Are *other preference structures* / *other axioms* maybe more appropriate, e.g., for applications in AI and MAS?
- *Communication complexity* of preference aggregation?
- What about preference aggregation for alternatives with an internal structure? (*combinatorial domains*)

Full Formalisation of Arrow's Theorem

Logic has long been used to *formally specify* computer systems, facilitating formal or even *automatic verification* of various properties. Can we apply this methodology also to *social choice* mechanisms?

Tang and Lin (2009) show that the *"base case"* of Arrow's Theorem with 2 agents and 3 alternatives can be fully modelled in *propositional logic:*

- Automated theorem provers can verify $\textsc{Arrow}(2,3)$ to be correct in < 1 second — that's $(3!)^{3! \times 3!} \approx 10^{28}$ SWFs to check
- Opens up opportunities for quick sanity checks of hypotheses regarding new possibility and impossibility theorems.

Our work using *first-order logic* tries to go beyond such base cases.

P. Tang and F. Lin. Computer-aided Proofs of Arrow's and other Impossibility Theorems. *Artificial Intelligence*, 173(11):1041–1053, 2009

U. Grandi and U. Endriss. *First-Order Logic Formalisation of Arrow's Theorem*. Proc. 2nd Workshop on Logic, Rationality and Interaction (LORI-2009).

Social Choice in Combinatorial Domains

Many social choice problems have a *combinatorial structure*, e.g., electing a *committee* of k members from amongst n candidates.

- no. of 3-member committees from 10 candidates: $\binom{10}{3} = 120$ (i.e., $120! \approx 6.7 \times 10^{198}$ possible rankings)
- So, we need good (*compact*) languages!

Simple logic-based languages can be used to model preferences:

- Example: $\{(p, 8), (q, 6), (p \wedge r, 7)\} \rightsquigarrow$ value of $\{p, \overline{q}, r\}$ is 15

Relevant research questions:

- Expressiveness? / Succinctness? / Complexity?

J. Uckelman, Y. Chevaleyre, U. Endriss, and J. Lang. Representing Utility Functions via Weighted Goals. *Mathem. Logic Quarterly*, 55(4):341–361, 2009.

Judgment Aggregation

Preferences are not the only structures we may wish to aggregate. JA studies the aggregation of judgments on related propositions.

	p	$p \rightarrow q$	q
Agent 1:	Yes	Yes	Yes
Agent 2:	No	Yes	No
Agent 3:	Yes	No	No
Majority:	Yes	Yes	No

In ongoing work we investigate the *computational complexity* of problems arising in JA, e.g., whether a given "agenda" is "safe".

C. List and C. Puppe. Judgment Aggregation: A Survey. *Handbook of Rational and Social Choice*. Oxford University Press, 2009.

U. Endriss, U. Grandi, and D. Porello. *Complexity of Judgment Aggregation*. Working Paper, ILLC, University of Amsterdam, 2009.

Last Slide

- Computational social choice studies problems of collective decision making from a "computational" point of view, i.e., using the kind of tools you acquire during the MoL.
- COMSOC website: http://www.illc.uva.nl/~ulle/COMSOC/ (workshop series, other events, mailing list, etc.)
- This is a good time to get into the field.
- COMSOC at the ILLC:
  - Course on Computational Social Choice (September 2010)
  - Computational Social Choice Seminar
- Papers, including several survey papers, are available from my website. If you have time for just one of them, maybe try the *AI Magazine* article cited below.

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine*, Winter 2008.