

Introduction to Computational Social Choice

The Endriss

Institute for Logic, Language and Computation
University of Amsterdam

The Endriss

1

Computational Social Choice

ESSLLI-2008

Voting Rules

- We'll discuss *voting rules* for selecting a *single winner* from a finite set of *candidates*. (The number of candidates is m .)
- A *voter* votes by submitting a *ballot*, e.g.: the name of a single candidate, a ranking of all candidates, or something else.
- A voting rule has to specify what makes a *valid ballot*, and how the preferences expressed via the ballots are to be *aggregated* to produce the election winner.
- All of the voting rules to be discussed allow for the possibility that two or more candidates come out on top (although this is unlikely for large numbers of voters). A complete system would also have to specify how to deal with such ties, but here we are going to ignore the issue of *tie-breaking*.

The Endriss

3

Computational Social Choice

ESSLLI-2008

Plurality with Run-Off

In the *plurality rule with run-off*, first each voter 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.

Used to elect the president in France (and heavily criticised after Le Pen came in second in the first round in 2002).

The Endriss

5

Computational Social Choice

ESSLLI-2008

Anonymity and Neutrality

On the positive side, both variants of the plurality rule (like most other rules) satisfy these two important properties:

- *Anonymity*: A voting rule is anonymous if it treats all *voters* the same — if two voters switch ballots the election outcome does not change.
- *Neutrality*: A voting rule is neutral if it treats all *candidates* the same — if the election winner switches names with some other candidate, then that other candidate will win.

Often the *tie-breaking* rule can be the source of violation of either anonymity (e.g. if one voter has the power to break ties) or neutrality (e.g. if the incumbent wins in case of a tie).

The Endriss

7

Lecture 2

Voting is a central topic in social choice theory. Today we will begin with an introduction to *voting theory*:

- Many different *voting procedures*, such as the plurality rule, the Borda count, approval voting, single transferable vote, ...
- Discussion of *properties* ("axioms") we would like to see satisfied and *paradoxes* generated by those procedures.

We will then highlight some applications of *complexity theory* to voting, e.g. complexity as a barrier against manipulation.

The Endriss

2

Computational Social Choice

ESSLLI-2008

Plurality Rule

Under the *plurality rule* (a.k.a. *simple majority*), each voter submits a ballot showing the name of one of the candidates standing. The candidate receiving the most votes wins.

This is the most widely used voting rule in practice.

- If there are only two candidates, then it is a very good rule. However, for more than two candidate there are some problems:
 - The information on voter preferences other than who their favourite candidate is gets ignored.
 - Encourages voters not to vote for their true favourite, if that candidate is perceived to have little chance of winning.

The Endriss

4

Computational Social Choice

ESSLLI-2008

Monotonicity

We would like a voting rule to satisfy *monotonicity*: if a particular candidate wins and a voter raises that candidate in their ballot, then that candidate should still win.

The *winner-tURNS-loser paradox* shows that plurality with run-off does *not* satisfy monotonicity:

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

B gets eliminated in the first round and C beats A 66:27 in the run-off. But if 4 of the voters from the first group *raise* C to the *top* (i.e. join the second group), then B will win.

The Endriss

6

Computational Social Choice

ESSLLI-2008

May's Theorem

As mentioned before, if there are only two candidates, then the plurality rule is a pretty good rule to use. Specifically:

Theorem 1 (May, 1952) *For two candidates, a voting rule is anonymous, neutral, and monotonic iff it is the plurality rule.*

Remark: In these slides we assume that there are no ties, but May's Theorem also works for an appropriate definition of monotonicity when ties are possible.

K.O. May: A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decisions. *Econometrica*, 20(4):680–684, 1952.

The Endriss

8

Proof Sketch

Clearly, plurality does satisfy all three properties. ✓

Now for the other direction:

For simplicity, assume the number of voters is odd (no ties).

Anonymous and neutrality \leadsto only number of votes matters.

Denote as A the set of voters voting for candidate a and as B those voting for b . Distinguish two cases:

- Whenever $|A| = |B| + 1$ then a wins. Then, by monotonicity, a wins whenever $|A| > |B|$ (that is, we have plurality). ✓
- There exist A, B with $|A| = |B| + 1$ but b wins. Now suppose one a -voter switches to b . By monotonicity, b still wins. But now $|B'| = |A| + 1$, which is symmetric to the earlier situation, so by neutrality a should win \leadsto contradiction. ✓

Ulle Endriss

9

Computational Social Choice

ESSLLI-2008

Positional Scoring Rules

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

Let m be the number of candidates. A *positional scoring rule* is given by a *scoring vector* $s = (s_1, \dots, s_m)$ with $s_1 \geq s_2 \geq \dots \geq s_m$.

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

The candidate with the highest score (sum of points) wins.

- The *Borda rule* is is the positional scoring rule with the scoring vector $(m-1, m-2, \dots, 0)$.
- The *plurality rule* is the positional scoring rule with the scoring vector $(1, 0, \dots, 0)$.

Ulle Endriss

11

Computational Social Choice

ESSLLI-2008

Positional Scoring violates 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* assigning strictly more points to a candidate placed 2nd than to a candidate placed 3rd ($s_2 > s_3$) makes B win:

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

This shows that *no positional scoring rule* (with a strictly descending scoring vector) will satisfy the *Condorcet principle*.

Ulle Endriss

13

Computational Social Choice

ESSLLI-2008

Approval Voting

In *approval voting*, a ballot may consist of any subset of the set of candidates. These are the candidates the voter approves of. The candidate receiving the most approvals wins.

Intuitive advantages of approval voting include:

- No need *not* to vote for a preferred candidate for strategic reasons, when that candidate has a slim chance of winning.
- Form of balloting seems like a good compromise between plurality (too simple) and Borda (too complex).

Approval voting has been used by several professional societies, such as the American Mathematical Society (AMS).

S.J. Brams and P.C. Fishburn. Approval Voting. *The American Political Science Review* 72(3):831-847, 1978.

Ulle Endriss

15

Borda Rule

Under the voting rule proposed by Jean-Charles de Borda, each voter submits a complete ranking of all the m candidates.

For each voter that places a candidate first, that candidate receives $m-1$ points, for each voter that places her 2nd she receives $m-2$ points, and so forth. The *Borda count* is the sum of all the points.

The candidate with the highest Borda count wins.

J.-C. de Borda. *Mémoire sur les élections au scrutin*. Histoire de l'Académie Royale des Sciences, Paris, 1781.

Ulle Endriss

10

Computational Social Choice

ESSLLI-2008

Condorcet Principle

Recall the Condorcet paradox:

Voter 1: $A \succ B \succ C$
 Voter 2: $B \succ C \succ A$
 Voter 3: $C \succ A \succ B$

A majority prefers A over B and a majority also prefers B over C , but then again a majority prefers C over A . Hence, no single candidate would beat any other candidate in pairwise contests.

In cases where there is such a candidate beating everyone else in a pairwise majority contest, we call her the *Condorcet winner*.

Observe that if there is a Condorcet winner, then it must be unique.

A voting rule satisfies the *Condorcet principle* if it elects the Condorcet winner whenever there is one.

Ulle Endriss

12

Computational Social Choice

ESSLLI-2008

Copeland Rule

The *Copeland rule* is defined as follows:

- Compute the *Copeland score* of each candidate C by awarding 1 point to C for every pairwise majority contest won and $\frac{1}{2}$ points for every draw.
- The candidate with the highest Copeland score is the winner.

Clearly, Copeland does satisfy the Condorcet principle.

Many more such *Condorcet-consistent* voting rules have been proposed in the literature, each taking a different slant on what it means to be the candidate closest to being a Condorcet winner.

A.H. Copeland. *A Reasonable Social Welfare Function*. Seminar on Mathematics in Social Sciences, University of Michigan, 1951.

Ulle Endriss

14

Computational Social Choice

ESSLLI-2008

Single Transferable Vote (STV)

Also known as the *Hare system*. To select a single winner, voters rank all candidates, and we repeat until there is a winner:

- If one of the candidates is the 1st choice for over 50% of the voters (*quota*), she wins.
- Otherwise, the candidate who is ranked 1st by the fewest voters (the *plurality loser*) gets *eliminated* from the race.
- Votes for eliminated candidates get *transferred*: delete removed candidates from ballots and “shift” rankings (e.g. if your 1st choice got eliminated, then your 2nd choice becomes 1st).

In practice, voters need not be required to rank all candidates (non-ranked candidates are assumed to be ranked lowest).

STV is used in several countries (e.g. Australia, New Zealand, ...).

Th. Hare. *The Machinery of Representation*. 1857.

Ulle Endriss

16

Manipulation: Plurality Rule

Suppose the *plurality rule* (as in most real-world situations) is used to decide the outcome of an election. Recall the Florida situation:

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

Bush will win the plurality contest.

It would have been in the interest of the Nader supporters to *manipulate*, i.e. to misrepresent their preferences.

 Title Endriss

17

 Computational Social Choice

ESSSLLI-2008

Setting and Notation

- Finite set A of *candidates* (alternatives); finite set $I = \{1, \dots, n\}$ of *voters* (individuals).
- A *preference ordering* is a strict linear order on A . The set of all such orderings is denoted \mathcal{P} . Each voter i has an *individual preference ordering* P_i . A *preference profile* $(P_1, \dots, P_n) \in \mathcal{P}^n$ consists of a preference ordering for each voter.
- The *top candidate* $\text{top}(P)$ of a preference ordering P is defined as the unique $x \in A$ such that xPy for all $y \in A \setminus \{x\}$.
- We write (P_{-i}, P') for the preference profile we obtain when we *replace* P_i by P' in the preference profile P .
- A *voting rule* is a function $f : \mathcal{P}^n \rightarrow A$ mapping preference profiles to winning candidates (so the P_i are used as ballots).

 Title Endriss

19

 Computational Social Choice

ESSSLLI-2008

Complexity as a Barrier against Manipulation

The Gibbard-Satterthwaite Theorem shows that manipulation is always possible. But how hard is it to find a manipulating ballot? 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 voting* and for the *Borda count*.
- 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.

 Title Endriss

21

 Computational Social Choice

ESSSLLI-2008

Manipulating the Plurality Rule

Recall the plurality rule:

- Each voter submits a ballot showing the name of one of the candidates. The candidate receiving the most votes wins.

The plurality rule is easy to manipulate (trivial):

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

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

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

 Title Endriss

23

The Gibbard-Satterthwaite Theorem

The Gibbard-Satterthwaite Theorem is widely regarded as the central result in voting theory. Broadly, it states that there can be no “reasonable” voting rule that would not be manipulable.

Our formal statement of the theorem follows Barberà (1983). We won’t prove it here. A proof that is similar to the one we have discussed for Arrow’s Theorem is given by Benoit (2000).

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

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

S. Barberà: *Strategy-proofness and Proximal Voters: A Direct Proof of the Gibbard-Satterthwaite Theorem*. *Intl. Economic Review*, 24(2):419–417, 1983.

J.-P. Benoit: *The Gibbard-Satterthwaite Theorem: A Simple Proof*. *Economic Letters*, 69:319–322, 2000.

 Title Endriss

18

 Computational Social Choice

ESSSLLI-2008

Statement of the Theorem

A voting rule f is *dictatorial* if the winner is always the top candidate of a particular voter (the dictator):

$$(\exists i \in I)(\forall \mathbf{P} \in \mathcal{P}^n)[f(\mathbf{P}) = \text{top}(P_i)]$$

A voting rule f is *manipulable* if it may give a voter an incentive to misrepresent their preferences:

$$(\exists \mathbf{P} \in \mathcal{P}^n)(\exists P' \in \mathcal{P})(\exists i \in I)[f(P_{-i}, P') P_i f(\mathbf{P})]$$

A voting rule that is not manipulable is also called *strategy-proof*.

Theorem 2 (Gibbard-Satterthwaite) *If $|A| > 2$, then every voting rule must be either dictatorial or manipulable.*

 Title Endriss

20

 Computational Social Choice

ESSSLLI-2008

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; candidate C .

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

We will be interested in the computational complexity of this problem in terms of the *number of candidates*.

If the $\text{MANIPULABILITY}(f)$ is computationally intractable, then manipulability may be considered less of a worry for voting rule f .

 Title Endriss

22

 Computational Social Choice

ESSSLLI-2008

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 C (the candidate to be made winner through manipulation) at the top of your declared preference ordering.
- Then inductively proceed as follows: Check if any of the remaining candidates can be put next into the preference ordering without preventing C from winning. If yes, do so. If no, terminate and say that manipulation is impossible.

After convincing ourselves that this algorithm is indeed correct, we also get $\text{MANIPULABILITY}(\text{Borda}) \in P$.

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.

 Title Endriss

24

Intractability of Manipulating STV

Recall STV: eliminate plurality losers until a candidate gets $> 50\%$

Theorem 3 (Bartholdi and Orlin, 1991) *Manipulation of STV for electing a single winner is NP-complete.*

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

- NP-membership is clear: checking whether a given ballot makes C win can be done in polynomial time.
 - NP-hardness: Bartholdi and Orlin (1991) give a reduction from 3-COVER. 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.

Ulle Endriss

25

 Computational Social Choice

ESSLI-2008

Even More on Complexity of Voting

- What is the *communication complexity* of different voting rules, i.e. how much information needs to be exchanged to determine the winner of an election? See Conitzer and Sandholm (2005).
- After having counted part of the vote, can we *compile* this information into a more compact form than just storing all the ballots? And how complex is it to reason about this information? See Chevaleyre et al. (2008).

V. Conitzer and T. Sandholm. *Communication Complexity of Common Voting Rules*. Proc. ACM Conference on Electronic Commerce 2005.

Y. Chevaleyre, J. Lang, N. Maudet, and G. Ravilly-Abadie. *Compiling the Votes of a Subelector*. Proc. CONISOC-2008.

Ulle Endriss

27

 Computational Social Choice

ESSLI-2008

Literature

For a definition of the voting procedures introduced (and many more), their properties and the paradoxes they generate, see:

- S.J. Brams and P.C. Fishburn. *Voting Procedures*. In K.J. Arrow et al. (eds.), *Handbook of Social Choice and Welfare*. Elsevier, 2002.

They also briefly cover the Gibbard-Satterthwaite Theorem and May's Theorem (more details are available in several textbooks).

For a nice introduction to work on (computational) complexity in voting, refer to this survey:

- P. Faliszewski, E. Hemaspaandra, L.A. Hemaspaandra, and J. Rothe. *A Richer Understanding of the Complexity of Election Systems*. Technical Report TR-2006-903, Dept. of Computer Science, University of Rochester, 2006.

Ulle Endriss

29

More on Complexity of Voting

Other questions that have been investigated include:

- What is the complexity of other forms of election manipulation, such as *bribery* ? See Faliszewski et al. (2006) for a survey.
 - What is the complexity of the *winner determination* problem? For *Dodgson's rule* (electing the candidate requiring the fewest “flips” in ballots to become a Condorcet winner) it is NP-hard (and not in NP). See Faliszewski et al. (2006) for references.
 - After some of the ballots have been counted, certain candidates may be *possible winners* or even *necessary winners* . How hard is it to check this? See e.g. Konczak and Lang (2005).
- P. Faliszewski, E. Hemaspaandra, L.A. Hemaspaandra, and J. Rothe. *A Richer Understanding of the Complexity of Election Systems*. Technical Report TR-2006-903, Dept. of Computer Science, University of Rochester, 2006.
- K. Konczak and J. Lang. *Voting Procedures with Incomplete Preferences*. Proc. Advances in Preference Handling 2005.

Ulle Endriss

26

 Computational Social Choice

ESSLI-2008

Summary

We have given an introduction to voting theory and seen several *voting procedures* and discussed their *properties* . Specifically:

- May's Theorem
- Gibbard-Satterthwaite Theorem

We have also seen that complexity theory offers an interesting perspective on voting procedures:

- Complexity can serve as a barrier against manipulation. But beware: this is only a *worst-case result* . Manipulation may well be easy on average (ongoing discussion).
- We have also mentioned other forms of control, winner determination, communication complexity, ...

Ulle Endriss

28

 Computational Social Choice

ESSLI-2008

Thursday Morning Quiz

If a candidate loses to every other candidates in pairwise majority contests, then that candidate is called the *Condorcet loser* .

A voting rule satisfies the *Condorcet loser principle* if it never elects a Condorcet loser.

Which of the following rules satisfies this principle, and why?

- *Plurality* : elect the candidate ranked first most often
- *STV* : eliminate plurality losers until someone gets $> 50\%$
- *Borda* : positional scoring rule with vector $(m-1, m-2, \dots, 0)$

Ulle Endriss

30