

Computational Social Choice: Autumn 2010

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 procedures that are strategy-proof. *That's very bad!*

We will consider three possible avenues to circumvent this problem:

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

Recap: Strategic Manipulation

We had seen two theorems that show that we cannot rule out strategic manipulation: any reasonable voting procedure will sometimes give a voter an incentive to misrepresent her preferences.

Theorem 1 (Gibbard-Satterthwaite) Any *resolute* voting procedure for ≥ 3 alternatives that is *surjective* and *strategy-proof* is *dictatorial*.

Theorem 2 (Duggan-Schwartz) Any voting procedure for ≥ 3 alternatives that is *nonimposed* and *immune to manipulation* by both *optimistic* and *pessimistic* voters is *weakly dictatorial*.

A. Gibbard. Manipulation of Voting Schemes: A General Result. *Econometrica*, 41(4):587–601, 1973.

M.A. Satterthwaite. Strategy-proofness and Arrow's Conditions. *Journal of Economic Theory*, 10:187–217, 1975.

J. Duggan and T. Schwartz. Strategic Manipulation w/o Resoluteness or Shared Beliefs: Gibbard-Satterthwaite Generalized. *Soc. Choice Welf.*, 17(1):85–93, 2000.

Approach 1: Domain Restrictions

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 procedures 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 wrt. \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.

Black's Median Voter Theorem

For simplicity, assume the number of voters is *odd*.

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

Theorem 3 (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.*

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

Proof Sketch

The candidate elected by the median-voter rule is a Condorcet winner:

Proof: Let x be the winner and compare x to some y to, say, the left of x . As x is the median, for more than half of the voters x is between y and their favourite, so they prefer x . ✓

Note that this also implies that a Condorcet winner exists.

As the Condorcet winner is (always) unique, it follows that, also, every Condorcet winner is a median-voter rule election winner. ✓

Strategy-Proofness

The following result is a corollary of Black's Theorem:

Theorem 4 (Strategy-proofness) *If an odd number of voters have preferences that are **single-peaked** wrt. a fixed left-to-right ordering \gg , then the **median-voter rule** (wrt. \gg) is **strategy-proof**.*

Direct 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 candidate to the right of the current winner (or the winner itself). Then the median/winner does not change.
- Nominate a candidate 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, misrepresenting preferences has either no effect or results in a worse outcome. ✓

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 strategy-proof voting procedures for single-peaked domains: median-voter rule + addition of “phantom peaks”
- Sen’s *triplewise value restriction* is a more powerful domain restriction that also guarantees strategy-proofness: for any triple of alternatives (x, y, z) , there exist one $x^* \in \{x, y, z\}$ and one value in $v^* \in \{\text{“best”}, \text{“middle”}, \text{“worst”}\}$ such that x^* never has value v^* wrt. (x, y, z) for any voter.

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.

Approach 2: Varying the Formal Framework

Varying the Formal Framework

The Gibbard-Satterthwaite and the Duggan-Schwartz Theorem say something about functions of the form $F : \mathcal{L}(\mathcal{X})^{\mathcal{N}} \rightarrow 2^{\mathcal{X}} \setminus \{\emptyset\}$ only.

It is thus conceivable, at least in principle, that “strategy-proofness” (suitably redefined), is possible for slightly different ways of modelling ballots and preferences.

We have to check what is possible and impossible for any choice of *ballot language* $\mathcal{B}(\mathcal{X})$ and any *class of preference structures* $\mathcal{P}(\mathcal{X})$.

We will briefly look into two examples:

- *Auctions*, where preferences and ballots are utility functions $u : \mathcal{X} \rightarrow \mathbb{R}$ (informally only).
- *Approval Voting*, where preferences are standard and ballots are sets of alternatives: $F : (2^{\mathcal{X}})^{\mathcal{N}} \rightarrow 2^{\mathcal{X}} \setminus \{\emptyset\}$.

Vickrey Auctions

(Note that we had already discussed this in the introductory lecture.)

Suppose we want to sell a single item in an auction.

- *First-price sealed-bid auction*: each bidder submits an offer in a sealed envelope (which encodes a utility function for this simple domain); highest bidder wins and pays what she offered
- *Vickrey auction*: each bidder submits an offer in a sealed envelope; highest bidder wins but pays the *second highest price*

In the Vickrey auction each bidder has an incentive to submit their *truthful valuation* of the item!

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

Approval Voting

Recall approval voting: voters can approve of any set of alternatives and the alternative(s) with the most approvals win(s)

If $\mathcal{B}(\mathcal{X}) = 2^{\mathcal{X}}$ but still $\mathcal{P}(\mathcal{X}) = \mathcal{L}(\mathcal{X})$, then what is “*truthful voting*”?

Replace this by the weaker notion of *sincere voting*:

- Ballot $b \in 2^{\mathcal{X}}$ is said to be *sincere* given preference order \succ if $x \succ y$ for all $x \in b$ and all $y \notin b$.

To study strategic manipulation for AV we also require:

- A way of *extending a preference* \succeq_i on \mathcal{X} to a preference $\hat{\succeq}_i$ on $2^{\mathcal{X}} \setminus \{\emptyset\}$, to be able to speak about the incentives of voters regarding election outcomes (which could be tied).
- Call a voting procedure *immune to insincere manipulation* if no voter who knows the other ballots ever has an incentive to vote insincerely. [if sincerity = truthfulness, this is strategy-proofness]

AV and Insincere Manipulation

Suppose we break ties using a *uniform* probability distribution.

A voter i might be an *expected-utility maximiser*:

- Voter i has a utility function $u_i : \mathcal{X} \rightarrow \mathbb{R}$, but all we know about u_i is that $u_i(x) > u_i(y)$ iff $x \succ_i y$ (we say: u_i is *compatible* with \succ_i).
- Voter i will prefer set X over Y if it has higher expected utility.

These assumptions give rise to a weak order $\hat{\succeq}_i$ on $2^{\mathcal{X}} \setminus \{\emptyset\}$:

- $X \hat{\succeq}_i Y$ iff there exists (“*pessimistic interpretation*”) a utility function u_i compatible with \succ_i such that $\frac{1}{|X|} \cdot \sum_{x \in X} u_i(x) > \frac{1}{|Y|} \cdot \sum_{y \in Y} u_i(y)$.

Theorem 5 *Approval voting with uniform tie-breaking is immune to insincere manipulation by expected-utility maximisers.*

Proof: Omitted.

U. Endriss. Vote Manipulation in the Presence of Multiple Sincere Ballots. Proc. TARK-2007.

Approach 3: Complexity Barriers

Complexity as a Barrier against Manipulation

The Gibbard-Satterthwaite Theorem shows that (in the standard model) strategic manipulation can never be rule out.

Idea: So it is always *possible* to manipulate; but may 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 procedures 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 procedures for which winner determination 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 procedures, 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 present 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 procedure 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?

In practice, a manipulator would have to solve MANIPULABILITY(F) for all alternatives, in order of her preference.

If the MANIPULABILITY(F) is computationally intractable, then manipulability may be considered less of a worry for procedure 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 not.

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

Recall STV: eliminate plurality losers until an alternative gets $> 50\%$

Theorem 6 (Bartholdi and Orlin, 1991) $\text{MANIPULABILITY}(STV)$ is NP-complete.

Proof sketch: We need to show NP-hardness and NP-membership.

- NP-membership is clear: checking whether a given ballot makes x win can be done in polynomial time (just try it).
- NP-hardness: by reduction from 3-COVER. The proof is long and somewhat tedious, but not very difficult. The basic idea is to build a large election instance introducing all sorts of constraints on the ballot of the manipulator, such that finding a ballot meeting those constraints solves a given instance of 3-COVER as a by-product.

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:

CONSTRUCTIVE MANIPULATION(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?

DESTRUCTIVE MANIPULATION(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 7 (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 8 (Conitzer et al., 2007) *Under the Borda rule, the destructive coalitional manip. problem with weighted voters is in P.*

Proof sketch: 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 procedure 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*. In press (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 procedures (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

While strategic manipulation is a major problem in voting, we have seen that there are several ways to (partially) circumvent it:

- *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.
- *Framework variations*: maybe the standard framework used in voting theory is not exactly what we need anyway, and maybe for a different framework some problems can be ruled out.
- *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?

Next we will address the problems that arise when the set of alternatives has a *combinatorial structure*, such as here:

- In a referendum, we may get asked to vote on several (possibly related) propositions (approving or disapproving each of them).
- When voting for a committee, we have to decide for each candidate standing whether she should get elected.