Logic and Social Choice

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

[Guest Lecture for "Logic, Language and Computation"
Master of Logic, 4 October 2021]

Social Choice Theory

Social choice theory is about methods for *collective decision making*, such as *political* decision making by groups of *economic* agents.

Its methodology ranges from the *philosophical* to the *mathematical* to the *algorithmic* (the latter in *computational social choice*).

Its findings are relevant to a variety of *applications*, such as these:

- How to choose a president given people's preferences?
- How to combine the website rankings of multiple search engines?
- How to fairly divide resources between several stake-holders?
- How to assign student doctors to hospitals?
- How to aggregate the views of different judges in a court case?

The most widely studied scenario is that of *voting* to select a single alternative, with n voters each ranking m alternatives.

Outline

- Examples for voting rules
- Example for a classical result: Gibbard-Satterthwaite Theorem
- Examples for recent work on logic and SCT in Amsterdam

The Plurality Rule

Under the *plurality rule*, the alternative ranked first most often wins:

Voter 1: $a \succ b \succ d \succ c$
Voter 2: $a \succ b \succ d \succ c$
Voter 3: $b \succ c \succ d \succ a$
Voter 4: $c \succ b \succ d \succ a$

The most widely used rule in practice (by far!). *But is it any good?*

The Borda Rule

Under the *Borda rule*, each voter gives $m-1$ points to the alternative she ranks first, $m-2$ to the alternative she ranks second, and so forth:

Voter 1:	$a \succ c \succ b$	$a: 2+2+1+1+0=6$
Voter 2:	$a \succ c \succ b$	$b: 0+0+2+2+1=5$
Voter 3:	$b \succ a \succ c$	$c: 1+1+0+0+2=4$
Voter 4:	$b \succ a \succ c$	
Voter 5:	$c \succ b \succ a$	

A clear advantage over the plurality rule is that we use much more of the information present in the profile to come to a decision.

But there still is a problem: *What if b challenges the winner a?*

The Condorcet Rule

Under the *Condorcet rule*, we run one-to-one majority contests between all pairs of alternatives and elect the one that performs best.

Nice idea. But: *Do you see the problem with this "definition"?*

Voter 1: $a \succ b \succ c$
Voter 2: $b \succ c \succ a$
Voter 3: $c \succ a \succ b$

The Axiomatic Method

So how do you decide what is the right voting rule to use?

The classical approach is to use the *axiomatic method:*

- identify good "axioms": normatively appealing high-level properties
- give mathematically rigorous definitions of your axioms
- explore the logical consequences of your definitions

The Model

The *voters* from a finite set $N = \{1, \ldots, n\}$, with $n \geqslant 2$, all rank the *alternatives* in a set A by supplying a *strict linear order* in $\mathcal{L}(A)$.

We are interested in *voting rules* of this form:

$$F : \mathcal{L}(A)^n \to A$$

Given a *profile* $(R_1, \ldots, R_n) \in \mathcal{L}(A)^n$ of rankings, one for each voter, we want our rule to return a single winning alternative ($\leadsto$ no *ties*).

Example

The voters in $N = \{1, \ldots, 5\}$ need to choose from $A = \{a, b, c\}$.
Suppose you know their true preferences:

Voter 1: $a \succ b \succ c$
Voter 2: $a \succ b \succ c$
Voter 3: $b \succ c \succ a$
Voter 4: $b \succ c \succ a$
Voter 5: $c \succ b \succ a$

Suppose we use the *plurality rule* (with "alphabetical" tie-breaking).

Exercise: *What is your advice to voter 5?*

Strategyproofness

Ideally, we would not want voters to have to lie about their preferences.

We encode this desideratum as an *axiom:*

> F is *strategyproof* (or *immune to manipulation*) if for no voter $i \in N$ there exist a profile $\boldsymbol{R} = (R_1, \ldots, R_n) \in \mathcal{L}(A)^n$ and an $R_i' \in \mathcal{L}(A)$ s.t. R_i ranks $F(R_i', \boldsymbol{R}_{-i})$ above $F(\boldsymbol{R})$.

Here R_i is agent i's *true preference* and R_i' her *untruthful ballot*.

Exercise: *Plurality is not strategyproof. Can you think of a rule that is?*

Notation: $(R_i', \boldsymbol{R}_{-i})$ is the profile obtained by replacing R_i in $\boldsymbol{R}$ by R_i'.

The Gibbard-Satterthwaite Theorem

Two more properties of voting rules F:

- F is *surjective* if for every $x \in A$ *some* profile $\boldsymbol{R}$ yields $F(\boldsymbol{R}) = x$.
- F is a *dictatorship* if there exists a voter $i \in N$ (the dictator) s.t. $F(\boldsymbol{R})$ is the top-ranked alternative in R_i for every profile $\boldsymbol{R}$.

Bad news:

Theorem 1 (Gibbard-Satterthwaite) *Any voting rule for 3 or more alternatives that is surjective and strategyproof must be a dictatorship.*

This was formulated and proved independently by Gibbard (1973) and Satterthwaite (1975), after having been "in the air" for some time.

A. Gibbard. Manipulation of Voting Schemes: A General Result. *Econometrica*, 41(4):587–601, 1973.

M.A. Satterthwaite. Strategy-proofness and Arrow's Conditions. *Journal of Economic Theory*, 10:187–217, 1975.

Remarks

The G-S Theorem says that for *$\geqslant$ 3 alternatives*, any voting rule F that is *surjective* and *strategyproof* must be *dictatorial*.

- make sure you really understand what "dictatorial" means
- *surprising* result + not applicable in case of *two* alternatives
- opposite also true: *dictatorial* $\Rightarrow$ *strategyproof*
- *random* rules don't count (but might be "strategyproof").

We will now prove the theorem under two additional assumptions:

- F is *neutral*, i.e., alternatives are treated symmetrically.
  [Note: neutrality $\Rightarrow$ surjectivity; so we won't make use of surjectivity.]
- There are *exactly 3 alternatives*.

For a full proof, using a similar approach, see, e.g.:

U. Endriss. Logic and Social Choice Theory. In A. Gupta and J. van Benthem (eds.), *Logic and Philosophy Today*, College Publications, 2011.

Proof (1): Independence and Blocking Coalitions

Notation: $N^{\boldsymbol{R}}_{x \succ y}$ is the set of voters who rank x above y in profile $\boldsymbol{R}$.

Claim: $F(\boldsymbol{R}) = x \;\; + \;\; N^{\boldsymbol{R}}_{x \succ y} = N^{\boldsymbol{R'}}_{x \succ y} \;\Rightarrow\; F(\boldsymbol{R'}) \neq y$ [independence]

Proof: From *strategyproofness*, by contradiction. Assume $F(\boldsymbol{R'}) = y$. Moving from $\boldsymbol{R}$ to $\boldsymbol{R'}$, there must be a *first* voter to affect the winner. So w.l.o.g., assume $\boldsymbol{R}$ and $\boldsymbol{R'}$ differ only w.r.t. voter i. Two cases:

- $i \in N^{\boldsymbol{R}}_{x \succ y}$: Suppose i's true preferences are as in profile $\boldsymbol{R'}$ (i.e., i prefers x to y). Then i has an incentive to vote as in $\boldsymbol{R}$. ✓
- $i \notin N^{\boldsymbol{R}}_{x \succ y}$: Suppose i's true preferences are as in profile $\boldsymbol{R}$ (i.e., i prefers y to x). Then i has an incentive to vote as in $\boldsymbol{R'}$. ✓

Some more terminology:

Call $C \subseteq N$ a *blocking coalition* for (x, y) if $C = N^{\boldsymbol{R}}_{x \succ y} \Rightarrow F(\boldsymbol{R}) \neq y$.

Thus: If $F(\boldsymbol{R}) = x$, then $C := N^{\boldsymbol{R}}_{x \succ y}$ is blocking for (x, y) [for any y].

Proof (2): Ultrafilters

From *neutrality:* all (x, y) must have *the same* blocking coalitions.

For any $C \subseteq N$, C *or* $\overline{C} := N \setminus C$ must be blocking.

> Proof: Assume C is not blocking; i.e., C is not blocking for (x, y). Then there exists a profile $\boldsymbol{R}$ with $N^{\boldsymbol{R}}_{x \succ y} = C$ but $F(\boldsymbol{R}) = y$. But we also have $N^{\boldsymbol{R}}_{y \succ x} = \overline{C}$. Hence, $\overline{C}$ is blocking for (y, x).

If C_1 and C_2 are blocking, then so is $C_1 \cap C_2$. [now we'll use $|A| = 3$]

> Proof: Consider a profile $\boldsymbol{R}$ with $C_1 = N^{\boldsymbol{R}}_{x \succ y}$, $C_2 = N^{\boldsymbol{R}}_{y \succ z}$, and $C_1 \cap C_2 = N^{\boldsymbol{R}}_{x \succ z}$. As C_1 is blocking, y cannot win. As C_2 is blocking, z cannot win. So x wins and $C_1 \cap C_2$ must be blocking.

The *empty coalition* is *not* blocking.

> Proof: Omitted (but not at all surprising).

Above properties (+ finiteness of N) imply that there's a *singleton* $\{i\}$ that is blocking. But that just means that i is a *dictator*! ✓

Logic and Social Choice Theory

The axiomatic method itself is an (informal) use of logic in SCT.

Next: Several more explicit uses of logic in research on SCT . . .

Formal Modelling of Social Choice Scenarios

Maybe the most obvious application of logic concerns the formal modelling of social choice scenarios. Examples for results obtained:

- We can model results such as the one just discussed in classical FOL (i.e., w/o resorting to second-order constructs)—*except* for the finiteness of the electorate (Grandi and Endriss, 2013).
- For a fixed number of voters, we can cast proofs such as the one just discussed as formal derivations in natural deduction in a tailor-made modal logic for SCT (Ciná and Endriss, 2016).

U. Grandi and U. Endriss. First-Order Logic Formalisation of Impossibility Theorems in Preference Aggregation. *Journal of Philosophical Logic*, 2013.

G. Ciná and U. Endriss. Proving Classical Theorems of Social Choice Theory in Modal Logic. *Journal of Autonomous Agents and Multiagent Systems*, 2016.

Automated Reasoning for Social Choice Theory

Logic-based automated reasoning has long been used to obtain results in a range of areas of mathematics. Can we do so also here? *Yes!*

For example, we have been able to use a *SAT solver* to help us find (and prove a central part of) a new *impossibility theorem* regarding voting rules to *elect committees* (Kluiving et al., 2020).

B. Kluiving, A. de Vries, P. Vrijbergen, A. Boixel, and U. Endriss. Analysing Irresolute Multiwinner Voting Rules with Approval Ballots via SAT Solving. ECAI-2020.

Judgment Aggregation

Suppose three robots are in charge of climate control for a building. They need to make judgments on p (*the temperature is below 17° C*), on q (*we should switch on the heating*), and on the "policy" $p \rightarrow q$.

	p	$p \rightarrow q$	q
Robot 1:	Yes	Yes	Yes
Robot 2:	No	Yes	No
Robot 3:	Yes	No	No

Exercise: *What is the correct collective decision on q?*

Some of our recent work has focused on *strategyproofness* in JA.

Z. Terzopoulou and U. Endriss. Strategyproof Judgment Aggregation under Partial Information. *Social Choice and Welfare*, 2019.

S. Botan and U. Endriss. Majority-Strategyproofness in Judgment Aggregation. AAMAS-2020.

Last Slide

We have seen:

- the Gibbard-Satterthwaite Theorem as an example for classical work in SCT: no good voting rule can avoid strategic behaviour
- examples for modern research directions in computational social choice (emphasising logic in particular)

For possible entry points into the field, see the papers cited below, my MoL course on COMSOC, or this website:

`http://research.illc.uva.nl/COMSOC/`

U. Endriss. Logic and Social Choice Theory. In A. Gupta and J. van Benthem (eds.), *Logic and Philosophy Today*. College Publications, 2011.

F. Brandt, V. Conitzer, and U. Endriss. Computational Social Choice. In G. Weiss (ed.), *Multiagent Systems*. MIT Press, 2013.