

Tutorial on Voting Theory

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

Table of Contents

Introduction	3
Voting Procedures and their Properties	8
Major Theorems in Voting Theory	29
Voting Theory and Computational Social Choice	59
Conclusion	83

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 *alternatives*. Each voter submits a *ballot*, based on which a *voting procedure* selects a (set of) *winner(s)* from amongst the alternatives.

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

Voting Theory and AI (1)

Voting theory has natural applications in AI:

- *Search Engines*: to determine the most important sites based on links (“votes”) + to aggregate the output of several search engines
- *Recommender Systems*: to recommend a product to a user based on earlier ratings by other users
- *Multiagent Systems*: to coordinate the actions of groups of autonomous software agents
- *AI Competitions*: to determine who has developed *the best* trading agent / SAT solver / RoboCup team

But not all of the classical assumptions will fit these new applications. So AI needs to develop *new models of voting* and *ask new questions*.

Voting Theory and AI (2)

Vice versa, techniques from AI, and computational techniques in general, are useful for advancing the state of the art in voting theory:

- *Algorithms and Complexity*: to develop algorithms for (complex) voting procedures + to understand the hardness of “using” them
- *Knowledge Representation*: to compactly represent the preferences of voters over large spaces of alternatives
- *Logic and Automated Reasoning*: to formally model problems in voting theory + to automatically verify (or discover) theorems

Indeed, you will find many papers on voting at AI conferences (e.g., IJCAI, AAAI, AAMAS) and many AI researchers participate in events dedicated to voting and social choice (particularly COMSOC).

Tutorial Overview

- Voting Procedures
 - Such as: Plurality, Borda, Approval, STV, Kemeny, . . .
 - Properties and Paradoxes
- Major Theorems in Voting Theory
 - The Axiomatic Method in Voting Theory
 - Characterisations: May's Theorem, Young's Theorem
 - Impossibilities: Arrow's Theorem, Gibbard-Satterthwaite Thm
 - Ways out: Black's Theorem
- Computational Social Choice
 - Introduction to the field
 - Examples for work involving voting

Voting Procedures and their Properties

Voting Procedures

We'll discuss procedures for n voters (or *individuals*, *agents*, *players*) to collectively choose from a set of m alternatives (or *candidates*):

- Each voter votes by submitting a *ballot*, e.g., the name of a single alternative, a ranking of all alternatives, or something else.
- The procedure defines what are *valid ballots*, and how to *aggregate* the ballot information to obtain a winner.

Remark 1: There could be *ties*. So our voting procedures will actually produce *sets of winners*. Tie-breaking is a separate issue.

Remark 2: Formally, *voting rules* (or *resolute* voting procedures) return single winners; *voting correspondences* return sets of winners.

Plurality Rule

Under the *plurality rule* each voter submits a ballot showing the name of one alternative. The alternative(s) receiving the most votes win(s).

Remarks:

- Also known as the *simple majority* rule (\neq *absolute majority* rule).
- This is the most widely used voting procedure in practice.
- If there are only two alternatives, then it is a very good procedure.

Criticism of the Plurality Rule

Problems with the plurality rule (for more than two alternatives):

- The information on voter preferences other than who their favourite candidate is gets ignored.
- Dispersion of votes across ideologically similar candidates.
- Encourages voters not to vote for their true favourite, if that candidate is perceived to have little chance of winning.

Plurality with Run-Off

Under the *plurality rule with run-off*, each voter initially votes for one alternative. The winner is elected in a second round by using the plurality rule with the two top alternatives from the first round.

Remarks:

- Used to elect the president in France.
- Addresses some of the noted problems: elicits more information from voters; realistic “second best” candidate gets another chance.
- Still: heavily criticised after Le Pen entered the run-off in 2002.

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.

Borda Rule

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

Remarks:

- Takes care of some of the problems identified for plurality voting, e.g., this form of balloting is more informative.
- Disadvantage (of any system where voters submit full rankings): higher elicitation and communication costs

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

Example

Consider again this 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

Our voting procedures give different winners:

- Plurality: Bush wins
- Plurality with run-off: Gore wins (Nader eliminated in round 1)
- Borda: Gore wins ($49 + 40 + 40 + 11 > 98 + 20 > 20 + 22$)
- Gore is also the *Condorcet winner* (wins any pairwise contest).

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 alternatives. Each alternative receives s_i points for every voter putting it at the i th position.

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

Remarks:

- 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 Condorcet Principle

An alternative that beats every other alternative 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 procedure 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* assigning strictly more points to a candidate placed 2nd than to a candidate placed 3rd ($s_2 > s_3$) makes B win:

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

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

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

Thus, *no positional scoring rule* (with a strictly descending scoring vector) will satisfy the *Condorcet principle*.

Copeland Rule

Some voting procedures (with ballots that are full rankings) have been designed specifically to meet the Condorcet principle.

The *Copeland rule* elects the alternative(s) that maximise(s) the difference between won and lost pairwise majority contests.

Remarks:

- The Copeland rule satisfies the Condorcet principle.
- Variations are possible: 0 points for every lost contest; 1 point for every won contest; α points (with possibly $\alpha \neq \frac{1}{2}$) for every draw

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

Tournament Solutions

The Copeland rule is an example for a *tournament solution*. There is an entire class of voting procedure that can be defined like this:

- ▶ Draw a directed graph where the alternatives are the vertices and there is an edge from A to B iff A beats B in a majority contest.

Many rules can be defined on such a *majority graph* (Laslier, 1997).

J.F. Laslier. *Tournament Solutions and Majority Voting*. Studies in Economic Theory, Springer-Verlag, 1997.

Kemeny Rule

Under the *Kemeny rule*, ballots are full rankings of the alternatives. An alternative wins if it is maximal in a ranking minimising the sum of disagreements with the ballots regarding pairs of alternatives.

That is:

- (1) For every possible ranking R , count the number of triples (i, x, y) s.t. R disagrees with voter i on the ranking of alternatives x and y .
- (2) Find all rankings R that have minimal score in the above sense.
- (3) Elect any alternative that is maximal in such a “closest” ranking.

Remarks:


- Satisfies the Condorcet principle.
- This will be hard to compute (more later).

J. Kemeny. Mathematics without Numbers. *Daedalus*, 88:571–591, 1959.

Voting Trees (Cup Rule, Sequential Majority)

If ballots are rankings, we can define a voting rule via a *binary tree*, with the alternatives labelling the leaves, and an alternative progressing to a parent node if it beats its sibling in a *majority contest*.

Two examples for such rules and a possible profile of ballots:



Remarks:

- Any such rule satisfies the Condorcet principle.
- Most such rules violate *neutrality* (= symmetry wrt. alternatives).

Single Transferable Vote (STV)

Also known as the *Hare system*. To select a single winner, it works as follows (voters submit ranked preferences for all candidates):

- 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 gets *eliminated* from the race.
- Votes for eliminated candidates get *transferred*: delete removed candidates from ballots and “shift” rankings (i.e., 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 (suitably generalised) is often used to elect committees.

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

Example

Elect one winner amongst four candidates, using STV (100 voters):

39 voters: $A \succ B \succ C \succ D$

20 voters: $B \succ A \succ C \succ D$

20 voters: $B \succ C \succ A \succ D$

11 voters: $C \succ B \succ A \succ D$

10 voters: $D \succ A \succ B \succ C$

(Answer: B wins)

Note that for 3 candidates, STV reduces to plurality voting with run-off, so it suffers from the same problems.

Approval Voting (AV)

In *approval voting*, a ballot is a *set* of alternatives (the ones the voter “approves” of). The alternative with the most approvals wins.

Remarks:

- Approval voting has been used by several professional societies, such as the American Mathematical Society (AMS).
- Intuitively, less cause *not* to vote for the most preferred candidate for strategic reasons when she has a slim chance of winning.
- Good compromise between plurality (too simple) and Borda (too complex) in terms of communication requirements.
- Only procedure we have seen where ballots *cannot* be modelled as linear orders over the set of alternatives.

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

Summary: Voting Procedures

We have seen a fair number of voting *procedures*:

- *Ballots* might be elements (plurality), rankings (e.g., Borda), or subsets (approval) of the set of alternatives. (Enough for AI?)
- Types of procedures:
 - *positional scoring rules*: Borda, (plurality)
 - based on the *majority graph*: Copeland, voting trees
 - based on the *weighted majority graph*: Kemeny
 - *staged procedures*: plurality with run-off, STV
 - *approval voting*

We have seen a few *properties* of voting procedures:

- *Monotonicity*, as violated by e.g. the *no-show paradox*
- *Strategic* issues, meaning people might not vote truthfully
- *Condorcet principle*: if an alternative wins all pairwise majority contests, then it should win the election

Other Voting Procedures

There are many more voting procedures, including these:

- ▶ the Antiplurality (or Veto) Rule, the Banks Rule, Black's Rule, Bucklin Voting, the Condorcet Rule, Coombs' Method, (General) Cumulative Voting, the Dodgson Rule, Even-and-Equal Cumulative Voting, Majority Judgment, Maximin Voting, Range Voting, the Slater Rule, and Tideman's Procedure.

Literature

Most textbooks on Social Choice Theory (some to be cited later) will introduce at least a small number of voting procedures.

This first part of the tutorial owes much to the handbook chapter of Brams and Fishburn (2002), who discuss a good number of voting procedures in detail (with a certain emphasis on approval voting).

Nurmi (1987) devotes an entire book to the analysis of the properties of different voting procedures.

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

H. Nurmi. *Comparing Voting Systems*. Kluwer Academic Publishers, 1987.

Major Theorems in Voting Theory

The Axiomatic Method

Most of the important classical results in voting theory are *axiomatic*. They formalise desirable properties as “axioms” and then establish:

- *Characterisation Theorems*, showing that a particular (class of) procedure(s) is the only one satisfying a given set of axioms
- *Impossibility Theorems*, showing that there exists *no* voting procedure satisfying a given set of axioms

We will see two examples each (+ one other thing).

Formal Framework

Basic terminology and notation:

- finite set of *voters* $\mathcal{N} = \{1, \dots, n\}$, the *electorate*
- (usually finite) set of *alternatives* $\mathcal{X} = \{x_1, x_2, x_3, \dots\}$
- Denote the set of *linear orders* on \mathcal{X} by $\mathcal{L}(\mathcal{X})$. *Preferences* are assumed to be elements of $\mathcal{L}(\mathcal{X})$. *Ballots* are elements of $\mathcal{L}(\mathcal{X})$.

A *voting procedure* is a function $F : \mathcal{L}(\mathcal{X})^n \rightarrow 2^{\mathcal{X}} \setminus \{\emptyset\}$, mapping profiles of ballots to nonempty sets of alternatives.

Remark: AV does not fit in this framework; everything else does.

Two Alternatives

When there are only *two alternatives*, then all the voting procedures we have seen coincide, and *intuitively* they do the “right” thing.

Can we make this intuition precise?

- ▶ Yes, using the axiomatic method.

Anonymity

A voting rule is *anonymous* if the *voters* are treated symmetrically: if two voters switch ballots, then the winners don't change.

Formally:

F is anonymous if $F(b_1, \dots, b_n) = F(b_{\pi(1)}, \dots, b_{\pi(n)})$ for any ballot profile (b_1, \dots, b_n) and any permutation $\pi : \mathcal{N} \rightarrow \mathcal{N}$.

Neutrality

A voting procedure is *neutral* if the *alternatives* are treated symmetrically.

Formally:

F is neutral if $F(\pi(\underline{b})) = \pi(F(\underline{b}))$ for any ballot profile \underline{b} and any permutation $\pi : \mathcal{X} \rightarrow \mathcal{X}$ (with π extended to ballot profiles and sets of alternatives in the natural manner).

Positive Responsiveness

A voting procedure satisfies the property of *positive responsiveness* if, whenever some voter raises a (possibly tied) winner x in her ballot, then x will become the unique winner.

Formally:

F satisfies positive responsiveness if $x \in F(\underline{b})$ implies $\{x\} = F(\underline{b}')$ for any alternative x and any two distinct profiles \underline{b} and \underline{b}' with $\underline{b}(x \succ y) \subseteq \underline{b}'(x \succ y)$ and $\underline{b}(y \succ z) = \underline{b}'(y \succ z)$ for all alternative y and z different from x .

Notation: $\underline{b}(x \succ y)$ is the set of voters ranking x above y in \underline{b}

May's Theorem

Now we can fully characterise the plurality rule:

Theorem 1 (May, 1952) *A voting procedure for **two** alternatives satisfies **anonymity**, **neutrality**, and **positive responsiveness** if and only if it is the **plurality rule**.*

Remark: In these slides we assume that there are no indifferences in ballots, but May's Theorem also works (with an appropriate definition of positive responsiveness) when ballots are weak orders.

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

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).

Plurality-style ballots are fully expressive for two alternatives.

Anonymity and neutrality \rightsquigarrow only number of votes matters.

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

- Whenever $|A| = |B| + 1$ then only a wins. Then, by PR, 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 PR, now only b wins. But now $|B'| = |A'| + 1$, which is symmetric to the earlier situation, so by neutrality a should win \rightsquigarrow contradiction. ✓

Characterisation Theorems

When there are more than two alternatives, then different voting procedures are really different. To choose one, we need to understand its properties: ideally, we get a *characterisation theorem*.

Maybe the best known result of this kind is Young's characterisation of the *positional scoring rules* (PSR) ...

Reminder:

- ▶ Every scoring vector $s = \langle s_1, \dots, s_m \rangle$ with $s_1 \geq s_2 \geq \dots \geq s_m$ and $s_1 > s_m$ defines a PSR: give s_i points to alternative x whenever someone ranks x at the i th position; the winners are the alternatives with the most points.

Reinforcement (a.k.a. Consistency)

A voting procedure satisfies *reinforcement* if, whenever we split the electorate into two groups and some alternative would win in both groups, then it will also win for the full electorate.

For a full formalisation of this concept we would need to be able to speak about a voting procedure F wrt. different electorates $\mathcal{N}, \mathcal{N}', \dots$

Formally (under natural refinements to our notation):

F satisfies reinforcement if $F^{\mathcal{N} \cup \mathcal{N}'}(\underline{b}) = F^{\mathcal{N}}(\underline{b}) \cap F^{\mathcal{N}'}(\underline{b})$ for any disjoint electorates \mathcal{N} and \mathcal{N}' and any ballot profile \underline{b} such that $F^{\mathcal{N}}(\underline{b}) \cap F^{\mathcal{N}'}(\underline{b}) \neq \emptyset$.

Continuity

A voting procedure is *continuous* if, whenever electorate \mathcal{N} elects a unique winner x , then for any other electorate \mathcal{N}' there exists a number k s.t. \mathcal{N}' together with k copies of \mathcal{N} will also elect only x .

Young's Theorem

We are now ready to state the theorem:

Theorem 2 (Young, 1975) *A voting procedure satisfies anonymity, neutrality, reinforcement, and continuity iff it is a positional scoring rule.*

Proof: Omitted (and difficult).

But it is not hard to verify the right-to-left direction.

H.P. Young. Social Choice Scoring Functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.

Impossibility Theorems

Another important type of result are *impossibility theorems*:

- showing that a certain combination of axioms is *inconsistent*
- alternative reading: a certain set of axioms *characterises* an obviously unattractive rule (directly violating a final axiom)

We first discuss Arrow' Theorem ...

Unanimity and the Pareto Condition

A voting procedure is *unanimous* if it elects only x whenever all voters say that x is the best alternative. Formally:

F is unanimous if whenever $\underline{b}(x \succ y) = \mathcal{N}$ for all $y \in \mathcal{N} \setminus \{x\}$ then $F(\underline{b}) = \{x\}$.

The *weak Pareto condition* is slightly less demanding. It is satisfied if an alternative y that is dominated by some other alternative x in all ballots cannot win. Formally:

F is weakly Pareto if $\underline{b}(x \succ y) = \mathcal{N}$ implies $y \notin F(\underline{b})$.

Independence of Irrelevant Alternatives (IIA)

A voting procedure is *irrelevant of independent alternatives* if, whenever y loses to some winner x and the relative ranking of x and y does not change in the ballots, then y cannot win (independently of any possible changes wrt. other, irrelevant, alternatives).

Formally:

F satisfies IIA if $x \in F(\underline{b})$ and $y \notin F(\underline{b})$ together with $\underline{b}(x \succ y) = \underline{b}'(x \succ y)$ imply $y \notin F(\underline{b}')$ for any profiles \underline{b} and \underline{b}' .

Remark: This variant of IIA (for voting rules) is due to Taylor (2005). Arrow's original formulation of IIA is for *social welfare functions*, where the outcome is a preference ordering.

A.D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005.

K.J. Arrow. *Social Choice and Individual Values*. 2nd edition. Cowles Foundation, Yale University Press, 1963.

Dictatorships

- A voting procedure is a *dictatorship* if there exists a voter such that the unique winner will always be the top-ranked alternative of that voter (the dictator).
- A voting procedure is *nondictatorial* if it is not a dictatorship.

Arrow's Theorem for Voting Procedures

This is widely regarded as *the* seminal result in Social Choice Theory. Kenneth J. Arrow received the Nobel Prize in Economics in 1972.

Theorem 3 (Arrow, 1951) *No voting procedure for ≥ 3 alternatives is weakly Pareto, IIA, and nondictatorial.*

Proof: Omitted.

This particular version of the theorem is proved by Taylor (2005).

Maybe the most accessible proof (of the standard formulation of the theorem) is the first proof in the paper by Geanakoplos (2005).

K.J. Arrow. *Social Choice and Individual Values*. 2nd edition. Cowles Foundation, Yale University Press, 1963.

A.D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005.

J. Geanakoplos. Three Brief Proofs of Arrow's Impossibility Theorem. *Economic Theory*, 26(1):211–215, 2005.

Remarks

- Note that this is a *surprising* result!
- Note that the theorem does *not* hold for *two* alternatives.
- We can interpret the theorem as a *characterisation* result:
A voting procedure for ≥ 3 alternatives satisfies the weak Pareto condition and IIA if and only if it is a dictatorship.
- *IIA* is the most debatable of the three axioms featuring in the theorem. Indeed, it is quite hard to satisfy.

Manipulation

Let's look once more at our favourite 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, the Nader supporters could *manipulate*: pretend they like Gore best and improve the result.

Ideally, there would be no need for voters to strategise in this way. Ideally, we'd like a procedure that is *strategy-proof*.

Strategy-Proofness

Recall: F is *resolute* if $F(\underline{b})$ is a singleton for any profile of ballots \underline{b} .

Let \succ_i be the true preference of voter i and let b_i be the ballot of i .

A resolute voting procedure is *strategy-proof* if there exist no profile $\underline{b} = (b_1, \dots, b_n)$ and no voter i s.t. $F(\underline{b}) \succ_i F(b_1, \dots, \succ_i, \dots, b_n)$, with \succ_i lifted from alternatives to singletons in the natural manner.

The Gibbard-Satterthwaite Theorem

A resolute voting procedure F is *surjective* if for any alternative x there exists a ballot profile \underline{b} such that $F(\underline{b}) = \{x\}$.

Theorem 4 (Gibbard-Satterthwaite) Any *resolute* voting procedure for ≥ 3 alternatives that is *surjective* and *strategy-proof* is *dictatorial*.

Remarks:

- Again, *surprising*. Again, not applicable for *two* alternatives.
- The opposite direction is clear: *dictatorial* \Rightarrow *strategy-proof*
- *Random* procedures don't count (but might be “strategy-proof”).

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.

Proof Sketch

One way of proving this involves the notion of a *pivotal voter*. Benoît (2000) gives a simple proof based on this idea.

The main steps are:

- show that when all voters rank x last, then x doesn't win
- show that when all voters rank x on top, then x wins
- observe that when we let voters switch x from bottom to top one by one, there must be a *pivotal voter* i causing x to win
- show that i can dictate x 's relative position wrt. any y
- repeat this for every alternative: each has a “local” dictator
- observe that, by definition, there can be only one dictator

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

Domain Restrictions

How can we circumvent these impossibilities?

- 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 5 (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. ✓

Consequences

If the number of voters is odd and their preferences (and ballots) are single-peaked wrt. a known order, then:

- The median voter rule (= electing the Condorcet winner) is *strategy-proof* (Gibbard-Satterthwaite fails).
- The median voter rule (= electing the Condorcet winner) is *weakly Pareto* and *IIA* (Arrow fails).

Summary: Major Theorems

We have seen some of the major theorems in Social Choice Theory pertaining to voting, using the *axiomatic method*:

- *May*: plurality for two alternatives is characterised by anonymity, neutrality and positive responsiveness
- *Young*: positional scoring rules are characterised by reinforcement
- *Arrow*: Pareto (unanimity) and independence lead to dictatorships
- *Gibbard-Satterthwaite*: strategy-proofness leads to dictatorships
- *Black*: single-peakedness solves most problems

Other classics to look out for:

- *McGarvey*: any majority graph can occur
- *Sen*: impossibility of a Paretian liberal
- *Sen*: triple-wise value restriction, generalising single-peakedness
- *Duggan-Schwartz*: G-S for irresolute voting procedures
- *Clarke* and *Groves*: strategy-proofness for quasi-linear preferences

Literature

The handbook edited by Arrow et al. (2002) is the authoritative reference work in Social Choice Theory.

Much more accessible, however, are the excellent textbooks by Gaertner (2009) and Taylor (2005).

Also nice is Part IV of Moulin (1988). This book is particularly good for topics at the interface of SCT and Welfare Economics.

K.J. Arrow, A.K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*. North-Holland, 2002.

W. Gaertner. *A Primer in Social Choice Theory*. Revised edition. LSE Perspectives in Economic Analysis. Oxford University Press, 2009.

A.D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005.

H. Moulin. *Axioms of Cooperative Decision Making*. Econometric Society Monographs. Cambridge University Press, 1988.

Voting Theory and Computational Social Choice

Computational Social Choice

Social choice theory studies mechanisms for collective decision making: voting, preference aggregation, fair division, 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

The COMSOC Research Community

- International Workshop on Computational Social Choice:
 - 1st edition: COMSOC-2006 in Amsterdam, December 2006
48 paper submissions and 80 participants (14 countries)
 - 2nd edition; COMSOC-2008 in Liverpool, September 2008
55 paper submissions and ~80 participants (~20 countries)
 - 3rd edition: COMSOC-2010 in Düsseldorf, September 2010
58 paper submissions
- Special issues in international journals:
 - *Mathematical Logic Quarterly*, vol. 55, no. 4, 2009
 - *Journal of Autonomous Agents and Multiagent Systems* (2010)
 - *Mathematical Social Sciences* (in preparation)
- Journals and conferences in AI, MAS, TCS, Logic, Econ, ...
- COMSOC website: <http://www.illc.uva.nl/COMSOC/>

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.

Examples

During the remainder of the tutorial, we will see some examples of application from methods originating in AI and Computer Science to (new) problems in voting and social choice:

- Automated Reasoning
- Complexity Theory
- Knowledge Representation

Logic and Automated Reasoning

Logic has long been used to *formally specify* computer systems, facilitating formal or even *automatic verification* of various properties.

Can we apply this methodology also to *social choice* mechanisms?

- What logic fits best?
- Which automated reasoning methods are useful?

Computer-aided Proof of Arrow's Theorem

Tang and Lin (2009) prove two inductive lemmas:

- If there exists an Arrovian aggregator for n voters and $m+1$ alternatives, then there exists one for n and m (if $n \geq 2$, $m \geq 3$).
- If there exists an Arrovian aggregator for $n+1$ voters and m alternatives, then there exists one for n and m (if $n \geq 2$, $m \geq 3$).

Tang and Lin then show that the “*base case*” of Arrow's Theorem with 2 agents and 3 alternatives can be fully modelled in *propositional logic*.

A SAT solver can *verify* $\text{ARROW}(2, 3)$ to be correct in < 1 second — that's $(3!)^{3! \times 3!} \approx 10^{28}$ aggregators [SWFs] to check.

Discussion: Opens up opportunities for quick sanity checks of hypotheses regarding new impossibility theorems.

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

Related Work

- Ågotnes et al. (2010) propose a modal logic to model preferences and their aggregation that can express Arrow's Theorem.
- Arrow's Theorem holds *iff* the set T_{ARROW} of FOL formulas (defined in the paper) has no finite models (Grandi and E., 2009).
- Nipkow (2009) formalises and verifies a known *proof* of Arrow's Theorem in the HOL proof assistant ISABELLE.

T. Ågotnes, W. van der Hoek, and M. Wooldridge. On the Logic of Preference and Judgment Aggregation. *J. Auton. Agents and Multiagent Sys.* In press (2010).

U. Grandi and U. Endriss. *First-order Logic Formalisation of Arrow's Theorem.* Proc. 2nd Internat. Workshop on Logic, Rationality and Interaction (LORI-2009).

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

Applications of Complexity Theory

One natural application of Computer Science to voting is to develop *algorithms* for computing the winners of complex voting procedures.

On the theoretical side, people have analysed the *complexity* of the *winner determination problem*.

Example: Checking whether a given alternative wins under the *Kemeny rule* is complete for parallel access to NP.

E. Hemaspaandra, H. Spakowski, and J. Vogel. The Complexity of Kemeny Elections. *Theoretical Computer Science*, 349:382–391, 2005.

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* 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 .

Manipulating the Plurality Rule

Recall the plurality rule:

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

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 declared preference ordering.
- Then inductively proceed as follows: Check if any of the remaining alternatives can be put next into the preference ordering without preventing x 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 $\text{MANIPULABILITY}(\text{Borda}) \in \text{P}$.

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) *Manipulation of 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.
- NP-hardness: by reduction from 3-COVER.

Discussion: NP is a worst-case notion. What about average complexity?

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

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. (2009) for a survey.
- 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*. In *Fundamental Problems in Computing*, Springer-Verlag, 2009.

K. Konczak and J. Lang. *Voting Procedures with Incomplete Preferences*. Proc. Advances in Preference Handling 2005.

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. (2009).

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 Subelectorate*. Proc. IJCAI-2009.

Voting in Combinatorial Domains

Besides the complexity-theoretic properties of voting procedures, another computational concern in voting is raised by the fact that the alternatives to vote for often have a *combinatorial structure*:

- Electing a committee of k members from amongst n candidates.
- During a referendum (in Switzerland, California, places like that), voters may be asked to vote on several propositions.

We will see an example and look into several possible approaches . . .

Based on J. Lang's "5 solutions". Read it in Chevaleyre et al. (2008).

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.

Example

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

- 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.

This is an instance of the *paradox of multiple elections*: the winning combination receives the fewest number of votes.

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

Basic Solution Attempts

- Solution 1: just vote for combinations directly
 - only feasible for *very* small problem instances
 - Example: 3-seat committee, 10 candidates $\rightsquigarrow \binom{10}{3} = 120$
- Solution 2: vote for top k combinations only (e.g., $k = 1$)
 - does address communication problem of Solution 1
 - possibly nobody gets more than one vote (tie-breaking decides)
- Solution 3: make a small preselection of combinations to vote on
 - does solve the computational problems
 - but who should select? (strategic control)

Combinatorial Vote

Idea: Ask voters to report their ballots using a compact preference representation language and apply your favourite voting procedure to the succinctly encoded ballots received.

Lang (2004) calls this approach *combinatorial vote*.

Discussion: This seems the most promising approach so far, although not too much is known to date what would be good choices for preference representation languages or voting procedures, or what algorithms to use to compute the winners. Also, complexity can be expected to be very high.

J. Lang. Logical Preference Representation and Combinatorial Vote. *Annals of Mathematics and Artificial Intelligence*, 42(1–3):37–71, 2004.

Example

Use the language defined by the *leximin ordering over prioritised goals* with the *Borda rule* (goals are labelled by their rank):

- Voter 1: $\{A:0, B:1\}$ induces order $AB \succ_1 A\bar{B} \succ_1 \bar{A}B \succ_1 \bar{A}\bar{B}$
- Voter 2: $\{A \vee \neg B:0\}$ induces order $A\bar{B} \sim_2 AB \sim_2 \bar{A}\bar{B} \succ_2 \bar{A}B$
- Voter 3: $\{\neg A:0, B:0\}$ induces order $\bar{A}B \succ_3 \bar{A}\bar{B} \sim_3 AB \succ_3 A\bar{B}$

As the induced orders need not be strict linear orders, we use a *generalisation of the Borda rule*: a candidate gets as many points as she dominates other candidates. So we get these Borda counts:

$$\begin{array}{ll}
 AB : 3 + 1 + 1 = 5 & \bar{A}B : 1 + 0 + 3 = 4 \\
 A\bar{B} : 2 + 1 + 0 = 3 & \bar{A}\bar{B} : 0 + 1 + 1 = 2
 \end{array}$$

So combination AB wins.

Combinatorial vote *proper* would be to compute the winner *directly* from the goal bases, without the detour via the induced orders.

Other Approaches

Vote on each issue separately *but* —

- identify *conditions* under which this does not lead to undesirable outcomes (“separable preferences”)
- find a *novel way of aggregating* the ballots to select a winner
 - Example: elect the combination minimising the maximal Hamming distance to any of the ballots (Brams et al., 2007)
- vote *sequentially* rather than simultaneously
 - Example: Lang and Xia (2009) use CP-nets to represent ballots and use the underlying graph as an agenda

S.J. Brams, D.M. Kilgour, and M.R. Sanver. A Minimax Procedure for Electing Committees. *Public Choice*, 132:401–420, 2007.

J. Lang and L. Xia. Sequential Composition of Voting Rules in Multi-issue Domains. *Mathematical Social Sciences*, 57(3):304–324, 2009.

Summary: Computational Social Choice

We have seen a small selection of samples of COMSOC research:

- Logic and automated reasoning for *verification* of results in SCT (also interesting: *formalisation*, *discovery*)
- Complexity theory to distinguish possibility from feasibility (for *manipulation*, *winner determination*, and more)
- KR for modelling social choice in *combinatorial domains*

There is a growing COMSOC research community out there, investigating these issues and much more:

- other questions in voting and preference aggregation
- fair division, stable matchings, judgment aggregation, ...

Literature

Chevaleyre et al. (2007) classify contributions in COMSOC wrt. the computational method used and the social choice problem addressed.

Faliszewski and Procaccia (2010) review work on the complexity of manipulation (the archetypical COMSOC problem).

Chevaleyre et al. (2008) give an introduction to social choice in combinatorial domains.

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

P. Faliszewski and A. Procaccia. *AI's War on Manipulation: Are We Winning?* *AI Magazine*. In press (2010).

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.

Conclusion

Last Slide

- We have seen: many voting procedures; classical theorems on voting in SCT; examples for recent work at the interface with AI
- Nice topic, particularly for AI people. Still lots to do.
- A website where you can find out more about Computational Social Choice (workshops, mailing list, PhD theses, etc.):

<http://www.illc.uva.nl/COMSOC/>

- These slides will remain available on the tutorial website, and more extensive materials can be found on the website of my Amsterdam course on Computational Social Choice:
 - <http://www.illc.uva.nl/~ulle/teaching/aaai-2010/>
 - <http://www.illc.uva.nl/~ulle/teaching/comsoc/>