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

The Gibbard-Satterthwaite Theorem tells us that there are no good voting rules that are strategyproof. *That's very bad!*

We are going to review three approaches for coping with this problem:

- *Domain restrictions:* excluding problematic profiles
- *Computational barriers:* making manipulation intractable
- *Informational barriers:* hiding information from manipulators

We'll conclude with briefly mentioning a few related topics.

Domain Restriction: Single-Peaked Preferences

Every voting rule can be manipulated, but not in all profiles. Can we do better if we restrict attention to specific (natural) profiles?

A profile $\boldsymbol{R} = (R_1, \ldots, R_n)$ is *single-peaked* if we can arrange the alternatives from left to right along some dimension $\triangleright$ such that R_i ranks x above y whenever x is between y and $\mathrm{top}(R_i)$ according to $\triangleright$.

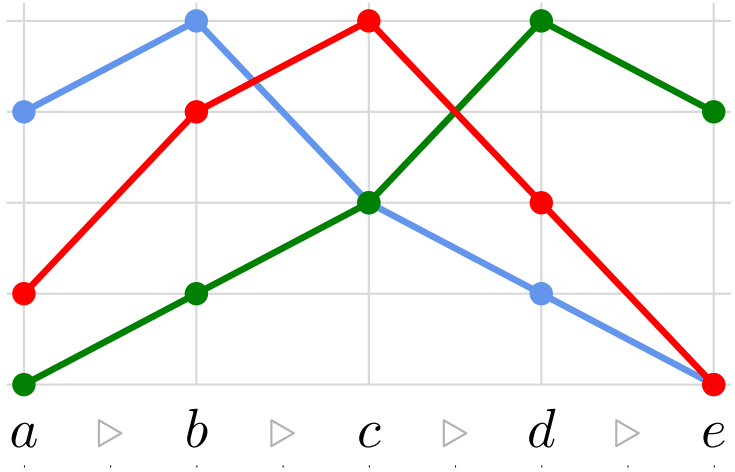

Sometimes a natural assumption: traditional political parties, spatial voting, agreeing on a number (e.g., legal drinking age), ...

Strategyproofness of the Median-Voter Rule

For a given dimension $\triangleright$, the *median-voter rule* asks each voter for her top alternative and elects the alternative proposed by the voter corresponding to the median w.r.t. $\triangleright$.

Folklore Theorem: *If an odd number of voters have preferences that are single-peaked w.r.t. $\triangleright$, then the median-voter rule is strategyproof.*

Proof: W.l.o.g., our manipulator's top alternative is *to the right* of the median (the winner). If she declares a peak further to the right, nothing will change. If she declares a peak further to the left, either nothing will change, or the new winner will be even worse. ✓

This is closely related to Black's *Median Voter Theorem*, showing that under the same conditions a Condorcet winner exists and is elected.

D. Black. On the Rationale of Group Decision-Making. *J. Polit. Econ.*, 1948.

Computational Barriers to Manipulation

Every voting rule can be manipulated in some profiles. But even when it is *possible* to manipulate, maybe actually doing so is *difficult*?

Tools from *complexity theory* can help make this idea precise:

> *If manipulation is computationally intractable for F, then we might consider F resistant (but not immune) to manipulation.*

Most interesting for voting rules for which computing election winners is tractable. Want to see a *complexity gap* between manipulation (undesired behaviour) and winner determination (desired functionality).

V. Conitzer and T. Walsh. Barriers to Manipulation in Voting. In F. Brandt et al. (eds.), *Handbook of Computational Social Choice*. CUP, 2016.

Classical Results

The seminal paper by Bartholdi, Tovey and Trick (1989) starts by showing that manipulation is in fact *easy* for a range of commonly used voting rules, and then presents one system (a variant of the Copeland rule) for which manipulation is NP-complete. Next:

- Easiness results for *plurality* and *Borda*
- Hardness result for *instant runoff* (Bartholdi and Orlin, 1991)

J.J. Bartholdi III, C.A. Tovey, and M.A. Trick. The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare*, 1989.

J.J. Bartholdi III and J.B. Orlin. Single Transferable Vote Resists Strategic Voting. *Social Choice and Welfare*, 1991.

Manipulability as a Decision Problem

We can cast the problem of strategic manipulability, for a given voting rule F, as a decision problem:

$\textsc{Manipulability}(F)$
Input: Ballots for all but one voter; alternative x.
Question: Is there a ballot for the final voter such that x wins?

To find out what the best winner achievable for the manipulator is, she has to solve $\textsc{Manipulability}(F)$ for all x, in order of her preference.

If $\textsc{Manipulability}(F)$ is intractable, then manipulability may be considered less of a worry for F.

Remark: An algorithm for this *decision problem* does not immediately also solve the *search problem* of finding the manipulating ballot.
As our focus here is on intractability results, this is ok.

Manipulating the Plurality Rule

Recall: under *plurality*, the alternative(s) ranked first most often win(s).

The plurality rule is easy to manipulate (trivial):

- Simply vote for x, the alternative to be made winner by means of manipulation. If manipulation is possible at all, this will work. Otherwise manipulation is not possible.

Thus: $\textsc{Manipulability}(\mathit{plurality})$ can be decided in *polynomial* time.

General: $\textsc{Manipulability}(F) \in \mathrm{P}$ for any rule F with polynomial winner determination problem and polynomial number of ballots.

Exercise: *Give another example for a voting rule this applies to!*

Manipulating the Borda Rule

Recall: under *Borda*, you submit a ranking of all alternatives and thereby award $m-k$ points to the alternative ranked in position k.

Remark: We now have superpolynomially-many possible ballots.

But Borda still is easy to manipulate. Use a *greedy algorithm:*

- Place x (the alternative to be made winner through manipulation) at the top of your ballot.
- Then inductively proceed as follows: Check if any of the remaining alternatives can be put next on the ballot without preventing x from winning. If yes, do so. (If no, manipulation is impossible.)

After convincing ourselves that this algorithm is indeed correct, we see that $\textsc{Manipulability}(Borda)$ can be decided in *polynomial* time.

J.J. Bartholdi III, C.A. Tovey, and M.A. Trick. The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare*, 1989.

Intractability of Manipulating Instant Runoff

Recall: *instant runoff* (a.k.a. *single transferable vote*) works by eliminating plurality losers until an alternative is the majority winner.

Theorem (Bartholdi and Orlin, 1991): *The* MANIPULABILITY *problem for instant runoff is NP-complete.*

Proof: Omitted. But try to get an intuition for why this is intractable.

For example, it is often not optimal to put the alternative x you want to win at the top of your ballot (by ranking y at the top, you may be able to eliminate z, which may be a stronger competitor than y).

J.J. Bartholdi III and J.B. Orlin. Single Transferable Vote Resists Strategic Voting. *Social Choice and Welfare*, 1991.

Coalitional Manipulation

It rarely is the case that a *single* voter really can make a difference.
So we should look into *manipulation by a coalition* of voters.

Variants of the problem:

- Ballots may be *weighted* or *unweighted*.

  Examples: countries in the EU, shareholders of a company

- Manipulation may be *constructive* (making alternative x win) or *destructive* (ensuring x does not win).

Decision Problems

Next, we consider two decision problems, for a given voting rule F:

ConstructiveManipulability(F)
Input: List of weighted ballots; set of weighted manipulators; $x \in A$.
Question: Are there ballots for the manipulators such that x wins?

DestructiveManipulability(F)
Input: List of weighted ballots; set of weighted manipulators; $x \in A$.
Question: Are there ballots for the manipulators such that x loses?

Exercise: *Which one seems harder? What's your intuition?*

Constructive Manipulation under Borda

In the context of coalitional manipulation with weighted voters, we can get hardness results for elections with small numbers of alternatives:

Theorem (Conitzer et al., 2007): *For weighted Borda and $m \geqslant 3$, the constructive coalitional manipulation problem is NP-complete.*

Proof: We have to prove NP-membership and NP-hardness:

- NP-membership: easy (if you guess ballots for the manipulators, we can check that it works in polynomial time)
- NP-hardness: for three alternatives by reduction from Partition (next slide); hardness for more alternatives follows

V. Conitzer, T. Sandholm, and J. Lang. When are Elections with Few Candidates Hard to Manipulate? *Journal of the ACM*, 2007.

Proof of NP-hardness

We use a reduction from the NP-complete Partition problem:

Partition
Input: $(w_1, \ldots, w_n) \in \mathbb{N}^n$
Question: Is there a set $I \subseteq \{1, \ldots, n\}$ s.t. $\sum_{i \in I} w_i = \frac{1}{2} \sum_{i=1}^{n} w_i$?

Let $K := \sum_{i=1}^{n} w_i$. Given an instance of Partition, we construct an election with $n + 2$ weighted voters and three alternatives:

- two voters with weight $\frac{1}{2}K - \frac{1}{4}$, voting $(x \succ y \succ z)$ and $(y \succ x \succ z)$
- a coalition of n voters with weights $w_1, \ldots, w_n$ who want z to win

Clearly, each manipulator should vote either $(z \succ x \succ y)$ or $(z \succ y \succ x)$.
Suppose there does exist a partition. Then they can vote like this:

- manipulators corresponding to elements in I vote $(z \succ x \succ y)$
- manipulators corresponding to elements outside I vote $(z \succ y \succ x)$

Scores: $2K$ for z; $\frac{1}{2}K + (\frac{1}{2}K - \frac{1}{4}) \cdot (2 + 1) = 2K - \frac{3}{4}$ for both x and y
If there is no partition, then either x or y will get at least 1 point more.
Hence, manipulation is feasible iff there exists a partition. ✓

Destructive Manipulation under Borda

Theorem (Conitzer et al., 2007): *For weighted Borda and $m \geqslant 3$, the destructive coalitional manipulation problem is in P.*

Proof: Let x be the alternative the manipulators want to lose.

For every $y \neq x$, simply try everyone ranking y at the top and x at the bottom. If none of these $m-1$ attempts work, nothing will. ✓

V. Conitzer, T. Sandholm, and J. Lang. When are Elections with Few Candidates Hard to Manipulate? *Journal of the ACM*, 2007.

Criticism

Such complexity results provide interesting insights into the dynamics of strategic manipulation. *But do they really offer protection?*

NP-hardness is a *worst-case* notion and cannot rule out the possibility that problem instances encountered in practice are easy to solve.

Research suggests that it might be impossible to find a voting rule that is *usually* hard to manipulation—for some definition of "usual". See Conitzer and Walsh (2016) for a discussion.

V. Conitzer and T. Walsh. Barriers to Manipulation in Voting. In F. Brandt et al. (eds.), *Handbook of Computational Social Choice*. CUP, 2016.

Informational Barriers to Manipulation

Suppose voter i has only *partial information* about the profile. If π is a function mapping any truthful profile $\boldsymbol{R}$ to the information $\pi(\boldsymbol{R})$ given to i, then i must consider possible any profile in this set:

$$\mathcal{W}_i^{\pi(\boldsymbol{R})} \quad = \quad \{\, \boldsymbol{R}' \in \mathcal{L}(A)^n \mid \pi(\boldsymbol{R}) = \pi(\boldsymbol{R}') \text{ and } R_i = R_i' \,\}$$

Example: π might be an *opinion poll* that returns, say, the winner of the election, or the plurality score of every alternative.

If i is cautious, she will play R_i' only if doing so is *strictly better* for her in at least one profile in $\mathcal{W}_i^{\pi(\boldsymbol{R})}$ and *no worse* in any of the others.

Remark: Interesting, still very much underexplored research direction.

A. Reijngoud and U. Endriss. Voter Response to Iterated Poll Information. AAMAS-2012.

Antiplurality Rule and Winner Information

One of the very few positive results available to date:

Theorem (Reijngoud and Endriss, 2012): *For $n \geqslant 2m - 2$, the antiplurality rule (with ties getting broken lexicographically) is strategyproof if voters only know the winner for the truthful profile.*

Proof: Suppose $m \geqslant 3$ (other cases: clear). Consider voter $i \in N$. Let x be voter i's *worst alternative*. Let $x^\star$ be the truthful *winner*.

Distinguish two cases:

- $x = x^\star$: Nothing she can do to change the outcome. ✓
- $x \neq x^\star$: If i manipulates by vetoing some $y \neq x$ (possibly $y = x^\star$), then x gains a point and $x^\star$ does not, so x *could* now win. ✓

For full details, see Annemieke Reijngoud's MoL thesis from 2011.

A. Reijngoud and U. Endriss. Voter Response to Iterated Poll Information. AAMAS-2012.

Iterative Voting

What happens if voters keep responding to the current preference profile by updating their own ballot? Will this converge?

Example: people revisiting their choice in a Doodle poll

Main result in the area shows that (under certain assumptions), we get convergence for *plurality*, unlike for most other rules (Meir et al., 2010).

R. Meir, M. Polukarov, J.S. Rosenschein, and N.R. Jennings. Convergence to Equilibria in Plurality Voting. AAAI-2010.

R. Meir. Iterative Voting. In U. Endriss (ed.), *Trends in Computational Social Choice*. AI Access, 2017.

Frequency of Manipulability

While essentially all voting rules can be manipulated, some intuitively are "more" manipulable than others, in the sense of there being more preference profiles where manipulation is possible.

To make such claims concrete and relevant, one would need to make assumptions regarding a probability distribution on profiles (or fix a dataset of real-world profiles of interest).

Probabilistic Social Choice

For all formal results so far we insisted on rules being deterministic.
But in some settings *randomisation* can make sense.

To study strategic behaviour we now need to *lift preferences* from alternatives to *lotteries over alternatives*. Different ways of doing that.

Some ways lead to similarly negative results as the G-S Theorem (basically: only randomised dictatorships work).

But for other assumptions more positive results are possible.

F. Brandt. Rolling the Dice: Recent Results in Probabilistic Social Choice. In U. Endriss (ed.), *Trends in Computational Social Choice*. AI Access, 2017.

Summary

Previously, we saw that *strategic manipulation* is a major problem in voting: essentially, only dictatorships are strategyproof.

Today we discussed approaches to *circumventing* this problem:

- *Domain restrictions:* if we can find a natural and large class of preference profiles (+ ballot restrictions) that make strategic manipulation impossible, then that will sometimes suffice.
- *Computational barriers:* maybe strategic manipulation is so difficult, in computational terms, so as to provide protection.
- *Informational barriers:* maybe strategic manipulation would require information about the profile the manipulator does not possess.

What next? Matters of information and communication in voting.