

Tutorial on Voting Theory

Ulle Endriss

Institute for Logic, Language and Computation
University of Amsterdam

[<http://www.illc.uva.nl/~ulle/teaching/secvote-2012/>]

Table of Contents

Introduction	3
Voting Rules and their Properties	6
Strategic Manipulation	16
Computational Social Choice	25
Conclusion	31

Introduction

Voting Theory

Voting theory (which is part of *social choice theory*) is the study of methods for conducting an election:

- ▶ A group of *voters* each have *preferences* over a set of *candidates*. Each voter submits a *ballot*, based on which a *voting rule* selects a (set of) *winner(s)* from amongst the candidates.

This is not a trivial problem. Remember Florida 2000 (simplified):

49%: Bush \succ Gore \succ Nader

20%: Gore \succ Nader \succ Bush

20%: Gore \succ Bush \succ Nader

11%: Nader \succ Gore \succ Bush

Tutorial Overview

- Voting Rules
 - Such as: Plurality, Borda, Approval, Copleand ...
 - Properties and Paradoxes
- Strategic Manipulation
 - The Axiomatic Method in Voting Theory
 - The Gibbard-Satterthwaite Theorem
- Computational Social Choice
 - Introduction to the field
 - Examples for work involving voting

Voting Rules and their Properties

Three Voting Rules

How should n *voters* choose from a set of m *candidates*?

- *Plurality*: elect the candidate ranked first most often (i.e., each voter assigns one point to a candidate of her choice, and the candidate receiving the most votes wins).
- *Borda*: each voter gives $m-1$ points to the candidate she ranks first, $m-2$ to the candidate she ranks second, etc., and the candidate with the most points wins.
- *Approval*: voters can approve of as many candidates as they wish, and the candidate with the most approvals wins.

Example

Suppose there are three *candidates* (A, B, C) and 11 *voters* with the following *preferences* (where boldface indicates *acceptability*, for AV):

5 voters think: **A** \succ B \succ C

4 voters think: **C** \succ B \succ A

2 voters think: **B** \succ **C** \succ A

Assuming the voters vote *sincerely*, who *wins* the election for

- the plurality rule?
- the Borda rule?
- approval voting?

Positional Scoring Rules

We can generalise the idea underlying the Borda rule as follows:

A *positional scoring rule* is given by a *scoring vector* $s = \langle s_1, \dots, s_m \rangle$ with $s_1 \geq s_2 \geq \dots \geq s_m$ and $s_1 > s_m$.

Each voter submits a ranking of the m candidates. Each candidate receives s_i points for every voter putting her at the i th position.

The candidates with the highest score (sum of points) win.

For instance:

- The *Borda rule* is the positional scoring rule with the scoring vector $\langle m-1, m-2, \dots, 0 \rangle$.
- The *plurality rule* is the positional scoring rule with the scoring vector $\langle 1, 0, \dots, 0 \rangle$.
- The *antiplurality* or *veto rule* is the positional scoring rule with the scoring vector $\langle 1, \dots, 1, 0 \rangle$.

The Condorcet Principle

A candidate that beats every other candidate in pairwise majority contests is called a *Condorcet winner*.

There may be no Condorcet winner; witness the *Condorcet paradox*:

Ann: $A \succ B \succ C$

Bob: $B \succ C \succ A$

Cesar: $C \succ A \succ B$

Whenever a Condorcet winner exists, then it must be *unique*.

A voting rule satisfies the *Condorcet principle* if it elects (only) the Condorcet winner whenever one exists.

M. le Marquis de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris, 1785.

Positional Scoring Rules violate Condorcet

Consider the following example:

3 voters: $A \succ B \succ C$

2 voters: $B \succ C \succ A$

1 voter: $B \succ A \succ C$

1 voter: $C \succ A \succ B$

A is the *Condorcet winner*; she beats both B and C 4 : 3. But any *positional scoring rule* makes B win (because $s_1 \geq s_2 \geq s_3$):

$$A: 3 \cdot s_1 + 2 \cdot s_2 + 2 \cdot s_3$$

$$B: 3 \cdot s_1 + 3 \cdot s_2 + 1 \cdot s_3$$

$$C: 1 \cdot s_1 + 2 \cdot s_2 + 4 \cdot s_3$$

Thus, *no positional scoring rule* for three (or more) candidates will satisfy the *Condorcet principle*.

Condorcet-Consistent Rules

Some voting rules have been designed specifically to meet the Condorcet principle.

- *Copeland*: elect the candidate that maximises the difference between won and lost pairwise majority contests.
- *Dodgson*: elect the candidate that is “closest” to being a Condorcet winner, where “closeness” between two profiles is measured in terms of the number of swaps of adjacent candidates in a voter’s ranking required to move from one to the other.

A problem with the latter is that it is *computationally intractable*.

E. Hemaspaandra, L.A. Hemaspaandra, and J. Rothe. Exact Analysis of Dodgson Elections: Lewis Carroll’s 1876 Voting System is Complete for Parallel Access to NP. *Journal of the ACM*, 44(6):806–825, 1997.

Plurality with Run-Off

One more voting rule:

- *Plurality with run-off*: each voter initially votes for one candidate; the winner is elected in a second round by using the plurality rule with the two top candidates from the first round.

Example: French presidential elections

The No-Show Paradox

Under plurality with run-off, it may be *better to abstain* than to vote for your favourite candidate! Example:

25 voters: $A \succ B \succ C$

46 voters: $C \succ A \succ B$

24 voters: $B \succ C \succ A$

Given these voter preferences, B gets eliminated in the first round, and C beats A 70:25 in the run-off.

Now suppose two voters from the first group abstain:

23 voters: $A \succ B \succ C$

46 voters: $C \succ A \succ B$

24 voters: $B \succ C \succ A$

A gets eliminated, and B beats C 47:46 in the run-off.

P.C. Fishburn and S.J. Brams. Paradoxes of Preferential Voting. *Mathematics Magazine*, 56(4):207-214, 1983.

Insights so far / What next?

We have seen:

- There are many *different voting rules* (all of them looking more or less reasonable at first sight).
- Those rules can do surprisingly badly in some cases (“*paradoxes*”).

This is why:

- We need to be precise in formulating our requirements (“*axioms*”).
- A major part of *social choice theory* concerns the formal study of voting rules and the axioms they do or do not satisfy.

We will now focus on one such axiom and its formal treatment.

Strategic Manipulation

Strategic Manipulation

Recall our initial example:

49%: Bush \succ Gore \succ Nader
20%: Gore \succ Nader \succ Bush
20%: Gore \succ Bush \succ Nader
11%: Nader \succ Gore \succ Bush

Under the *plurality* rule, Bush will win the election.

Note that the Nader supporters have an incentive to *manipulate* by misrepresenting their preferences and vote for Gore instead of Nader (in which case Gore rather than Bush will win).

► Can we find a voting rule that avoids this problem?

Notation and Terminology

Set of n *voters* $\mathcal{N} = \{1, \dots, n\}$ and set of m *candidates* \mathcal{X} .

Both (true) *preferences* and (reported) *ballots* are modelled as linear orders on \mathcal{X} . $\mathcal{L}(\mathcal{X})$ is the set of all such linear orders.

A *profile* $\mathbf{R} = (R_1, \dots, R_n)$ fixes one preference/ballot for each voter.

We are looking for a *resolute voting rule* $F : \mathcal{L}(\mathcal{X})^{\mathcal{N}} \rightarrow \mathcal{X}$, mapping any given *profile* of ballots to a (single) *winning* candidate.

Strategy-Proofness

Notation: (\mathbf{R}_{-i}, R'_i) is the profile obtained by replacing R_i in \mathbf{R} by R'_i .

F is *strategy-proof* (or *immune to manipulation*) if for no individual $i \in \mathcal{N}$ there exist a profile \mathbf{R} (including the “truthful preference” R_i of i) and a linear order R'_i (representing the “untruthful” ballot of i) such that $F(\mathbf{R}_{-i}, R'_i)$ is ranked above $F(\mathbf{R})$ according to R_i .

In other words: under a strategy-proof voting rule no voter will ever have an incentive to misrepresent her preferences.

The Gibbard-Satterthwaite Theorem

Two more properties of resolute voting rules F :

- F is *surjective* if for any candidate $x \in \mathcal{X}$ there exists a profile \mathbf{R} such that $F(\mathbf{R}) = x$.
- F is a *dictatorship* if there exists a voter $i \in \mathcal{N}$ (the dictator) such that $F(\mathbf{R}) = \text{top}(R_i)$ for any profile \mathbf{R} .

Gibbard (1973) and Satterthwaite (1975) independently proved:

Theorem 1 (Gibbard-Satterthwaite) *Any resolute voting rule for ≥ 3 candidates that is *surjective* and *strategy-proof* is a *dictatorship*.*

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.

Remarks

The G-S Theorem says that for ≥ 3 candidates, any resolute voting rule F that is *surjective* and *strategy-proof* is a *dictatorship*.

- a *surprising* result + not applicable in case of *two* candidates
- The opposite direction is clear: *dictatorial* \Rightarrow *strategy-proof*
- *Random* procedures don't count (but might be "strategy-proof").

We will now prove the theorem under two additional assumptions:

- F is *neutral*, i.e., candidates are treated symmetrically.
[Note: neutrality \Rightarrow surjectivity; so we won't make use of surjectivity.]
- There are *exactly 3 candidates*.

For a full proof, using a similar approach, see, e.g.:

U. Endriss. Logic and Social Choice Theory. In A. Gupta and J. van Benthem (eds.), *Logic and Philosophy Today*, College Publications, 2011.

Proof (1)

Notation: $N_{x \succ y}^{\mathbf{R}}$ is the set of voters who rank x above y in profile \mathbf{R} .

Claim: If $F(\mathbf{R}) = x$ and $N_{x \succ y}^{\mathbf{R}} = N_{x \succ y}^{\mathbf{R}'}$, then $F(\mathbf{R}') \neq y$.

Proof: From *strategy-proofness*, by contradiction. Assume $F(\mathbf{R}') = y$. Moving from \mathbf{R} to \mathbf{R}' , there must be a *first* voter to affect the winner. So w.l.o.g., assume \mathbf{R} and \mathbf{R}' differ only wrt. voter i . Two cases:

- $i \in N_{x \succ y}^{\mathbf{R}}$: Suppose i 's true preferences are as in profile \mathbf{R}' (i.e., i prefers x to y). Then i has an incentive to vote as in \mathbf{R} . ✓
- $i \notin N_{x \succ y}^{\mathbf{R}}$: Suppose i 's true preferences are as in profile \mathbf{R} (i.e., i prefers y to x). Then i has an incentive to vote as in \mathbf{R}' . ✓

Some more terminology:

Call $C \subseteq \mathcal{N}$ a *blocking coalition* for (x, y) if $C = N_{x \succ y}^{\mathbf{R}} \Rightarrow F(\mathbf{R}) \neq y$.

Thus: If $F(\mathbf{R}) = x$, then $C := N_{x \succ y}^{\mathbf{R}}$ is blocking for (x, y) [for any y].

Proof (2)

From *neutrality*: all (x, y) must have *the same* blocking coalitions.

For any $C \subseteq \mathcal{N}$, C or $\bar{C} := \mathcal{N} \setminus C$ must be blocking.

Proof: Assume C is not blocking; i.e., C is not blocking for (x, y) .

Then there exists an \mathbf{R} with $N_{x \succ y}^{\mathbf{R}} = C$ but $F(\mathbf{R}) = y$.

But we also have $N_{y \succ x}^{\mathbf{R}} = \bar{C}$. Hence, \bar{C} is blocking for (y, x) .

If C_1 and C_2 are blocking, then so is $C_1 \cap C_2$.

Proof: Consider a profile \mathbf{R} with $C_1 = N_{x \succ y}^{\mathbf{R}}$, $C_2 = N_{y \succ z}^{\mathbf{R}}$, and $C_1 \cap C_2 = N_{x \succ z}^{\mathbf{R}}$. As C_1 is blocking, y cannot win. As C_2 is blocking, z cannot win. So x wins and $C_1 \cap C_2$ must be blocking.

The *empty coalition* is *not* blocking.

Proof: Omitted (but not at all surprising).

Above three properties imply that there must be a *singleton* $\{i\}$ that is blocking. But that just means that i is a *dictator*! ✓

Single-Peakedness

The G-S Thm shows that no “reasonable” voting rule is strategy-proof.

The classical way to circumvent this problem are *domain restrictions*.

The most important domain restriction is due to Black (1948):

- Definition: A profile is *single-peaked* if there exists a “left-to-right” ordering \gg on the candidates such that any voter ranks x above y if x is between y and her top candidate wrt. \gg .
Think of spectrum of political parties.
- Result: Fix a dimension \gg . Assuming that all profiles are single-peaked wrt. \gg , the *median-voter rule* is strategy-proof.

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

Computational Social Choice

Computational Social Choice

Social choice theory studies mechanisms for collective decision making: voting, preference aggregation, fair division, two-sided matching, . . .

- Precursors: Condorcet, Borda (18th century) and others
- serious scientific discipline since 1950s

Computational social choice adds a computational perspective to this, and also explores the use of concepts from social choice in computing.

- “classical” papers: ~1990 (Bartholdi et al.)
- active research area with regular contributions since ~2002
- name “COMSOC” and biannual workshop since 2006

Next: three examples for research directions in COMSOC

Complexity as a Barrier against Manipulation

By the Gibbard-Satterthwaite Theorem, any voting rule for ≥ 3 candidates can be manipulated (unless it is dictatorial).

Idea: So it's always *possible* to manipulate, but maybe it's *difficult*!

Tools from *complexity theory* can be used to make this idea precise.

- For *some* procedures this does *not* work: if I know all other ballots and want X to win, it is *easy* to compute my best strategy.
- But for *others* it does work: manipulation is *NP-complete*.

Recent work in COMSOC has expanded on this idea:

- NP is a worst-case notion. What about average complexity?
- Also: complexity of winner determination, control, bribery, ...

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.

P. Faliszewski, E. Hemaspaandra, and L.A. Hemaspaandra. Using Complexity to Protect Elections. *Communications of the ACM*, 55(11):74–82, 2010.

Automated Reasoning for Social Choice Theory

Logic has long been used to formally specify computer systems, facilitating verification of properties. Can we apply this methodology also here? Yes:

- Verification of a (known) proof of the Gibbard-Satterthwaite Theorem in the HOL proof assistant ISABELLE (Nipkow, 2009).
- Fully automated proof of Arrow's Theorem for 3 candidates via a SAT solver or constraint programming (Tang and Lin, 2009).
- Automated search for new impossibility theorems in *ranking sets of objects* using a SAT solver (Geist and E., 2011).

T. Nipkow. Social Choice Theory in HOL. *Journal of Automated Reasoning*, 43(3):289–304, 2009.

P. Tang and F. Lin. Computer-aided Proofs of Arrow's and other Impossibility Theorems. *Artificial Intelligence*, 173(11):1041–1053, 2009.

C. Geist and U. Endriss. Automated Search for Impossibility Theorems in Social Choice Theory: Ranking Sets of Objects. *J. of Artif. Intell. Res.*, 40:143-174, 2011.

Social Choice in Combinatorial Domains

Suppose 13 voters are asked to each vote *yes* or *no* on three issues; and we use the plurality rule for each issue independently:

- 3 voters each vote for YNN, NYN, NNY.
- 1 voter each votes for YYY, YYN, YNY, NYY.
- No voter votes for NNN.

But then NNN wins: 7 out of 13 vote *no* on each issue (*paradox!*).

What to do instead? The number of candidates is *exponential* in the number of issues (e.g., $2^3 = 8$), so even just representing the voters' preferences is a challenge (\rightsquigarrow *knowledge representation*).

S.J. Brams, D.M. Kilgour, and W.S. Zwicker. The Paradox of Multiple Elections. *Social Choice and Welfare*, 15(2):211–236, 1998.

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine*, 29(4):37–46, 2008.

Computational Social Choice

Research can be broadly classified along two dimensions —

The kind of *social choice problem* studied, e.g.:

- electing a winner given individual preferences over candidates
- aggregating individual judgements into a collective verdict
- fairly dividing a cake given individual tastes

The kind of *computational technique* employed, e.g.:

- algorithm design to implement complex mechanisms
- complexity theory to understand limitations
- logical modelling to fully formalise intuitions
- knowledge representation techniques to compactly model problems
- deployment in a multiagent system

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. A Short Introduction to Computational Social Choice. Proc. SOFSEM-2007.

Conclusion

Last Slide

Tried to give an introduction to *voting theory* (\subseteq social choice theory) and to hint at recent development in *computational social choice*.

Main points:

- many different voting rules available
- surprising phenomena require careful formal modelling
- there's scope for new ideas from computer scientists

These slides and more extensive materials from my Amsterdam course on COMSOC are available online

- <http://www.illc.uva.nl/~ulle/teaching/secvote-2012/>
- <http://www.illc.uva.nl/~ulle/teaching/comsoc/>