Computational Social Choice: Spring 2017

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

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

Domain Restriction: Single-Peaked Preferences

We only discuss the oldest and most famous domain restriction . . .

A profile $(\succ_1, \ldots, \succ_n)$ is *single-peaked* if there exists a "left-to-right" ordering $\gg$ on the alternatives such that $x \succ_i y$ for voter i whenever x is $\gg$-between y and $\mathrm{top}(\succ_i)$.

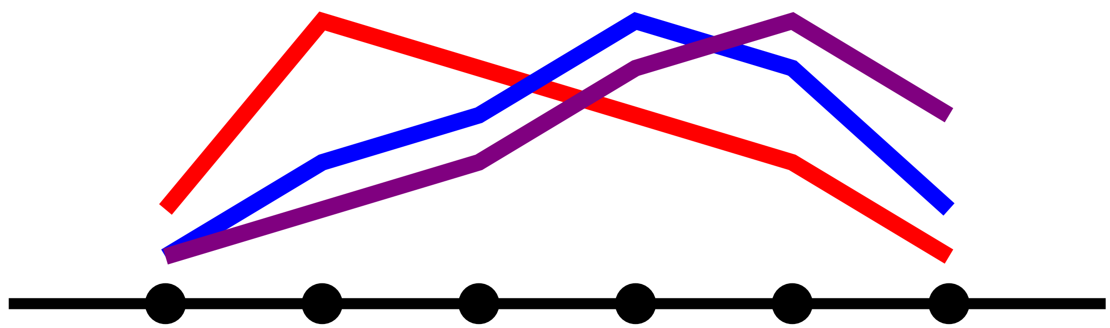

Sometimes a natural assumption: traditional political parties, agreeing on a number (e.g., legal drinking age), . . .

D. Black. On the Rationale of Group Decision-Making. *The Journal of Political Economy*, 56(1):23–34, 1948.

Strategyproofness of the Median-Voter Rule

For a given left-to-right ordering $\gg$, 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. $\gg$.

Theorem 1 *If an odd number of voters have preferences that are single-peaked w.r.t. a fixed left-to-right ordering $\gg$, then the median-voter rule (w.r.t. $\gg$) 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. *The Journal of Political Economy*, 56(1):23–34, 1948.

Complexity as a Barrier against 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 F might be considered *resistant* (albeit still not *immune*) to manipulation.
- Even if standard voting rules turn out to be easy to manipulate, it might still be possible to *design new ones* that are resistant.

Remark: This approach is most interesting for voting rules for which computing election winners is tractable. At least, we would like 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:

- We first present a couple of these easiness results, namely for *plurality* and for the *Borda rule*.
- We then mention a result from a follow-up paper by Bartholdi and Orlin (1991): the manipulation of *STV* is *NP-complete*.

J.J. Bartholdi III, C.A. Tovey, and M.A. Trick. The Computational Difficulty of Manipulating an Election. *Soc. Choice and Welfare*, 6(3):227–241, 1989.

J.J. Bartholdi III and J.B. Orlin. Single Transferable Vote Resists Strategic Voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

Manipulability as a Decision Problem

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

> $\textsc{Manipulability}(F)$
> **Instance:** 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: This simple formulation of the *decision problem* cannot be used to solve the *search problem* of computing 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: $\textsc{Manipulability}(\textit{plurality})$ 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.

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

J.J. Bartholdi III, C.A. Tovey, and M.A. Trick. The Computational Difficulty of Manipulating an Election. *Soc. Choice and Welfare*, 6(3):227–241, 1989.

Intractability of Manipulating STV

Recall: *Single Transferable Vote* (STV) works by eliminating plurality losers until an alternative is ranked first by $> 50\%$ of the voters.

Theorem 2 (Barthold and Orlin, 1991) $\textsc{Manipulability}(STV)$ *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*, 8(4):341–354, 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)
Instance: List of weighted ballots; set of weighted manipulators; $x \in X$.
Question: Are there ballots for the manipulators such that x wins?

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

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 3 (Conitzer et al., 2007) *For the Borda rule, the constructive coalitional manipulation problem with weighted voters is NP-complete for $\geqslant 3$ alternatives.*

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*, 54(3), Article 14, 2007.

Proof of NP-hardness

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

PARTITION
Instance: $(w_1, \ldots, w_n) \in \mathbb{N}^n$
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 <u>iff</u> there exists a partition. ✓

Destructive Manipulation under Borda

Theorem 4 (Conitzer et al., 2007) *For the Borda rule, the destructive coalitional manip. problem with weighted voters 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*, 54(3), Article 14, 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 enountered 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 a suitable 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.

Manipulation under Partial Information

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

$$\mathcal{W}_i^{\pi(\boldsymbol{\succ})} \quad = \quad \{\boldsymbol{\succ}' \in \mathcal{L}(X)^n \mid \pi(\boldsymbol{\succ}) = \pi(\boldsymbol{\succ}') \text{ and } \succ_i = \succ_i'\}$$

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 manipulate using $\succ_i^\star$ instead of $\succ_i$ only if both:

- $F(\succ_i^\star, \boldsymbol{\succ}'_{-i}) \succ_i F(\succ_i, \boldsymbol{\succ}'_{-i})$ *for some* $\boldsymbol{\succ}'_{-i}$ with $(\succ_i, \boldsymbol{\succ}'_{-i}) \in \mathcal{W}_i^{\pi(\boldsymbol{\succ})}$
- $F(\succ_i^\star, \boldsymbol{\succ}'_{-i}) \succcurlyeq_i F(\succ_i, \boldsymbol{\succ}'_{-i})$ *for all* $\boldsymbol{\succ}'_{-i}$ with $(\succ_i, \boldsymbol{\succ}'_{-i}) \in \mathcal{W}_i^{\pi(\boldsymbol{\succ})}$

V. Conitzer, T. Walsh, and L. Xia. Dominating Manipulations in Voting with Partial Information. Proc. AAAI-2011.

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

Example: Manipulation under Zero Information

When does lack of information constitute a barrier to manipulation?

Contrary to intuition, some (strange) voting rules *can be manipulated* even if you have *no information* at all:

> Take the voting rule that elects the Condorcet winner if it exists, and otherwise the bottom alternative of voter 1.

If voter 1's true preferences are $x \succ_1 y \succ_1 z$, she can never do worse by voting $x \succ z \succ y$, and she does better if the others vote like this:

$$x \succ_2 y \succ_2 z$$

$$y \succ_3 x \succ_3 z$$

$$y \succ_4 x \succ_4 z$$

Antiplurality Rule and Winner Information

One of the very few positive results available to date:

Theorem 5 (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. Proc. AAMAS-2012.

Summary

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

Today we have 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.
- *Complexity barriers:* maybe strategic manipulation will turn out to 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? Other aspects of informational requirements in voting.