Topics in Computational Social Choice 2026

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

Plan for Today

- Discussion of first exercise sheet
- Introduction to voting theory and Arrow's Theorem
- Paper assignment for next week's presentations
- (optional) COMSOC Video Seminar in the afternoon

Three Voting Rules

Suppose n *voters* choose from a set of m *alternatives* by stating their preferences in the form of *linear orders* over the alternatives.

Here are three *voting rules* (there are many more):

- *Plurality:* elect the alternative ranked first most often (i.e., each voter assigns 1 point to an alternative of her choice, and the alternative receiving the most points wins)
- *Plurality with runoff:* run a plurality election and retain the two front-runners; then run a majority contest between them
- *Borda:* each voter gives $m-1$ points to the alternative she ranks first, $m-2$ to the alternative she ranks second, etc.; and the alternative with the most points wins

Exercise: *Do you know real-world elections where these rules are used?*

Example: Choosing a Beverage for Lunch

Consider this election, with nine *voters* having to choose from three *alternatives* (namely what beverage to order for a common lunch):

2 Germans:	Beer	$\succ$	Wine	$\succ$	Milk
3 French people:	Wine	$\succ$	Beer	$\succ$	Milk
4 Dutch people:	Milk	$\succ$	Beer	$\succ$	Wine

Recall that we saw three different voting rules:

- Plurality
- Plurality with runoff
- Borda

Exercise: *For each of the rules, which beverage wins the election?*

Even More Voting Rules

There are many more voting rules. Some more examples:

- *Copeland:* Score alternatives in terms of the number of pairwise majority contests they win and lose.
- *Kemeny:* Compute the weighted majority graph on the set of alternatives and then find the “closest” linear order.
- *IRV/STV:* Repeatedly eliminate the plurality loser.

Positional Scoring Rules

We can generalise the idea underlying the Borda rule as follows:

A *positional scoring rule* (PSR) is defined by a so-called *scoring vector* $\boldsymbol{s} = (s_1, \ldots, s_m) \in \mathbb{R}^m$ with $s_1 \geqslant s_2 \geqslant \cdots \geqslant s_m$ and $s_1 > s_m$.

Each voter submits a ranking of the m alternatives. Each alternative receives s_i points for every voter putting it at the ith position.

The alternative(s) with the highest score (sum of points) win(s).

Examples:

- *Borda rule* = PSR with scoring vector $(m-1, m-2, \ldots, 0)$
- *Plurality rule* = PSR with scoring vector $(1, 0, \ldots, 0)$
- *Veto rule* = PSR with scoring vector $(0, \ldots, 0, -1)$
- For any $k < m$, *k-approval* = PSR with $(\underbrace{1, \ldots, 1}_{k}, 0, \ldots, 0)$

Exercise: *Name the rule induced by $\boldsymbol{s} = (9, 7, 5)$! General idea?*

Condorcet Extensions

An alternative that beats every other alternative in pairwise majority contests is called a *Condorcet winner*. Sometimes there is no CW:

$$a \succ b \succ c$$
$$b \succ c \succ a$$
$$c \succ a \succ b$$

This is the famous *Condorcet Paradox*.

The *Condorcet Principle* says that, if it exists, only the CW should win. Voting rules that satisfy this principle are called *Condorcet extensions*.

Exercise: *Show that Copeland and Kemeny are Condorcet extensions.*

Positional Scoring Rules and the Condorcet Principle

Consider this example with three alternatives and seven voters:

3 voters:	$a \succ b \succ c$
2 voters:	$b \succ c \succ a$
1 voter:	$b \succ a \succ c$
1 voter:	$c \succ a \succ b$

So a is the *Condorcet winner:* a beats both b and c (with 4 out of 7). But any *positional scoring rule* makes b win (because $s_1 \geqslant s_2 \geqslant s_3$):

$$\begin{aligned} a&: \quad 3 \cdot s_1 + 2 \cdot s_2 + 2 \cdot s_3 \\ b&: \quad 3 \cdot s_1 + 3 \cdot s_2 + 1 \cdot s_3 \\ c&: \quad 1 \cdot s_1 + 2 \cdot s_2 + 4 \cdot s_3 \end{aligned}$$

Thus, *no positional scoring rule* for three (or more) alternatives can possibly satisfy the *Condorcet Principle*.

Fishburn's Classification

Can classify voting rules on the basis of the *information* they require. The best known such classification is due to Fishburn (1977):

- *C1:* Winners can be computed from the *majority graph* alone.
  Examples: Copeland, Slater
- *C2:* Winners can be computed from the *weighted majority graph* (but not from the majority graph alone).
  Examples: Kemeny, Ranked Pairs, Borda
- *C3:* All other voting rules.
  Examples: Young, Dodgson, IRV/STV

Remark: Fishburn originally intended this for Condorcet extensions only, but the concept also applies to all other voting rules.

P.C. Fishburn. Condorcet Social Choice Functions. *SIAM Journal on Applied Mathematics*, 1977.

The Axiomatic Method

So many voting rules! How do you choose?

One approach is to use the *axiomatic method* to identify voting rules of *normative* appeal. Next:

- Formal model voting
- Examples for axioms
- Example for a characterisation result: May's Theorem
- Example for an impossibility result: Arrow's Theorem

The Model

Fix a finite set $A = \{a, b, c, \ldots\}$ of *alternatives*, with $|A| = m \geqslant 2$.

Let $\mathcal{L}(A)$ denote the set of all strict linear orders R on A. We use elements of $\mathcal{L}(A)$ to model (true) *preferences* and (declared) *ballots*.

Each member i of a finite set $N = \{1, \ldots, n\}$ of *voters* supplies us with a ballot R_i, giving rise to a *profile* $\boldsymbol{R} = (R_1, \ldots, R_n) \in \mathcal{L}(A)^n$.

A *voting rule* (or *social choice function*) for N and A selects (ideally) one or (in case of a tie) more winners for every such profile:

$$F : \mathcal{L}(A)^n \to 2^A \setminus \{\emptyset\}$$

If $|F(\boldsymbol{R})| = 1$ for all profiles $\boldsymbol{R}$, then F is called *resolute*.

Most natural voting rules are *irresolute* and have to be paired with a *tie-breaking rule* to always select a unique election winner.

Examples: random tie-breaking, lexicographic tie-breaking

Axioms: Anonymity and Neutrality

Two basic fairness requirements for a voting rule F:

- F is *anonymous* if $F(R_1, \ldots, R_n) = F(R_{\pi(1)}, \ldots, R_{\pi(n)})$ for any profile $(R_1, \ldots, R_n)$ and any permutation $\pi : N \to N$.
- F is *neutral* if $F(\pi(\boldsymbol{R})) = \pi(F(\boldsymbol{R}))$ for any profile $\boldsymbol{R}$ and any permutation $\pi : A \to A$ (with π extended to profiles and sets of alternatives in the natural manner).

In other words:

- Anonymity is symmetry w.r.t. voters.
- Neutrality is symmetry w.r.t. alternatives.

Consequences of Axioms

For this slide only, let us restrict attention to voting rules for scenarios with just *two voters* ($n = 2$) and *two alternatives* ($m = 2$).

Exercise: *Show that there exists no resolute voting rule that is 'fair' in the sense of being both anonymous and neutral.*

Exercise: *But there still are a couple of irresolute voting rules that are both anonymous and neutral. Give some examples!*

Axiom: Positive Responsiveness

Notation: Write $N^{\boldsymbol{R}}_{x \succ y} = \{i \in N \mid (x,y) \in R_i\}$ for the set of voters who rank alternative x above alternative y in profile $\boldsymbol{R}$.

A (not necessarily resolute) voting rule satisfies *positive responsiveness* if, whenever some voter raises a (possibly tied) winner $x^\star$ in her ballot, then $x^\star$ will become the *unique* winner. Formally:

> F is *positively responsive* if $x^\star \in F(\boldsymbol{R})$ implies $\{x^\star\} = F(\boldsymbol{R'})$ for any alternative $x^\star$ and any two *distinct* profiles $\boldsymbol{R}$ and $\boldsymbol{R'}$ s.t. $N^{\boldsymbol{R}}_{x^\star \succ y} \subseteq N^{\boldsymbol{R'}}_{x^\star \succ y}$ and $N^{\boldsymbol{R}}_{y \succ z} = N^{\boldsymbol{R'}}_{y \succ z}$ for all $y, z \in A \setminus \{x^\star\}$.

Thus, this is a *monotonicity* requirement (there are others as well).

May's Theorem

When there are only *two alternatives*, then all the voting rules we have seen coincide with what is known as the *simple majority rule*.

Exercise: *Formulate a simple definition of the simple majority rule!*

Good news:

May's Theorem: *A voting rule for two alternatives satisfies the axioms of anonymity, neutrality, and positive responsiveness if and only if that rule is the simple majority rule.*

This provides a good justification for using this rule (arguing in favour of majority directly is harder than arguing for anonymity etc.).

K.O. May. A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decisions. *Econometrica*, 1952.

Proof Sketch

Clearly, the simple majority rule satisfies all three properties. ✓

Now for the other direction:

Assume the number of voters is *odd* $\rightsquigarrow$ no ties. (other case: similar)

There are two possible ballots: $a \succ b$ and $b \succ a$.
Anonymity $\rightsquigarrow$ only *number of ballots* of each type matters.

Consider all possible profiles $\boldsymbol{R}$. Distinguish two cases:

- Whenever $|N^{\boldsymbol{R}}_{a \succ b}| = |N^{\boldsymbol{R}}_{b \succ a}| + 1$, then only a wins.
  By *PR*, a wins whenever $|N^{\boldsymbol{R}}_{a \succ b}| > |N^{\boldsymbol{R}}_{b \succ a}|$. By *neutrality*, b wins otherwise. But this is just what the simple majority rule does. ✓

- There exists a profile $\boldsymbol{R}$ with $|N^{\boldsymbol{R}}_{a \succ b}| = |N^{\boldsymbol{R}}_{b \succ a}| + 1$, yet b wins.
  Suppose one a-voter switches to b, yielding $\boldsymbol{R}'$. By *PR*, now only b wins. But now $|N^{\boldsymbol{R}'}_{b \succ a}| = |N^{\boldsymbol{R}'}_{a \succ b}| + 1$, which is symmetric to the earlier situation, so by *neutrality* a should win. Contradiction. ✓

Resolute Social Choice Functions

For the remainder of today, we focus on *resolute SCFs:*

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

The next result we are going to see, Arrow's Theorem, originally got formulated for so-called *social welfare functions* instead:

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

This change in framework does not affect the essence of the result.

Axiom: The Pareto Principle

A resolute voting rule $F : \mathcal{L}(A)^n \to A$ is called (weakly) *Paretian* if, whenever all voters rank alternative x above y, then y cannot win:

$$N^{\boldsymbol{R}}_{x \succ y} = N \text{ implies } y \neq F(\boldsymbol{R})$$

Axiom: Independence of Irrelevant Alternatives

If alternative x wins and y does not, then x is *socially preferred* to y.
If both x and y lose, then we cannot say.

Whether x is socially preferred to y should *depend* only on the relative rankings of x and y in the profile (not on other, irrelevant, alternatives).

These considerations motivate our next axiom:

> F is called *independent* if, for any two profiles $\boldsymbol{R}, \boldsymbol{R}' \in \mathcal{L}(A)^n$ and any two distinct alternatives $x, y \in A$, it is the case that $N^{\boldsymbol{R}}_{x \succ y} = N^{\boldsymbol{R}'}_{x \succ y}$ and $F(\boldsymbol{R}) = x$ imply $F(\boldsymbol{R}') \neq y$.

Thus, if x prevents y from winning in $\boldsymbol{R}$ and the relative rankings of x and y remain the same, then x also prevents y from winning in $\boldsymbol{R}'$.

Arrow's Impossibility Theorem

A resolute SCF F is a *dictatorship* if there exists an $i \in N$ such that $F(\boldsymbol{R}) = \text{top}(R_i)$ for every profile $\boldsymbol{R}$. Voter i is the dictator.

The seminal result in SCT, here adapted from SWFs to SCFs:

Arrow's Theorem: *Any resolute SCF for $m \geqslant 3$ alternatives that is Paretian and independent must be a dictatorship.*

Exercise: *Redo Exercise 3 from the first exercise sheet!*

Remarks:

- Common misunderstanding: dictatorship $\neq$ 'local dictatorship'
- Impossibility reading: independence + Pareto + nondictatoriality
- Characterisation reading: dictatorship = independence + Pareto

K.J. Arrow. *Social Choice and Individual Values.* John Wiley and Sons, 2nd edition, 1963. First edition published in 1951.

Proof Plan

For full details, consult my review paper, which includes proofs both for SWFs and SCFs (the latter within the proof for the *M-S Thm*).

Let F be a SCF for $\geqslant 3$ alternatives that is Paretian and independent.

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

We proceed as follows:

- *Pareto* condition $=$ N is decisive for all pairs of alternatives
- C with $|C| \geqslant 2$ *decisive* for all pairs $\Rightarrow$ some $C' \subset C$ as well
- By induction: there's a decisive coalition of size 1 (= *dictator*).

Remark: Observe that this only works for finite sets of voters. (*Why?*)

The step in the middle of the list is known as the *Contraction Lemma*. To prove it, we first require another lemma . . .

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

Contagion Lemma

Recall: $C \subseteq N$ *decisive* for (x, y) if $C \subseteq N^{\boldsymbol{R}}_{x \succ y} \;\Rightarrow\; y \neq F(\boldsymbol{R})$

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

Claim: C weakly decisive for $(x, y) \;\Rightarrow\; C$ decisive for *all* pairs (x', y').

Proof: Suppose x, y, x', y' are all distinct (other cases: similar).

Consider a profile where individuals express these preferences:

- Members of C: $x' \succ x \succ y \succ y'$
- Others: $x' \succ x$, $y \succ y'$, and $y \succ x$ (note: x'-vs.-y' not specified)
- All rank x, y, x', y' above all other alternatives.

From C being weakly decisive for (x, y): y must lose.
From Pareto: x must lose (to x') and y' must lose (to y).
Thus, x' must win (and y' must lose). By independence, y' will still lose when everyone changes their non-x'-vs.-y' rankings.

Thus, for *every* profile $\boldsymbol{R}$ with $C \subseteq N^{\boldsymbol{R}}_{x' \succ y'}$, we get $y' \neq F(\boldsymbol{R})$. ✓

Contraction Lemma

Claim: If $C \subseteq N$ with $|C| \geqslant 2$ is a coalition that is decisive on all pairs of alternatives, then so is some nonempty coalition $C' \subset C$.

Proof: Take any nonempty C_1, C_2 with $C = C_1 \cup C_2$ and $C_1 \cap C_2 = \emptyset$.

Recall that there are $\geqslant 3$ alternatives. Consider this profile:

- Members of C_1: $x \succ y \succ z \succ$ *rest*
- Members of C_2: $y \succ z \succ x \succ$ *rest*
- Others: $z \succ x \succ y \succ$ *rest*

As $C = C_1 \cup C_2$ is decisive, z cannot win (it loses to y). Two cases:

(1) The winner is x: Exactly C_1 ranks $x \succ z \Rightarrow$ By independence, in any profile where exactly C_1 ranks $x \succ z$, z will lose (to x) $\Rightarrow C_1$ is weakly decisive on (x, z). So by Contagion Lemma: C_1 is decisive on all pairs.

(2) The winner is y, i.e., x loses (to y). Exactly C_2 ranks $y \succ x \Rightarrow \cdots \Rightarrow$ C_2 is decisive on all pairs.

Hence, one of C_1 and C_2 will always be decisive. ✓

Summary and Outlook

This has been an introduction to *voting theory*, with an emphasis on the *axiomatic method*, including two *seminal results* (May, Arrow).

Forthcoming lectures:

- Thursday: strategic manipulation in voting (+ other topics)
- Friday: introduction to fair division (FACT lecture)
- Friday: voting with approval ballots
- Monday: experiments in COMSOC