Axiomatic Social Choice

Strategic Manipulation

Zoi Terzopoulou

GATE

University of Lyon — Saint-Etienne

Plan for Today

So far we have assumed that voters will report ballots that reflect their sincere preferences. However, they may also lie to benefit from an outcome they like more. This is called strategic behaviour. We will see:

- Gibbard-Satterthwaite Theorem (1973/1975), showing that this cannot be avoided.
- Conditions and ways out of this impossibility.
- Extensions to incomplete preferences and ballots.
- Iterative voting.

Example

Under the plurality rule the alternative ranked first most often wins.

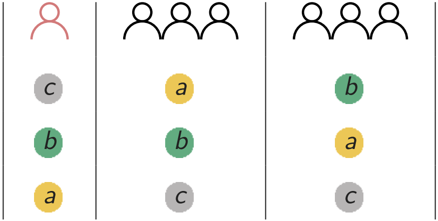

What would you do?
Is there a better voting rule to avoid this problem?

Definitions

For now, we look at resolute voting rules $F : \mathcal{L}(A)^n \to A$.
We need to distinguish:

- The preference a voter has over the alternatives.
- The ballot the voter reports.

We assume that both are elements of $\mathcal{L}(A)$ (but in general they don't need to be). They coincide when the voter is truthful.

F is strategyproof (or immune to manipulation) if for no voter $i \in N$ there exist a profile $\boldsymbol{P} = (\succ_1, \ldots, \succ_i, \ldots, \succ_n)$ and a ballot $\succ'$ such that:

$$F(\boldsymbol{P}_{-i}, \succ_i') \succ_i F(\boldsymbol{P}_{-i}, \succ_i)$$

Voter i prefers the outcome obtained by reporting the untruthful ballot $\succ_i'$ to the outcome obtained by reporting the truthful preference $\succ_i$.

We write $(\boldsymbol{P}_{-i}, \succ_i') = (\succ_1, \ldots, \succ_i', \ldots, \succ_n)$.

Why Strategyproofness?

- Voters should not have to waste resources pondering over what other voters will do and trying to figure out how to respond. Most often, they will not be able to figure this out perfectly.

- If voters strategise, then the final outcome of the election will be based on a skewed ballot profile, and the winner may not be a good compromise overall.

The Gibbard-Satterthwaite Theorem

Recall that F is called surjective if for every alternative $a \in A$ there exists a profile $\boldsymbol{P}$ such that $F(\boldsymbol{P}) = a$.

Gibbard and Satterthwaite independently proved that:

Theorem. *For $m \geqslant 3$, any voting rule that is surjective and strategyproof is a dictatorship.*

A. Gibbard. Manipulation of Voting Schemes: A General Result. Econometrica, 1973.

M.A. Satterthwaite. Strategy-proofness and Arrow's Conditions. Journal of Economic Theory, 1975.

Proof Sketch

We will prove that the Gibbard-Satterthwaite Theorem is a corollary of the Muller-Satterthwaite Theorem (even if G-S is an older result).

For $m \geqslant 3$, any surjective and str. monotonic rule is a dictatorship.
We will show that strategyproofness implies strong monotonicity.

U. Endriss. Logic and Social Choice Theory. Logic and Philosophy Today, 2011.

S. Barberá. Strategy-Proofness and Pivotal Voters: A Direct Proof the Gibbard-Satterthwaite Theorem. International Economic Review, 1983.

J.-P. Benoît. The Gibbard-Satterthwaite Theorem: A Simple Proof. Economic Letters, 2000.

Strategyproofness implies Strong Monotonicity

- SP: No voter prefers the outcome of an untruthful ballot.
- SM: If $F(\boldsymbol{P}) = a$ then $F(\boldsymbol{P'}) = a$ when for all b, $N^{\boldsymbol{P}}_{a\succ b} \subseteq N^{\boldsymbol{P'}}_{a\succ b}$

We will prove the contrapositive. Assume that F is not SM.
Then, there exist $a, a' \in A$ with $a \neq a'$ and profiles $\boldsymbol{P}, \boldsymbol{P'}$ s.t.:

- $N^{\boldsymbol{P}}_{a\succ b} \subseteq N^{\boldsymbol{P'}}_{a\succ b}$ for all $b \in A$ $(*)$
- $F(\boldsymbol{P}) = a$ and $F(\boldsymbol{P'}) = a'$

Moving from $\boldsymbol{P}$ to $\boldsymbol{P'}$ successively, there is a first voter changing the winner. Wlog assume that $\boldsymbol{\succ}$ and $\boldsymbol{\succ'}$ only differ on i.

One of the following two cases holds:

- $i \in N^{\boldsymbol{P'}}_{a\succ a'}$: if i's truthful preference is as in $\boldsymbol{P'}$, she can benefit from voting instead as in $\boldsymbol{P}$ and SP fails. ✓
- $i \notin N^{\boldsymbol{P'}}_{a\succ a'}$: $(*)$ implies $i \notin N^{\boldsymbol{P}}_{a\succ a'}$. So $i \in N^{\boldsymbol{P}}_{a'\succ a}$. If i's truthful pref. is as in $\boldsymbol{P}$, she can benefit from voting as in $\boldsymbol{P'}$ and SP fails. ✓

Overview

How the three main impossibility theorems relate to each other:

Arrow's Theorem
(Pareto + Independent $\Rightarrow$ Dictatorial)

$\Downarrow$

Muller-Satterthwaite Theorem
(Surjective + Strongly Monotonic $\Rightarrow$ Dictatorial)

$\Downarrow$

Gibbard-Satterthwaite Theorem
(Surjective + Strategyproof $\Rightarrow$ Dictatorial)

The Condition of Complete Information

The classical work on strategic manipulation in voting assumes that the manipulation has full information about the ballots of other voters. This makes sense sometimes:

- In small committees (e.g., students electing the classroom representative), this is more realistic.
- Many elections are preceeded by polls, which allow voters to have fairly accurate information about the votes of others.
- We may need to use a voting rule that is safe against strategic manipulation in the worst case.

Incomplete Information

An information function I tells the information that every voter has, given a (truthful) profile $\boldsymbol{P}$ and a voting rule F. For example:

- Winner information: $I_i(\boldsymbol{P}) = F(\boldsymbol{P})$ for every $i \in N$
- Zero information: $I_i(\boldsymbol{P}) = \emptyset$ for every $i \in N$

Every voter considers some possible profiles, given her information:

$$W_i^{\boldsymbol{P}} = \{\boldsymbol{P'} \in \mathcal{L}(A)^n \text{ such that } I_i(\boldsymbol{P'}) = I_i(\boldsymbol{P})\}$$

F is strategyproof under partial information if for no voter $i \in N$ there exist a profile $\boldsymbol{P} = (\succ_1, \ldots, \succ_i, \ldots, \succ_n)$ and a ballot $\succ'$ such that:

- $F(\boldsymbol{P'}_{-i}, \succ_i') \succ_i F(\boldsymbol{P'}_{-i}, \succ_i)$ for some $\boldsymbol{P'} \in W_i^{\boldsymbol{P}}$
- $F(\boldsymbol{P''}_{-i}, \succ_i') \succeq_i F(\boldsymbol{P''}_{-i}, \succ_i)$ for all other $\boldsymbol{P''} \in W_i^{\boldsymbol{P}}$

Strategyproofness under Incomplete Information

Recall the plurality rule. What would you do now?

c	?	?
b	?	?
a	?	?

Theorem: *When $n > 2m - 2$, any scoring rule, paired with the lexicographic tie-breaking rule, is immune to zero-manipulation.*

Theorem: *When $n > 3$, any Condorcet-consistent voting rule, paired with the lexicographic tie-breaking rule, is immune to zero-manipulation.*

A. Reijngoud & U. Endriss. Voter Response to Iterated Poll Information. AAMAS, 2012.

The Condition of Complete Preferences

Tell the truth, the whole truth, and nothing but the truth.

Voters may lie by:
flipping, omitting, or adding preferences.

Scoring rules are safe when and only when a voter cannot increase the score difference of a preferred alternative over another one.

Kruger and Terzopoulou. Strategic Manipulation with Incomplete Preferences: Possibilities and Impossibilities for Positional Scoring Rules. *AAMAS*, 2020.

Some More Scoring Rules

- Domination scoring: $w_{\succ}(x) = |\{y \in A : x \succ y\}|$
- Veto scoring:

$$w_{\succ}(x) = \begin{cases} 1 & \text{if there is some } y \in A \text{ such that } x \succ y, \\ 0 & \text{otherwise} \end{cases}$$

- Cumulative scoring:

$$w_{\succ}(x) = \begin{cases} 0 & \text{if there is no } y \in A \text{ such that } x \succ y, \\ \sum_{y \in A : x \succ y} + 1 & \text{otherwise} \end{cases}$$

Example

Consider acyclic (not necessarily transitive) preferences.

$$\begin{array}{ccccc} 1 & \rightarrow & 1 & \rightarrow & 0 \\ \uparrow & \nwarrow & & & \\ 3 & \rightarrow & 2 & \rightarrow & 1 \\ \downarrow & & & \nearrow & \downarrow \\ 0 & & 1 & & 0 \end{array} \qquad \begin{array}{ccccc} 2 & \rightarrow & 1 & \rightarrow & 0 \\ \uparrow & \nwarrow & & & \\ 7 & \rightarrow & 4 & \rightarrow & 1 \\ \downarrow & & & \nearrow & \downarrow \\ 0 & & 2 & & 0 \end{array} \qquad \begin{array}{ccccc} 1 & \rightarrow & 1 & \rightarrow & 0 \\ \uparrow & \nwarrow & & & \\ 1 & \rightarrow & 1 & \rightarrow & 1 \\ \downarrow & & & \nearrow & \downarrow \\ 0 & & 1 & & 0 \end{array}$$

Scores given by domination, cumulative, and veto scoring.

Strategyproofness under Incomplete Preferences

- Omission: Cumulative
- Addition: Veto
- Omission + addition: Constant (only!)
- Flipping: ?

J. Kruger & Z. Terzopoulou. Strategic Manipulation with Incomplete Preferences: Possibilities and Impossibilities for Positional Scoring Rules. AAMAS, 2020.

Manipulation in Iteration

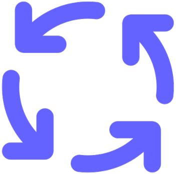

Agents may modify their opinions iteratively.

We want equilibria, where noone has an incentive to change ballot.

R. Meir et Al. Convergence to Equilibria in Plurality Voting.
AAAI, 2010.

Example

2-approval & deterministic tie-breaking order $a; b; c; d$.

agent 1: $a \longrightarrow c \longrightarrow b \longrightarrow d$

agent 2: $d \longrightarrow c \longrightarrow b \longrightarrow a$

Example – Allowing only Complete Preferences

2-approval & deterministic tie-breaking order $a; b; c; d$.

agent 1: $a \longrightarrow c \longrightarrow b \longrightarrow d \dashrightarrow a \longrightarrow b \longrightarrow c \longrightarrow d$

agent 2: $d \longrightarrow c \longrightarrow b \longrightarrow a$

Example – Allowing only Complete Preferences

2-approval & deterministic tie-breaking order $a; b; c; d$.

agent 1: $a \longrightarrow c \longrightarrow b \longrightarrow d \dashrightarrow a \longrightarrow b \longrightarrow c \longrightarrow d$

agent 2: $d \longrightarrow c \longrightarrow b \longrightarrow a \dashrightarrow d \longrightarrow b \longrightarrow c \longrightarrow a$

Example – Allowing only Complete Preferences

2-approval & deterministic tie-breaking order a; b; c; d.

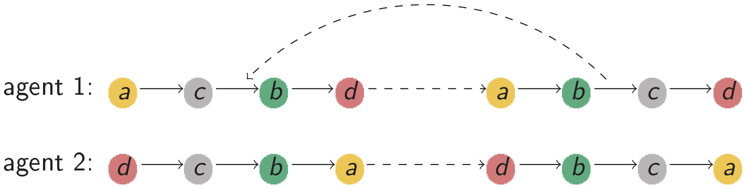

Example – Allowing only Complete Preferences

2-approval & deterministic tie-breaking order $a; b; c; d$.

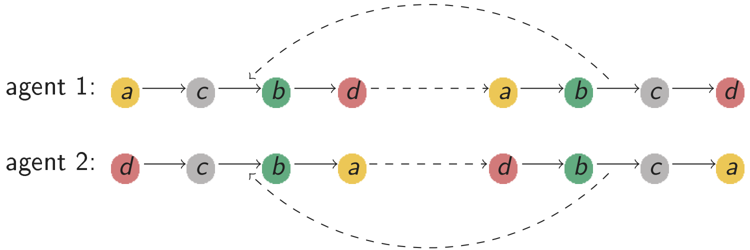

There is no equilibrium!

Example – Allowing Incomplete Preferences

2-approval & deterministic tie-breaking order $a; b; c; d$.

agent 1: $a \longrightarrow c \longrightarrow b \longrightarrow d \dashrightarrow a \longrightarrow b \longrightarrow c \longrightarrow d$

agent 2: $d \longrightarrow c \longrightarrow b \longrightarrow a$

Example

2-approval & deterministic tie-breaking order $a; b; c; d$.

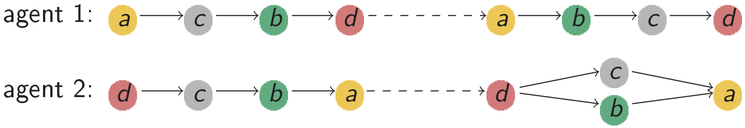

Assume that omitting is cheaper than flipping a preference pair.

Example – Allowing Incomplete Preferences

2-approval & deterministic tie-breaking order a; b; c; d.

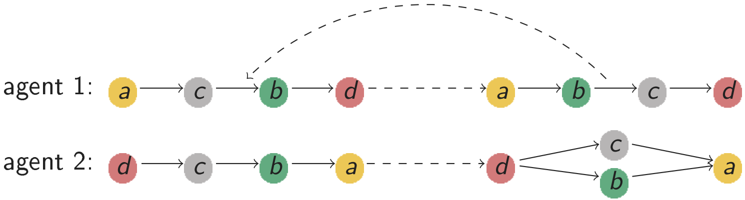

The equilibrium winner is c!

Summary

We have seen the problem of strategic manipulation in voting, where voters may report untruthful preferences to obtain a better outcome.

- Gibbard-Satterthwaite proved an impossibility theorem.
- Under incomplete information, the impossibility can be avoided.
- Under incomplete preferences, more manipulation moves exist.
- Equilibria of iterative voting depend on the allowed moves.

Thank you, and stay in touch!

PhD and Postdoc positions in Saint-Etienne :)