

Tutorial on Voting Theory

Ulle Endriss

Institute for Logic, Language and Computation
University of Amsterdam

[<http://www.illc.uva.nl/~ulle/teaching/easss-2013/>]

Table of Contents

Introduction	3
Voting Rules and their Properties	8
Classical Theorems in Voting Theory	28
Voting Theory and Computational Social Choice	48
Conclusion	55

Introduction

Opening Example

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

Questions:

- Who wins?
- Is that a fair outcome?
- What would your advice to the Nader-supporters have been?

Social Choice and AI (1)

Voting theory (which is part of *social choice theory*, mostly developed in Economics) has a number of natural applications in AI:

- *Multiagent Systems*: to aggregate the beliefs + to coordinate the actions of groups of autonomous software agents
- *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
- *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* and *ask new questions*.

Social Choice and AI (2)

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

- *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 individual agents over large spaces of alternatives
- *Logic and Automated Reasoning*: to formally model problems in social choice + to automatically verify (or discover) theorems

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

F. Brandt, V. Conitzer, and U. Endriss. Computational Social Choice. In G. Weiss (ed.), *Multiagent Systems*. MIT Press, 2013.

Tutorial Outline

The aim of this tutorial is to show you enough voting theory to enable you to appreciate current research on voting in the multiagent systems community. There will be three main topics:

- Voting Rules and their Properties
- Classical Theorems in Voting Theory
- Voting Theory and Computational Social Choice

Our focus will be on the *theory* of voting, which you need to master before it makes sense to contemplate applications (e.g., in MAS).

Voting Rules and their Properties

Three Voting Rules

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

Here are three *voting rules* (there are many more):

- *Plurality*: elect the alternative ranked first most often (i.e., each voter assigns 1 point to an alternative of her choice, and the alternative receiving the most points wins)
- *Plurality with runoff*: run a plurality election and retain the two front-runners; then run a majority contest between them
- *Borda*: each voter gives $m-1$ points to the alternative she ranks first, $m-2$ to the alternative she ranks second, etc.; and the alternative with the most points wins

Example

Consider this election with nine *voters* having to choose from three *alternatives* (namely what drink to order for a common lunch):

4 *Dutchmen*: Milk \succ Beer \succ Wine
2 *Germans*: Beer \succ Wine \succ Milk
3 *Frenchmen*: Wine \succ Beer \succ Milk

Which beverage *wins* the election for

- the plurality rule?
- plurality with runoff?
- the Borda rule?

Formal Framework

Finite set of n *voters* (or *individuals* or *agents*) $\mathcal{N} = \{1, \dots, n\}$.

Finite set of m *alternatives* (or *candidates*) \mathcal{X} .

Each voter expresses a *preference* over the alternatives by providing a linear order on \mathcal{X} (her *ballot*). $\mathcal{L}(\mathcal{X})$ is the set of all such linear orders.

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

A *voting rule* or (*social choice function*) is a function F mapping any given profile to a nonempty set of winning alternatives:

$$F : \mathcal{L}(\mathcal{X})^n \rightarrow 2^{\mathcal{X}} \setminus \{\emptyset\}$$

F is called *resolute* if there is always a unique winner: $|F(\mathbf{R})| \equiv 1$.

Aside: Nonstandard voting methods such as *approval voting* (award 1 point each to any candidate you like!) or *range voting* (distribute 100 points amongst all candidates!) do not fit into this framework.

But *plurality* does: it just ignores everything but the top-ranked alternative.

Single Transferable Vote (STV)

STV is a *staged procedure* that generalises the idea at the core of plurality with runoff. It is often used to elect committees. For a single-winner election, ask each voter to rank all candidates and then:

- If one candidate is the 1st choice for over 50% of the voters (*quota*), then that candidate wins.
- Otherwise, the candidate that is ranked 1st by the fewest voters (the *plurality loser*) gets *eliminated* from the race.
- Votes for eliminated candidates get *transferred*: delete eliminated 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 is used in several countries (e.g., Australia, New Zealand, ...).

Example

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

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

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

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

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

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

Who wins?

The No-Show Paradox

Under plurality with runoff (and thus under STV), 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 runoff.

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 runoff.

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

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(s) with the highest score (sum of points) win(s).

Examples:

- *Borda rule* = PSR with scoring vector $\langle m-1, m-2, \dots, 0 \rangle$
- *Plurality rule* = PSR with scoring vector $\langle 1, 0, \dots, 0 \rangle$
- *Antiplurality rule* = PSR with scoring vector $\langle 1, \dots, 1, 0 \rangle$
- For any $k \leq m$, *k-approval* = PSR with $\langle \underbrace{1, \dots, 1}_k, 0, \dots, 0 \rangle$

Note that k -approval and approval voting are two very different rules!

The Condorcet Principle

The Marquis de Condorcet was a public intellectual working in France during the second half of the 18th century.

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: $B \succ C \succ A$

Cindy: $C \succ A \succ B$

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.

PSR's 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) alternatives will satisfy the *Condorcet principle*.

Copeland Rule

Under the *Copeland rule* each alternative gets +1 point for every won pairwise majority contest and -1 point for every lost pairwise majority contest. The alternative with the most points wins.

Remark 1: The Copeland rule satisfies the Condorcet principle.

Remark 2: All we need to compute the Copeland winner for an election is the *majority graph* (with an edge from alternative A to alternative B if A beats B in a pairwise majority contest).


Exercise: How can you characterise the *Condorcet winner* (if it exists) in graph-theoretical terms in a given *majority graph*?

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

Voting Trees (Cup Rule, Sequential Majority)

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*. (Common assumption: each alternative must show up at least once.)

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



Remarks:

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

The (Weak) Pareto Principle

Vilfredo Pareto was an Italian economist active around 1900.

In economics, an outcome X is called *Pareto efficient* if there is no other outcome Y such that some agents are better off and no agent is worse off when we choose Y rather than X .


Pareto principle: never choose an outcome that is not Pareto efficient.

Weak Pareto principle: never choose an outcome X when there is an other outcome Y strictly preferred by all agents.

Remark: In our context, where all preferences are strict (nobody equally prefers two distinct alternatives), the two principles coincide.

Voting Trees Violate Pareto

Despite being such a weak (and highly desirable) requirement, the (weak) Pareto principle is violated by some rules based on voting trees:



Consider this profile with three agents:

Ann: $A \succ B \succ C \succ D$

Bob: $B \succ C \succ D \succ A$

Cindy: $C \succ D \succ A \succ B$

D wins! (despite being dominated by C)

What happened? To understand the essence of this paradox, note how it is constructed from the Condorcet paradox, with every occurrence of C being replaced by $C \succ D$...

Slater Rule

One more rule that is based on the *majority graph*:

Under the *Slater rule*, we pick a ranking R of the alternatives that minimises the number of edges in the majority graph we have to turn around before we obtain R ; we then elect the top element in R .

(If there is more than one R that minimises the distance to the majority graph, then we get several winners.)

Remark: Computing Slater winners is *NP-hard* (as this is equivalent to the classical MINIMUM FEEDBACK ARC SET problem).

P. Slater. Inconsistencies in a Schedule of Paired Comparisons. *Biometrika*, 48(3–4):303–312, 1961.

Kemeny Rule

Under the *Kemeny rule* 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 (Exercise: why?).
- Knowing the majority graph is *not* enough for this rule.
- Hard to compute: complete for parallel access to NP.

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

E. Hemaspaandra, H. Spakowski, and J. Vogel. The Complexity of Kemeny Elections. *Theoretical Computer Science*, 349(3):382-391, 2005.

Classification of Condorcet Extensions

A *Condorcet extension* is a voting rule that respects the Condorcet principle. Fishburn suggested the following classification:

- *C1*: Rules for which the winners can be computed from the *majority graph* alone. Example:
 - *Copeland*: elect the candidate that maximises the difference between won and lost pairwise majority contests
- *C2*: Non-C1 rules for which the winners can be computed from the *weighted majority graph* alone. Example:
 - *Kemeny*: elect top candidates in rankings that minimise the sum of the weights of the edges we need to flip
- *C3*: All other Condorcet extensions. Example:
 - *Young*: elect candidates that minimise number of voters to be removed before those candidates become Condorcet winners

P.C. Fishburn. Condorcet Social Choice Functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.

Summary: Voting Rules

We have seen a fair number of voting rules, in particular:

- *Positional scoring rules*: Borda, plurality, antiplurality, k -approval
- *Condorcet extensions*:
 - based on the *majority graph*: Copeland, voting trees, Slater
 - based on the *weighted majority graph*: Kemeny
- *Staged procedures*: plurality with runoff, STV
- *Nonstandard methods*: approval voting

We have seen some surprising phenomena (“*paradoxes*”):

- seemingly reasonable rules might all elect different winners
- voting for a candidate might damage her (no-show paradox)
- majority might disagree with the outcome (violation of Condorcet)
- everyone might disagree with the outcome (violation of Pareto)

So: we need some theory to understand all of this better!

Literature: Voting Rules

Most textbooks on Social Choice Theory (some to be cited later on) introduce at least a few voting rules, as does our introductory chapter on Computational Social Choice (Brandt et al., 2013).

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

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

F. Brandt, V. Conitzer, and U. Endriss. Computational Social Choice. In G. Weiss (ed.), *Multiagent Systems*. MIT Press, 2013.

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.

What next?

The rest of the tutorial will be devoted to two main topics:

- Classical Theorems in Voting Theory

Keywords: axiomatic method as a means of characterising voting rules by their properties; strategic behaviour in elections and the Gibbard-Satterthwaite Theorem

- Voting Theory and Computational Social Choice

Keywords: complexity of voting; automated verification of results; compact representation of large voting problems

Classical Theorems in Voting Theory

The Axiomatic Method

We have seen many different voting rules. It is not obvious how to choose the “right” one. We can approach this problem by formulating *axioms* expressing desirable properties (often related to fairness).

Possible results:

- *Characterisation theorems*: certain axioms fully characterise a given voting rule
- *Impossibility theorems*: certain axioms cannot be satisfied together

Reminder: Formal Framework

Finite set of n *voters* (or *individuals* or *agents*) $\mathcal{N} = \{1, \dots, n\}$.

Finite set of m *alternatives* (or *candidates*) \mathcal{X} .

Each voter expresses a *preference* over the alternatives by providing a linear order on \mathcal{X} (her *ballot*). $\mathcal{L}(\mathcal{X})$ is the set of all such linear orders.

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

A *voting rule* or (*social choice function*) is a function F mapping any given profile to a nonempty set of winning alternatives:

$$F : \mathcal{L}(\mathcal{X})^n \rightarrow 2^{\mathcal{X}} \setminus \{\emptyset\}$$

F is called *resolute* if there is always a unique winner: $|F(\mathbf{R})| \equiv 1$.

Anonymity and Neutrality

A voting rule F is *anonymous* if *individuals* are treated symmetrically:

$$F(R_1, \dots, R_n) = F(R_{\pi(1)}, \dots, R_{\pi(n)})$$

for any profile \mathbf{R} and any permutation $\pi : \mathcal{N} \rightarrow \mathcal{N}$

A voting rule F is *neutral* if *alternatives* are treated symmetrically:

$$F(\pi(\mathbf{R})) = \pi(F(\mathbf{R}))$$

for any profile \mathbf{R} and any permutation $\pi : \mathcal{X} \rightarrow \mathcal{X}$

(with π extended to profiles and sets in the natural manner)

Remarks:

- not every voting rule will satisfy every axiom we state here
- axioms are meant to be *desirable* properties (always arguable)

Exercise: show that you cannot get both A and N for *resolute* rules!

Positive Responsiveness

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

A (not necessarily resolute) voting rule satisfies *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 $x^* \in F(\mathbf{R})$ implies $\{x^*\} = F(\mathbf{R}')$ for any alternative x^* and any two *distinct* profiles \mathbf{R} and \mathbf{R}' with $N_{x^* \succ y}^{\mathbf{R}} \subseteq N_{x^* \succ y}^{\mathbf{R}'}$ and $N_{y \succ z}^{\mathbf{R}} = N_{y \succ z}^{\mathbf{R}'}$ for all $y, z \in \mathcal{X} \setminus \{x^*\}$.

May's Theorem

When there are only *two alternatives*, the *plurality rule* is usually called the *simple majority rule*. Intuitively, it does the “right” thing. Can we make this intuition precise? *Yes!*

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

Proof: next slide

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

Proof Sketch

Clearly, simple majority does satisfy all three properties. ✓

Now for the other direction:

Assume the number of voters is *odd* (other case: similar) \rightsquigarrow no ties.

There are two possible ballots: $a \succ b$ and $b \succ a$.

Anonymity \rightsquigarrow only number of ballots of each type matters.

Denote as A the set of voters voting $a \succ b$ and as B those voting $b \succ a$. Distinguish two cases:

- Whenever $|A| = |B| + 1$ then only a wins. Then, by PR, a wins whenever $|A| > |B|$ (which is exactly the simple majority rule). ✓
- 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. ✓

More than Two Alternatives

For more than two alternatives, our three axioms are not sufficient anymore to characterise a specific rule.

For example, both *plurality* and *Borda* satisfy all of them.

But *plurality with runoff* violates positive responsiveness:

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

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

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

B is eliminated in the first round and C beats A 14:7 in the runoff.

But if 2 of the voters in the first group *raise C to the top* (i.e., if they join the second group), then B wins (beating C 11:10 in the runoff).

Young's Theorem

Another seminal result (which we won't discuss in detail here) is *Young's Theorem*. It provides a characterisation of the *positional scoring rules*. The core axiom is *reinforcement* (a.k.a. *consistency*):

- ▶ F satisfies *reinforcement* if, whenever we split the electorate into two groups and some alternative would win for both groups, then it will also win for the full electorate:

$$F(\mathbf{R}) \cap F(\mathbf{R}') \neq \emptyset \quad \Rightarrow \quad F(\mathbf{R} \oplus \mathbf{R}') = F(\mathbf{R}) \cap F(\mathbf{R}')$$

Young showed that F is a (*generalised*) *positional scoring rule* iff it satisfies *anonymity*, *neutrality*, *reinforcement*, and a technical condition known as *continuity*.

Here, “generalised” means that the scoring vector need not be decreasing.

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

Strategic Manipulation

Suppose the *plurality rule* is used to decide an election: the candidate ranked first most often wins.

Recall our Florida 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

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

Is there a better voting rule that avoids this dilemma?

Strategy-Proofness

Convention: Until further notice, we shall deal with *resolute* voting rules F and write $F(\mathbf{R}) = x$ instead of $F(\mathbf{R}) = \{x\}$.

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.

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

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 2 (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)

Recall: $N_{x \succ y}^{\mathbf{R}}$ 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 $F(\mathbf{R}') \neq y$. [independence]

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 a profile \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 $C_1 \cap C_2$. [now we'll use $|\mathcal{X}| = 3$]

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*! ✓

Arrow's Theorem

You may have heard of Arrow's Theorem, the first and most famous impossibility theorem in social choice theory.

The original theorem is about *social welfare functions*, which are aggregators of the form $F : \mathcal{L}(\mathcal{X})^n \rightarrow \mathcal{L}(\mathcal{X})$, but there is also a version for voting rules:

Any resolute voting rule for ≥ 3 candidates that is *independent* and *Paretian* is a *dictatorship*.

Proof: Omitted (but similar to our proof of the G-S Thm).

K.J. Arrow. *Social Choice and Individual Values*. John Wiley and Sons, 2nd edition, 1963. First edition published in 1951.

Interlude: Mechanism Design

You may have heard that the *Vickrey Auction* is strategy-proof:

*One item on auction. Each bidder makes a bid (offering a price).
The highest bidder wins, but pays the second highest price. No
incentive to misrepresent true price.*

We *could* think of this as an election:

- Candidates = set of all pairs of the form $(winner, price)$.
- Preference order determined by bid. E.g., if *Alice* bids €5:

$$\begin{aligned} \dots \succ (Alice, 3) \succ (Alice, 4) \succ (Alice, 5) \sim (Bob, 17) \sim \dots \\ \dots \sim (Carla, 2) \succ \dots \succ (Alice, 6) \succ (Alice, 7) \succ \dots \end{aligned}$$

That is: very specific preferences (“*quasi-linear utilities*”)

This *domain restriction* is the underlying reason why we can turn the general impossibility (Gibbard-Satterthwaite) into a specific possibility (Vickrey).

Single-Peakedness

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

Another 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.

Summary: Classical 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
- *Gibbard-Satterthwaite*: strategy-proofness leads to dictatorships
- *Arrow*: Pareto (unanimity) and independence lead to dictatorships
- *Black*: single-peakedness solves most problems

Observe that also the *impossibility theorems* above may be considered *characterisation theorems* (namely, of dictatorships).

Literature: Classical Theorems

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

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

For a succinct exposition of most of the results covered here, with full proofs, see my expository paper cited below.

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.

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

Voting Theory and 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 the domain of *ranking sets of objects* using a SAT solver (Geist and E., 2011).

T. Nipkow. Social Choice Theory in HOL. *J. Autom. Reas.*, 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. 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 simple majority 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! (on each issue, 7 out of 13 vote *no*)

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.

Summary: 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.

Literature: Computational Social Choice

Work on COMSOC has been presented at all recent editions of AAMAS, IJCAI, AAI, and ECAI, and also regularly appears in JAIR, AIJ, and JAAMAS. Still, the best source for the latest developments in the field are the proceedings of the biannual COMSOC workshop.

A possible entry point into the field is the chapter cited below.

F. Brandt, V. Conitzer, and U. Endriss. Computational Social Choice. In G. Weiss (ed.), *Multiagent Systems*. MIT Press, 2013.

Conclusion

Last Slide

- We have seen: many voting rules; classical theorems on voting in SCT; examples for recent work at the interface with AI
- Nice topic, particularly for AI and MAS 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/easss-2013/>
 - <http://www.illc.uva.nl/~ulle/teaching/comsoc/>