

Computational Social Choice: Autumn 2013

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

Plan for Today

The Gibbard-Satterthwaite Theorem tells us that there aren't any reasonable voting rules that are strategy-proof. *That's very bad!*

We will consider three possible avenues of dealing with this problem:

- Changing the formal framework a little (one slide only)
- Restricting the domain (the classical approach)
- Making strategic manipulation computationally hard

Changing the Framework

The Gibbard-Satterthwaite Theorem applies when both preferences and ballots are linear orders. The problem persists for several variations. But:

- In a framework with *money*, if preferences and ballots are modelled as (quasi-linear) *utility functions* $u : \mathcal{X} \rightarrow \mathbb{R}$, we can design strategy-proof mechanisms. Example: *Vickrey Auction* (winner pays second price)
- In the context of *approval voting* (ballots $\in 2^{\mathcal{X}}$, preferences $\in \mathcal{L}(\mathcal{X})$), under certain conditions we can ensure that no voter has an incentive to vote *insincerely* (weak variant of strategy-proofness).
- More generally, for any *preference language* and *ballot language*, we can define a notion of *sincerity* and study incentives to be sincere.

W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16(1):8–37, 1961.

U. Endriss. Sincerity and Manipulation under Approval Voting. *Theory and Decision*. 74(3):335–355, 2013.

U. Endriss, M.-S. Pini, F. Rossi, and K.B. Venable. Preference Aggregation over Restricted Ballot Languages: Sincerity and Strategy-Proofness. Proc. IJCAI-2009.

Domain Restrictions

- Note that we have made an implicit *universal domain* assumption: *any* linear order may come up as a preference or ballot.
- If we *restrict* the domain (possible ballot profiles + possible preferences), more voting rules will satisfy more axioms ...

Single-Peaked Preferences

An electorate \mathcal{N} has *single-peaked* preferences if there exists a “left-to-right” ordering \gg on the alternatives such that any voter prefers x to y if x is between y and her top alternative w.r.t. \gg .

The same definition can be applied to profiles of ballots.

Remarks:

- Quite natural: classical spectrum of political parties; decisions involving agreeing on a number (e.g., legal drinking age); ...
- But certainly not universally applicable.

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

Strategy-Proofness of the Median-Voter Rule

For a given left-to-right ordering \succcurlyeq , the *median-voter rule* asks each voter for their top alternative and elects the alternative proposed by the voter corresponding to the median w.r.t. \succcurlyeq .

Theorem 1 (Strategy-proofness) *If an odd number of voters have preferences that are **single-peaked** w.r.t. a fixed left-to-right ordering \succcurlyeq , then the **median-voter rule** (w.r.t. \succcurlyeq) is **strategy-proof**.*

Proof: W.l.o.g., suppose our manipulator's top alternative is to the right of the median (the winner). She has two options:

- Nominate some other alternative to the right of the current winner (or the winner itself). Then the median/winner does not change.
- Nominate an alternative to the left of the current winner. Then the new winner will be to the left of the old winner, which—by the single-peakedness assumption—is worse for our manipulator.

Thus, manipulation has either no effect or results in a worse outcome. ✓

Black's Median Voter Theorem

The result on strategy-proofness presented here is closely related to Black's original result for single-peaked domains:

Theorem 2 (Black's Theorem, 1948) *If an odd number of voters submit **single-peaked** ballots, then there exists a **Condorcet winner** and it will get elected by the **median-voter** rule.*

Proof: Very similar to the preceding result.

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

More on Domain Restrictions

This is a big topic in SCT. We have only scratched the surface here.

- It suffices to enforce single-peakedness for *triples* of alternatives.
- Moulin (1980) gives a *characterisation* of the class of voting rules that are strategy-proof for single-peaked domains: median-voter rule + addition of “phantom peaks”
- Sen’s *triplewise value restriction* is more powerful and also guarantees strategy-proofness: for any three alternatives (x, y, z) , there exist an $x^* \in \{x, y, z\}$ and a value $v^* \in \{\text{“best”}, \text{“middle”}, \text{“worst”}\}$ such that x^* never has value v^* w.r.t. (x, y, z) for any voter.
- The aforementioned strategy-proofness result for the *Vickrey auction* may also be cast as a result for a domain restriction.

H. Moulin. On Strategy-Proofness and Single Peakedness. *Public Choice*, 35(4):437–455, 1980.

A.K. Sen. A Possibility Theorem on Majority Decisions. *Econometrica*, 34(2):491–499, 1966.

Complexity as a Barrier against Manipulation

The Gibbard-Satterthwaite Theorem shows that strategic manipulation cannot be ruled out.

Idea: So it's always *possible* to manipulate; but maybe it's also *difficult*? Tools from *complexity theory* can 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.
- This approach is most interesting for voting rules for which the problem of computing election winners is tractable. At least, we want to see a *complexity gap* between manipulation (undesired behaviour) and winner determination (desired functionality).

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:

MANIPULABILITY(F)

Instance: Set of ballots for all but one voter; alternative x .

Question: Is there a ballot for the final voter such that x wins?

A manipulator has to solve MANIPULABILITY(F) for all alternatives, in order of her preference. (Note that in practice the manipulator does not just want a yes/no answer, but the manipulating ballot.)

If MANIPULABILITY(F) is computationally intractable, then manipulability may be considered less of a worry for F .

Remark: We assume that the manipulator knows all the other ballots. This unrealistic assumption is intentional: if manipulation is intractable even under such favourable conditions, then all the better.

Manipulating the Plurality Rule

Recall 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.

That is, we have $\text{MANIPULABILITY}(\textit{plurality}) \in \text{P}$.

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

Manipulating the Borda Rule

Recall Borda: submit a ranking (super-polynomially many choices!) and give $m-1$ points to 1st ranked, $m-2$ points to 2nd ranked, etc.

The Borda rule is also 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 also get $\text{MANIPULABILITY}(\textit{Borda}) \in \text{P}$.

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

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

Theorem 3 (Bartholdi and Orlin, 1991) $\text{MANIPULABILITY}(STV)$ is NP-complete.

Proof: Omitted.

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 will rarely be the case that a *single* voter 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 a *unique* or *tied* winner) or *destructive* (ensuring x does not win).

Decision Problems

On the following slides, we will consider two decision problems, for a given voting rule F :

CONSTRUCTIVEMANIPULABILITY(F)

Instance: Set of weighted ballots; set of weighted manipulators; $x \in \mathcal{X}$.

Question: Are there ballots for the manipulators such that x wins?

DESTRUCTIVEMANIPULABILITY(F)

Instance: Set of weighted ballots; set of weighted manipulators; $x \in \mathcal{X}$.

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 4 (Conitzer et al., 2007) *Under the **Borda** rule, the **constructive** coalitional manipulation problem with weighted voters is **NP-complete** for ≥ 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 will use a reduction from the NP-complete PARTITION problem:

PARTITION

Instance: $(w_1, \dots, w_n) \in \mathbb{N}^n$

Question: Is there a set $I \subseteq \{1, \dots, n\}$ 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, \dots, 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 5 (Conitzer et al., 2007) *Under 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. The following algorithm will find a manipulation, if one exists:

For each alternative $y \neq x$, try letting all manipulators rank y first, x last, and the other alternatives in any fixed order.

If x loses in one of these $m-1$ elections, then manipulation is possible; otherwise it is not.

Correctness of the algorithm follows from the fact that (a) the best we can do about x is not to give x any points and, (b) if any other alternative y has a chance of beating x , she will do so if we give y a maximal number of points. ✓

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.

Worst-Case vs. Average-Case Complexity

NP-hardness is only a *worst-case* notion. Do NP-hardness barriers provide sufficient protection against manipulation?

What about the *average complexity* of strategic manipulation?

Some recent work suggests that it might be impossible to find a voting rule that is *usually* hard to manipulation—for a suitable definition of “usual”. See Faliszewski and Procaccia (2010) for a discussion.

P. Faliszewski and A.D. Procaccia. AI's War on Manipulation: Are We Winning? *AI Magazine*, 31(4):53–64, 2010.

Controlling Elections

Strategic manipulation is not the only undesirable form of behaviour in voting we may want to contain by means of complexity barriers . . .

People have studied the computational complexity of a range of different types of *control* in elections:

- Adding or removing *candidates*.
- Adding or removing *voters*.
- Redefining *districts* (if your party is likely to win district A with an 80% majority and lose district B by a small margin, you might win both districts if you carefully redraw the district borders . . .).

See Faliszewski et al. (2009) for an introduction to this area.

P. Faliszewski, E. Hemaspaandra, L.A. Hemaspaandra, and J. Rothe. *A Richer Understanding of the Complexity of Election Systems*. In *Fundamental Problems in Computing*, Springer-Verlag, 2009.

Bribery in Elections

Bribery is the problem of finding $\leq K$ voters such that a suitable change of their ballots will make a given candidate x win.

- Connection to *manipulation*: in the (coalitional) manipulation problem the names of the voters changing ballot are part of the input, while for the bribery problem we need to choose them.
- Several *variants* of the bribery problem have been studied: when each voter has a possibly different “price”; when bribes depend on the extent of the change in the bribed voter’s ballot; etc.

People have studied the complexity of several variants of the bribery problem for various voting rules (e.g., Faliszewski et al., 2009).

P. Faliszewski, E. Hemaspaandra, and L.A. Hemaspaandra. How Hard is Bribery in Elections? *Journal of Artificial Intelligence Research*, 35:485–532, 2009.

Summary

Previously, we have seen that *strategic manipulation* is a major problem in voting: essentially, only dictatorships are strategy-proof.

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 be sufficiently hard computationally to provide protection.

A related question, which we have not addressed, deals with the *frequency of manipulability*, using either empirical methods or devising formal models regarding the distribution of voter preferences.

What next?

In the remaining lectures on voting, we will go more deeply into questions of a computational nature:

- *Information and communication*: What can we say about the status of an election when we only have incomplete information regarding preferences/ballots?
- *Combinatorial domains*: How can we conduct elections on outcomes with multiple attributes, given that the number of outcomes is exponential in the number of attributes in this case?