Computational Social Choice 2026

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

http://www.illc.uva.nl/~ulle/teaching/comsoc/2026/

Plan for Today

We first review a large number of concrete examples for *voting rules* proposed in the literature and used in practice.

To put some order into this large space of rules, we then shall:

- look into approaches to *classifying* voting rules
- review some *axioms* to differentiate between voting rules

For full details see Zwicker (2016).

W.S. Zwicker. Introduction to the Theory of Voting. In F. Brandt et al. (eds.), *Handbook of Computational Social Choice*. CUP, 2016.

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

Examples for Voting Rules

Borda | Plurality | Antiplurality | k-Approval | (Approval Voting)
Instant Runoff | Plurality with Runoff | Black
Cup | Copeland | Slater | Kemeny | Dodgson | Young
Ranked Pairs | Split Cycle | Simpson | Bucklin
(Range Voting) | (Cumulative Voting) | (Majority Judgment)

Exercise: *What is it that the rules in brackets have in common?*

Classifying Voting Rules

How can we put some order into this zoo of voting rules? Attempts:

- important family of *positional scoring rules* (operational definition)
- important family of *Condorcet extensions* (axiomatic perspective)
- classifying rules in terms of the *information* they require

Remark: Some of the rules we saw do not fit into our formal model (different ballot format), so they also do not fit these classifications.

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 $(1, 0, \ldots, 0)$
- *Antiplurality* (*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.*

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, Instant Runoff

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.

Computational Complexity

For some voting rules computing outcomes is computationally easy (in P) and for some it is intractable (NP-hard or worse).

Exercise: *Which of Borda, Copeland, Kemeny, Ranked Pairs look hard?*

M. Brill and F. Fischer. The Price of Neutrality for the Ranked Pairs Method. AAAI-2012.

The Axiomatic Method

Can also classify voting rules on the basis of the axioms they satisfy.

Examples for some axioms we might want to see satisfied:

- *Participation Principle:* It should be in the best interest of voters to participate; voting truthfully should be no worse than abstaining.
- *Pareto Principle:* There should be no alternative that every voter strictly prefers to the alternative selected by the voting rule.
- *Condorcet Principle:* If there is an alternative that is preferred to every other alternative by a majority of voters, then it should win.

Satisfying these seemingly basic requirements is harder than it looks . . .

The No-Show Paradox

Under *plurality with runoff* (and under *instant runoff*), it can be better to abstain than to participate and vote for your favourite alternative!

$$\begin{aligned} &25 \text{ voters:} \quad a \succ b \succ c \\ &46 \text{ voters:} \quad c \succ a \succ b \\ &24 \text{ voters:} \quad b \succ c \succ a \end{aligned}$$

Given these voter preferences, b gets eliminated in the first round, and c beats a 70:25 in the runoff.

Now suppose two voters from the first group abstain:

$$\begin{aligned} &23 \text{ voters:} \quad a \succ b \succ c \\ &46 \text{ voters:} \quad c \succ a \succ b \\ &24 \text{ voters:} \quad b \succ c \succ a \end{aligned}$$

Now a gets eliminated, and b beats c 47:46 in the runoff.

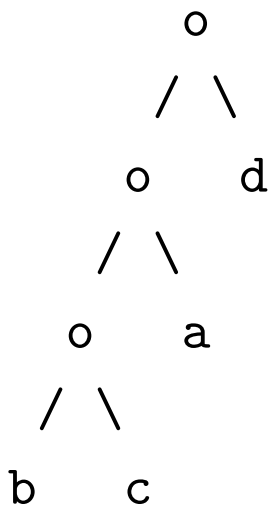
P.C. Fishburn and S.J Brams. Paradoxes of Preferential Voting. *Mathematics Magazine*, 1983.

Cup Rules and the Pareto Principle

The (weak) *Pareto Principle* requires that we should never elect an alternative that is strictly dominated in *every* voter's ballot.

In a *cup rule* the alternatives label the leaves of a *binary tree*, with an alternative progressing if it beats its sibling in a *majority contest*.

Cup rules do not always guarantee Pareto efficiency!



Consider this profile with three voters:

Ann: $\mathtt{a} \succ \mathtt{b} \succ \mathtt{c} \succ \mathtt{d}$

Bob: $\mathtt{b} \succ \mathtt{c} \succ \mathtt{d} \succ \mathtt{a}$

Cindy: $\mathtt{c} \succ \mathtt{d} \succ \mathtt{a} \succ \mathtt{b}$

d *wins!* (despite being dominated by c)

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 $s_1 \geqslant s_2 \geqslant s_3$):

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

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

Summary

We reviewed a large number of *voting rules* and observed:

- they explore different *intuitions* about how voting 'should' work
- they differ in view of the *profile information* they require
- they differ in view of their *computational requirements*
- they differ in view of the *normative principles* they satisfy

What next? Different approaches to characterising a voting rule.