Games Manipulators Play

Umberto Grandi

Department of Mathematics
University of Padova

23 January 2014

[Joint work with Edith Elkind, Francesca Rossi and Arkadii Slinko]

Gibbard-Satterthwaite Theorem

All (reasonable) voting rules are susceptible to strategic voting:

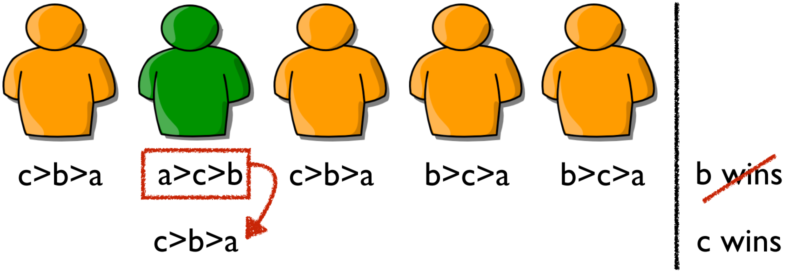

A. Gibbard. Manipulation of Voting Schemes: A General Result. Econometrica, 1973.
M. Satterthwaite. Strategy-proofness and Arrows Conditions... Journal of Economic Theory, 1975.

But manipulation may not be so easy...

A large literature points at the difficulties of strategic voting:

- In the economic literature, critiques to the unrealistic assumption of perfect information: a single voter needs to know the exact preference distribution to manipulate
- In the computer science literature, an analysis of the computational complexity of computing manipulation strategies: easy for some rules (e.g., scoring rules), hard for others such as STV
- ...and...

One step further: GS games

Manipulation may require coordination:

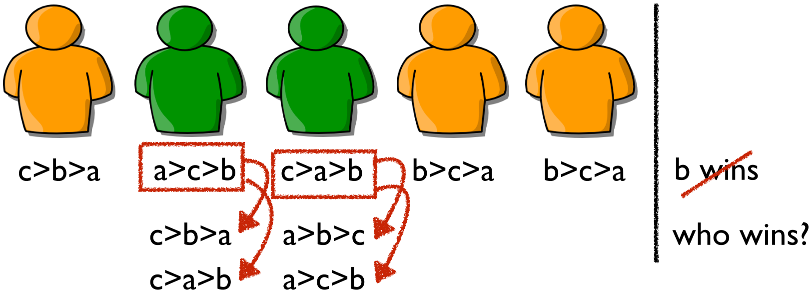

In this work we study the strategic game that manipulators have to play in situations of strategic voting: Gibbard-Satterthwaite games.

Outline

1. Basic definitions: voting theory, manipulation, GS Theorem
2. Gibbard-Satterthwaite Games
3. Special case: 2×2 GS games
   - Simple classification of 2×2 games
   - Full characterisation for plurality
   - Borda can implement the full spectrum of 2×2 GS games
4. General case: k-approval
   - Plurality (1-approval) always have a NE
   - Mild assumptions for 2 and 3-approval to have a NE
   - ≥ 4-approval do not have NE
5. How to compute weakly dominant strategies? (ongoing)

Plurality, k-Approval and Borda

- A set $\mathcal{N}$ of voters
- A set $\mathcal{X}$ of candidates
- A profile of preferences (linear orders) $\boldsymbol{P} = (P_1, \dots, P_n)$

Voting rules aggregate individual preferences in a (set of) winning candidate(s).

Positional Scoring Rules

A positional scoring rule (PSR) is defined by a vector $(s_1, \dots, s_m)$. The voting rule gives s_j points to candidates ranked in position j in one individual preference, and elect the candidates with the maximal score.

We focus on some particular PSRs:

- Plurality, with vector (1,0,. . . ,0)
- k-approval, with vector $(1, \dots, 1, 0, \dots, 0)$ with exactly k 1s
- Borda, with vector $(m-1, m-2, \dots, 0)$

We always use lexicographic tie-breaking.

Strategic Voting - Notation

- A manipulation strategy for voter i at $\boldsymbol{P}$ is a linear order P_i' such that:

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

Gibbard-Satterthwaite Theorem tells us that for any (reasonable) voting rule there always is a profile in which a player has a manipulation strategy.

- A manipulator at $\boldsymbol{P}$ is any voter who has a manipulation strategy at profile $\boldsymbol{P}$
- Let $V_{\boldsymbol{P}} \subset \mathcal{N}$ be the set of GS manipulators at profile $\boldsymbol{P}$.

We want to study the game that is played by GS manipulators

Definition

Given a voting rule F and a profile of preferences $\boldsymbol{P} = (P_1, \ldots, P_n)$, a Gibbard-Satterthwaite game (GS game) is a normal-form game

$$\mathcal{G} = \langle V_{\boldsymbol{P}}, \{S_i \mid i \in V_{\boldsymbol{P}}\}, \{\succeq_i \mid i \in V_{\boldsymbol{P}}\} \rangle$$

Where:

- $V_{\boldsymbol{P}}$ is the set of GS-manipulators at profile $\boldsymbol{P}$;
- $S_i = \{s, i_1, \ldots, i_\ell\}$ where $s = P_i$ and $\{i_1, \ldots, i_\ell\}$ is a (possibly empty) set of undominated manipulation strategies available to player i;
- $\succeq_i$ is a preference relation defined over joint strategy profiles:

$$(s_1, \ldots, s_k) \succeq_i (s'_1, \ldots, s'_k) \text{ if and only if}$$

$$F(s_1, \ldots, s_k) P_i F(s'_1, \ldots, s'_k) \text{ or } F(s_1, \ldots, s_k) = F(s'_1, \ldots, s'_k)$$

Example

Ingredients: 2 voters, 5 alternatives, Borda rule, lexicographic tie-breaking.

v_1	v_2
a	b
c	c
d	d
b	a
e	e

c wins

v_2
b
d
c
e
a

b wins

v_1
a
d
c
e
b

d wins

	Sincere	Insincere
Sincere	3,3	1,4
Insincere	4,1	2,2

Prisoner's dilemma

v_1
a
e
d
c
b

a wins

	Sincere	Insincere
Sincere	3,3	1,4
Insincere	4,1	4,1

Game Analysis

Quick notation: a Nash equilibria (NE) is a strategy profile where no voter has incentive to deviate, and a weakly dominant strategy (WDS) for a voter is a strategy that in all profiles leads to a better outcome than any of her strategies.

We want to study how "hard" the manipulation game can be:

There is a NE in WDS. Easy to solve: each manipulator plays a WDS.

There is a unique NE. Also easy: individuals play the unique NE strategies.

There are multiple NEs. Requires coordination: players need to communicate to choose the NE that they want to achieve.

There is no NE. Better not play it?

Related Work

This framework is closely related to:

- **Iterative voting**: best-response dynamics on full voting games – all voters are players, all strategies are possible (but restrictions have been considered), the process converges to a NE.
- **Voting games studied in economics**: focus on mixed strategies, tackling the problem of complete information.

GS games impose three restrictions on full voting games:

1. only manipulators are players
2. only undominated manipulation strategies
3. possibly subsets of manipulation strategies

Simple manipulation games: 2 manipulators, 1 strategy

Two manipulators using a greedy algorithm to compute a manipulation move:

	Sincere	Insincere
Sincere	a,b	e,f
Insincere	c,d	g,h

Lemma

If any two payoff pairs have equal first (second) components, then their second (first) components should also be equal (e.g., if $a = e$, then also $b = f$).

Lemma

A 2×2 GS game for the plurality rule has at most three distinct payoff pairs.

Corollary: the Prisoner's Dilemma is not a GS game for plurality.

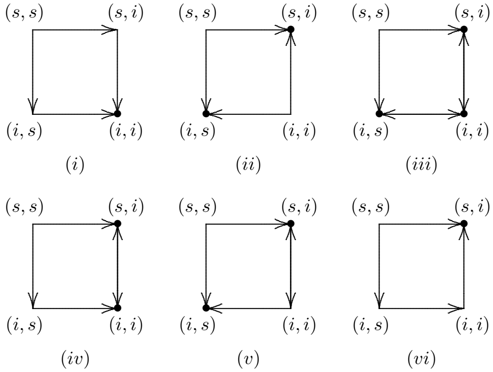

Proposition

All 2×2 GS games have at least one NE.

Corollary: Battle of the Sexes, Matching Pennies, are not GS games.

Full Classification

Proposition

All game forms are GS games for the Borda rule.

Proof. By examples. The initial example showed that (i) and $(vi)'$ is representable with Borda. For instance:

	Sincere	Insincere
Sincere	3,3	1,4
Insincere	4,1	2,2

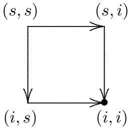

(i)

Proposition

(i), (v) and (vi) are not GS games for Plurality. All others are.

Proof idea. Remember there can only be 3 different outcomes.

(Preliminary) Conclusions

We can already draw a number of considerations:

- 2×2 games are not too hard, they always have a NE.
- Some of them are! Prisoner's dilemma cases, but also coordination games.
- Except for one case, games for plurality are solvable in WDS:

  How do we compute WDS? Is it hard to find them?

General case: the question

Having a complete charactherisation is out of range, even restricting to 2 strategies per player...

A simpler question:
is there at least one NE?

Our answer:

- Yes, for plurality
- Yes, for 2 and 3-approval under restrictive assumptions
- No, for $\geq$4-approval (even under restrictive assumptions)

Plurality always have a NE

Two strategies are equivalent if they yield the same outcome in all profiles:

Lemma

Optimal manipulation strategies for plurality are equivalent.

Proposition

Every GS game for plurality has at least one Nash equilibrium.

Proof idea. A candidate is called competitive at profile $\boldsymbol{P}$ if there is a manipulation in its favour. Take the competitive candidate with the highest plurality score (or tie-breaking position) and have one voter manipulate in its favour and all other sincere. $(s, \ldots, s, i, s, \ldots, s)$ is a NE.

This result can be obtained as a corollary of a known result in iterative voting:

Meir Et Al. Convergence to Equilibria in Plurality Voting. AAAI-2010.

k-approval: General results

Let P be a profile and a is the winner at P. Manipulation strategies in a GS game for k-approval can be divided into two categories:

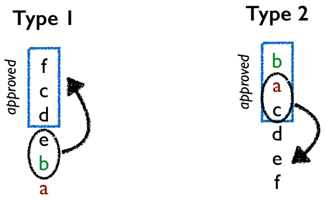

Lemma

The set of competitive candidates can be partitioned as follows:

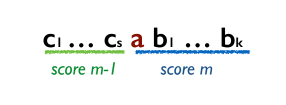

with candidates in lexicographic order.

k-approval: Restrictions on manipulation strategies

Definition

A manipulation of type 1 in favour of candidate x is called sound if it does not increase the score of candidates that are preferred to x.

Definition

A manipulation in favour of x is called minimal if:

type 1: *x is swapped with the alternative in k-th position.*

type 2: *a minimal number of ℓ approved alternatives is swapped with the top ℓ not approved alternatives (including x).*

A GS game is sound (resp. minimal) if all manipulation strategies are sound (resp. minimal).

2 and 3-approval: Existence results

Proposition

Every sound GS game for 2-approval has at least one NE.

Proof idea. Non trivial. Several cases, but if there is a type 1 minimal manipulation then the winner is the lexicographically highest competitive candidate and the profile again looks like $(s, \ldots, s, i, s, \ldots, s)$.

Proposition

Every minimal GS game for 3-approval has at least one NE.

Proof idea. Non trivial. Type 1 manipulation solved as previously. Type 2 manipulation can be done in favour of only two candidates, and the proof is done by cases.

4-approval: no NE!

Proposition

There exists a minimal GS game for 4-approval that does not have a NE

Proof ideas. An example. The game has 6 voters and 8 candidates. By minimality each voter has only two strategies, so we can picture is as the following cube:

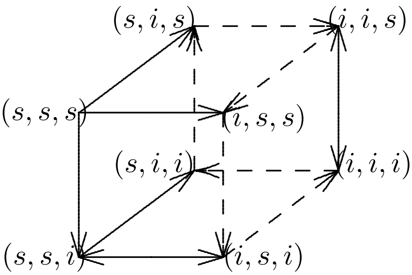

Corollary: iterative voting for k-approval ($k \geq 4$) does not converge!

Ongoing work

Weakly dominant strategies: simplest games, individually justifiable moves...

How hard is to find whether manipulating is a WDS
for Plurality, k-approval, Borda?

Some preliminary results:

- Plurality: WDS for c iff my top candidate ($\neq c$) is not competitive or there is no other voters that can make my top candidate the winner. Polynomial to check.
- 2-approval: non trivial. But still polynomial to check under minimality.
- 4-approval: co-NP-hard!

Conclusions

GS games are the kind of games that manipulators
need to solve when facing a situation of strategic voting.

Assumptions are strong:

- Only manipulators are players
- Only a subset of undominated manipulation strategies

A good starting point to draw conclusions on more complex voting games:

1. There is no guarantee that strategic voting is easy to perform! Even in simple (2×2) cases it may require coordination. (Similar argument to the computational complexity of manipulation)

2. Interesting assumptions can be devised to restrict strategic behaviour (minimality, fairness...) depending on how voters compute their strategies: this affects the difficulty of strategic voting.

3. computing WDS is the next step!