

Games Gibbard-Satterthwaite Manipulators Play

Umberto Grandi


Department of Mathematics
University of Padova

25 June 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. Results:
 - Full characterisation of 2×2 GS games
 - Computational properties of weak dominance
 - Existence of NE for general case

Plurality, k-Approval and Borda

- A set of n voters
- A set C of candidates
- A profile of preferences (linear orders) $V = (v_1, \dots, v_n)$

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 $(1, 0, \dots, 0)$
- 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 profile V is a linear order v'_i such that
player i prefers $\mathcal{R}(V_{-i}, v'_i)$ to $\mathcal{R}(V)$

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 **GS manipulator** at V is any voter who has a manipulation strategy. Let $N(V, \mathcal{R})$ be the set of **GS manipulators** at profile V for rule \mathcal{R}
- A **GS manipulation** is undominated (optimal) manipulation strategy

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

Gibbard-Satterthwaite Games

Definition

Given a voting rule \mathcal{R} and a profile of preferences $V = (v_1, \dots, v_n)$, a Gibbard-Satterthwaite game (GS game) is a normal-form game

$$\mathcal{G} = \langle N(V, \mathcal{R}), \{S_i \mid i \in N(V, \mathcal{R})\}, \{\succeq_i \mid i \in N(V, \mathcal{R})\} \rangle$$

Where:

- $N(V, \mathcal{R})$ is the set of **GS-manipulators**;
- $S_i = \{s, i_1, \dots, i_\ell\}$ where $s = v_i$ and $\{i_1, \dots, i_\ell\}$ is a (possibly empty) set of **GS manipulation strategies** available to player i ;
- \succeq_i is a **preference relation** defined over joint strategy profiles:

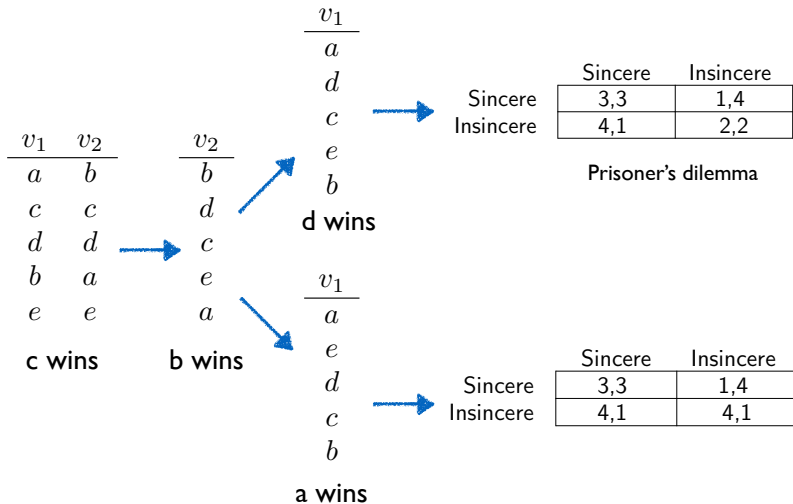
$(s_1, \dots, s_k) \succeq_i (s'_1, \dots, s'_k)$ if and only if

i prefers $\mathcal{R}(s_1, \dots, s_k)$ to $\mathcal{R}(s'_1, \dots, s'_k)$ or $\mathcal{R}(s_1, \dots, s_k) = \mathcal{R}(s'_1, \dots, s'_k)$

We use the position of a candidate as the payoff (=Borda score).

Example

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



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" manipulation games can be:

There is a NE in WDS. Easy to solve: each manipulator plays a WDS (generalisation of **safe manipulation**)

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**: convergent best-response dynamics on **full voting games**
– all voters are players, all strategies are possible [MeirEtAl, AAAI-2010...]
- **Voting games studied in economics**: focus on mixed strategies, tackling the problem of complete information [Myerson and Weber 1993 ...]
- **Truth-biased equilibria**: rule out undesirable equilibria by adding a small price of manipulation [Obratsova et al. SAGT 2013]
- **Safe manipulation**: is manipulation weakly dominant even if manipulators have the same preferences? [Slinko and White, SCW 2013]

Simple manipulation games: 2 manipulators, 1 strategy each

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



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



Proposition

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


Two games are solvable in WDS, one is a **coordination game**

What happens with countermanipulators

We can use the results in iterative voting:

- Plurality always has a NE (from convergence)
- Borda or 2-approval may have **no NE!**

4 candidates, 2 voters, and the tie breaking rule is alphabetical ($a \succ b \succ c \succ d$)



Meir et al. Convergence to equilibria of plurality voting. AAAI-2010
Lev and Rosenshien. Convergence of iterative voting. AAMAS-2012

Computational properties of weak dominance

How hard is to find out whether
manipulating weakly dominates truth-telling?

Results for k -approval:

- 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 mildly restrictive assumptions on players' strategies (minimality)
- 3-approval: hard (conjecture)
- 4-approval: co-NP-hard!

Existence of a NE in general GS games

Having a complete characterisation 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 (minimality)
- No, for 4-approval (even under restrictive assumptions)

Conclusions

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

Assumptions are strong but realistic:

- **Only manipulators** are players (strategic behaviour triggered by the possibility of manipulation)
- Only a subset of **undominated manipulation strategies** (greedy algorithms, weak dominance, minimal changes)

Conclusions:

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, soundness...) depending on how voters compute their strategies: this affects the difficulty of strategic voting.

k -approval: Restrictions on manipulation strategies

Definition

A "positive manipulation" 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:

"positive": x is swapped with the alternative in k -th position.

"negative": 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).

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 it as the following cube:



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