Logic and Social Choice Theory

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

[http://www.illc.uva.nl/~ulle/teaching/esslli-2013/]

Plan for Today

Yesterday we have focussed on *preference aggregation* and analysed *social welfare functions* (mapping *preference profiles* to *preferences*).

Today we switch to *voting rules*, which map *preference profiles* to *winning alternatives*. Specific topics:

- more voting rules, Fishburn's classification of Condorcet extensions
- May's Theorem as an example for a characterisation result
- strategic behaviour and the Gibbard-Satterthwaite Theorem

Voting Rules

Rules we have seen already:

- *Plurality:* elect the alternative ranked first most often
- *Plurality with runoff:* run a plurality election and retain the two front-runners; then run a majority contest between them
- *Borda:* positional scoring rule with vector $\langle m-1, m-2, \ldots, 0\rangle$
- *Sequential majority:* run a series of pairwise majority contests, always promoting winners to the next stage

All of them map profiles of individual preferences (linear orders on the alternatives) to (sets of) winning alternatives.

An important rule that does not fit into this schema:

- *Approval:* voters can approve of as many alternatives as they wish, and the alternative with the most approvals wins ($\neq$ k-approval)

Copeland Rule

Most of the rules discussed so far violate the Condorcet principle . . .

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).

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

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.
- 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.

Formal Framework

Finite set of n *voters* (or *individuals* or *agents*) $\mathcal{N} = \{1, \ldots, n\}$.
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* $\boldsymbol{R} = (R_1, \ldots, R_n)$ 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:

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

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

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 fix a given rule
- *Impossibility theorems:* certain axioms cannot be satisfied together

Anonymity and Neutrality

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

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

for any profile $\boldsymbol{R}$ and any permutation $\pi : \mathcal{N} \to \mathcal{N}$

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

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

for any profile $\boldsymbol{R}$ and any permutation $\pi : \mathcal{X} \to \mathcal{X}$

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

Remark: You cannot get both A and N for *resolute* rules.

Positive Responsiveness

<u>Notation:</u> $N^{\boldsymbol{R}}_{x \succ y}$ is the set of voters ranking x above y in profile $\boldsymbol{R}$.

A (not necessarily resolute) voting rule satisfies *positive responsiveness* if, whenever some voter raises a (possibly tied) winner $x^\star$ in her ballot, then $x^\star$ will become the *unique* winner. Formally:

> F satisfies *positive responsiveness* if $x^\star \in F(\boldsymbol{R})$ implies $\{x^\star\} = F(\boldsymbol{R'})$ for any alternative $x^\star$ and any two *distinct* profiles $\boldsymbol{R}$ and $\boldsymbol{R'}$ with $N^{\boldsymbol{R}}_{x^\star \succ y} \subseteq N^{\boldsymbol{R'}}_{x^\star \succ y}$ and $N^{\boldsymbol{R}}_{y \succ z} = N^{\boldsymbol{R'}}_{y \succ z}$ for all $y, z \in \mathcal{X} \setminus \{x^\star\}$.

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).

Strategic Manipulation

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

Recall yesterday's Florida example:

49%:	Bush $\succ$ Gore $\succ$ Nader
20%:	Gore $\succ$ Nader $\succ$ Bush
20%:	Gore $\succ$ Bush $\succ$ Nader
11%:	Nader $\succ$ $\underline{\text{Gore}}$ $\succ$ $\underline{\text{Bush}}$

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: For the remainder of today, we shall deal with *resolute* voting rules F and write $F(\boldsymbol{R}) = x$ instead of $F(\boldsymbol{R}) = \{x\}$.

F is *strategy-proof* (or *immune to manipulation*) if for no individual $i \in \mathcal{N}$ there exist a profile $\boldsymbol{R}$ (including the "truthful preference" R_i of i) and a linear order R_i' (representing the "untruthful" ballot of i) such that $F(\boldsymbol{R}_{-i}, R_i')$ is ranked above $F(\boldsymbol{R})$ 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: $(\boldsymbol{R}_{-i}, R_i')$ is the profile obtained by replacing R_i in $\boldsymbol{R}$ 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 $\boldsymbol{R}$ such that $F(\boldsymbol{R}) = x$.
- F is a *dictatorship* if there exists a voter $i \in \mathcal{N}$ (the dictator) such that $F(\boldsymbol{R}) = \text{top}(R_i)$ for any profile $\boldsymbol{R}$.

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

Theorem 2 (Gibbard-Satterthwaite) *Any resolute voting rule for $\geqslant 3$ 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 *$\geqslant$ 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): Independence and Blocking Coalitions

Recall: $N^{\boldsymbol{R}}_{x \succ y}$ is the set of voters who rank x above y in profile $\boldsymbol{R}$.

Claim: If $F(\boldsymbol{R}) = x$ and $N^{\boldsymbol{R}}_{x \succ y} = N^{\boldsymbol{R'}}_{x \succ y}$, then $F(\boldsymbol{R'}) \neq y$. [independence]

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

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

Some more terminology:

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

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

Proof (2): Ultrafilters

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

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

> Proof: Assume C is not blocking; i.e., C is not blocking for (x, y).
> Then there exists a profile $\boldsymbol{R}$ with $N^{\boldsymbol{R}}_{x \succ y} = C$ but $F(\boldsymbol{R}) = y$.
> But we also have $N^{\boldsymbol{R}}_{y \succ x} = \overline{C}$. Hence, $\overline{C}$ 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 $\boldsymbol{R}$ with $C_1 = N^{\boldsymbol{R}}_{x \succ y}$, $C_2 = N^{\boldsymbol{R}}_{y \succ z}$, and $C_1 \cap C_2 = N^{\boldsymbol{R}}_{x \succ z}$. 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 properties (+ finiteness of $\mathcal{N}$) imply that there's a *singleton $\{i\}$* that is blocking. But that just means that i is a *dictator*! ✓

Single-Peakedness

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

The classical way to circumvent this problem are *domain restrictions*.
The most important domain restriction is due to Black (1948):

- Definition: A profile is *single-peaked* if there exists a "left-to-right" ordering $\gg$ on the candidates such that any voter ranks x above y if x is between y and her top candidate wrt. $\gg$.

  Think of spectrum of political parties.

- Result: Fix a dimension $\gg$. Assuming that all profiles are single-peaked wrt. $\gg$, the *median-voter rule* is strategy-proof.

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

Complexity as a Barrier against Manipulation

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*, 553(11):74–82, 2010.

Summary

This has been an introduction to voting theory. We have seen several voting rules (Borda, Copeland, Kemeny, . . .) and covered:

- *Fishburn's classification* of Condorcet extensions in terms of information requirements: (weighted) majority graph etc.
- *May's Theorem:* simple majority rule for two alternatives fully characterised by anonymity, neutrality, and positive responsiveness
- *Gibbard-Satterthwaite Theorem:* impossible to do nondictatorial resolute voting with $\geqslant 3$ possible winners that is strategy-proof
- Circumventing the impossibility: *single-peakedness*, *complexity*