Gibbard–Satterthwaite Games for k-Approval Voting Rules

Umberto Grandi[1], Daniel Hughes[2], Francesca Rossi[3], and Arkadii Slinko[*†4]

[1]*IRIT, University of Toulouse, France*
[2]*Dept of Mathematics, The University of Auckland, New Zealand*
[3]*Dept of Mathematics, University of Padova, Italy*
[4]*Dept of Mathematics, The University of Auckland, New Zealand*

Abstract

The Gibbard-Satterthwaite theorem implies the existence of voters, called manipulators, who can change the election outcome in their favour by voting strategically. When a given preference profile admits several such manipulators, voting becomes a game played by these voters, who have to reason strategically about each other's actions. To complicate the game even further, countermanipulators may try to counteract potential actions of manipulators. Previously, voting manipulation games have been studied mostly for the Plurality rule. We extend this to k-Approval voting rules. However, unlike previous studies, we assume that voters are boundedly rational and do not think beyond manipulating or countermanipulating. We classify all 2-by-2 games that can be encountered by these simple strategic voters, and we investigate the complexity of arbitrary voting manipulation games, identifying conditions on strategy sets that guarantee the existence of a Nash equilibrium in pure strategies.

JEL classification: D72.

Keywords: strategic voting, voting manipulation game, Nash equilibrium.

*Corresponding author. Email addresses: umberto.grandi@irit.fr (U. Grandi), dhug729@aucklanduni.ac.nz (D. Hughes), frossi@math.unipd.it (F. Rossi), a.slinko@auckland.ac.nz (A. Slinko).

†Part of this work was done when Arkadii Slinko was visiting the Institute for Mathematical Sciences (Singapore), University of Padova (Italy) and University of Toulouse (France). He is grateful for the support from Marsden Fund 3706352 of The Royal Society of New Zealand. We are grateful to Edith Elkind who took part in this project at the initial stage. This paper builds on initial results presented at the *24th International Joint Conference on Artificial Intelligence (IJCAI-2015)* (Elkind et al., 2015a).

1 Introduction

Voting is a common method of preference aggregation, which enables the participating agents to identify the best candidate given the individual agents' rankings of the candidates. However, no "reasonable" voting rule is immune to manipulation: as shown by Gibbard (1973) and Satterthwaite (1975), if there are at least three candidates, then any onto, non-dictatorial voting rule admits a preference profile (a collection of voters' rankings) where some voter would be better off by submitting a ranking that differs from his truthful one or, in other words, his truthful vote is not the best response to the votes of other voters. We call such voter a Gibbard-Satterthwaite manipulator or GS-manipulator for short. When such a manipulator is unique, he[1] then has a disproportional influence on the election outcome. However, in the presence of multiple manipulators, their attempt to manipulate the election simultaneously in an uncoordinated fashion (and we assume that no coordination devices exist) may result in an outcome that differs not just from the outcome under truthful voting, but also from the outcome that any of the GS-manipulators could anticipate. This may be due to possible complex interference among the different manipulative votes, and may deter some of the voters (especially risk-averse ones) from manipulating. When we also include in consideration those voters who cannot manipulate themselves but can prevent others from manipulating—so-called countermanipulators—the situation becomes even more complex and can be described only in game-theoretic terms. Let us illustrate the kind of situation on which we focus by an example.

Example 1. *Table 1 describes a voting situation. There are four voters and five candidates a, b, c, d, and w. Each voter ranks the candidates from the most to the least preferred, with the most preferred candidates on the left of the lists in Table 1. Suppose that the rule used to elect the winner is the Plurality rule, which declares the candidate with the highest number of first positions the winner, and that ties are resolved following a predetermined ordering over the candidates, in this case $w > a > b > c > d$.*

Voters	Preferences
Voter 1	$a\,b\,c\,d\,w$
Voter 2	$b\,a\,c\,d\,w$
Voter 3	$w\,c\,a\,d\,b$
Voter 4	$d\,w\,a\,b\,c$

Table 1: A preference profile. The most preferred candidates of voters are in the left-most positions, followed by the less preferred candidates in a complete ranking.

[1]In the paper we refer to candidates as females and voters as males.

In the situation described in Table 1 the winner is w, thanks to the tie-breaking. This result is the worst possible for the first two voters, who each have the opportunity to manipulate. Voters 1 and 2 are the only GS-manipulators at this profile. Voter 1's insincere strategy consists in voting for b instead of a making b the winner; voter 2 can vote insincerely in favour of a instead of b making candidate a the winner. If both manipulate at the same time, their efforts will cancel out. If we zoom in on the situation when voter 1 and voter 2 are strategically choosing between their sincere vote and their manipulation vote, while other voters are not strategic, we see that voters 1 and 2 are playing an anti-coordination game.

Let us now extend our analysis to all the voters. We observe that voter 3 is happy and does not have reasons for strategizing. Voter 4 does not have any incentive to manipulate: the current winner w is in his second position, so giving her more support will not change the outcome. For instance, the voter does not have any incentive to vote for w instead of d. However, this move is a very strong countermanipulation if voter 4 fears any strategic move from any of the first three players: giving additional support to w makes any manipulation impossible, ending all strategic considerations. Moreover, in game-theoretic terms this move weakly dominates his sincere vote.

As shown in the above example, even for such a simple voting rule as Plurality, a single profile can give us a number of different games depending on which voters are strategic and which are not. A non-strategic voter has only his sincere vote in his strategy set, while a strategic voter has more than one strategy. In this paper we are interested in the properties of the normal-form games that arise under k-Approval voting rules (and Plurality corresponds to 1-Approval). These rules are simple enough to allow for a classification of voting manipulations, but complex enough to admit the realization of non-trivial games.

As is usually the case in any initial investigation, it is customary to assume the full information framework, in which everybody's sincere preferences are publicly known as well as their strategy sets. However, the voting intentions of the voters remain private to those voters.

An important novel feature of the class of games considered in this paper, which distinguishes them from voting games that have been considered in prior literature (see Section 2 for related literature survey), is the introduction of types of players which are characterised by their strategy sets. There are several reasons for the introduction of types. Firstly, by reducing players' strategy sets we can model their bounded rationality. Secondly, the knowledge of the sincere profile does not allow to unambiguously decide who is strategic and who is not. Voters may be able to manipulate but reject this on moral grounds, or they may be unable to calculate their manipulation. On the other hand, a voter may not be able to manipulate but can take preventive measures from a disastrous effect of someone else's manipulation (such as voter 4 in Example 1). The introduction of "personalised" strategy sets allows us to bring into a spotlight and to study in isolation various aspects of strategic manipulation, e.g., the interaction of Gibbard-Satterthwaite manipulators (e.g., voters 1 and 2 in Example 1) or the interaction of a manipulator and a countermanipulator.

Given a voting rule, every profile of voters' preferences gives rise to a number of games that can be played; we will call them *voting manipulation games*. We treat the strategy sets as a state of nature whose move makes the structure of the game a public knowledge. One of the most interesting cases is when only GS-manipulators are strategic (no countermanipulators exist), and such games will be called *Gibbard-Satterthwaite games* or *GS-games*. The purpose of our work is to investigate the game-theoretic complexity of voting manipulation games, showing that manipulation in practice might be prevented by, e.g., the absence of Nash equilibria, or the existence of multiple equilibria, which in the absence of coordination mechanisms is a source of strategic complexity.

The simplest non-trivial example of our framework involves two players each having two strategies: one sincere and one insincere, we call them 2-by-2 games. They can be both GS-manipulators or, alternatively, one can be a GS-manipulator and another a countermanipulator. Observe that this does not imply that the election from which similar games arise has two voters only, but simply only two voters at the given profile are strategic. We begin by answering the question of which 2-by-2 games can be represented as voting manipulation games. As existing classifications of games turn out to be too fine-grained for our purposes, we develop a simple coarser classification, and observe that the definition of GS-games imposes certain restrictions on players' preferences. Combining this observation with symmetry arguments, we arrive at 6 basic types of 2-by-2 games played by two manipulators. We then show that, while all six games can be obtained as GS-games under the 2-Approval voting rule, for the Plurality rule (i.e., 1-Approval) only four of them are realizable. In the same spirit we also obtain a classification of 2-by-2 manipulator and countermanipulator games.

Since in the presence of countermanipulators a pure Nash equilibrium may not exist even in a very simple 2-by-2 game, we bring under the spotlight the situation when all players are GS-manipulators. We study the existence of pure strategy Nash equilibria in such games. We show that every GS-game for Plurality has a Nash equilibrium, and identify necessary and sufficient conditions for the existence of Nash equilibria for 2-Approval games. It appears that a mild rationality condition which we call soundness assumption is necessary and sufficient. We also find sufficient conditions for the existence of Nash equilibria of 3-Approval games, assuming that manipulating voters choose manipulation strategies which are in some sense minimal. However, we show that this minimality assumption fails to ensure the existence of Nash equilibria for 4-Approval games.

The paper is organised as follows. In Section 2 we discuss related work, and in Sections 3 we introduce voting manipulation games. The main contributions of the paper are presented in Section 4, in which we classify 2-by-2 voting manipulation games, and in Section 5, where we study the existence of Nash equilibria in arbitrarily large GS-games for k-Approval voting rules.

Section 6 provides a discussion of the results presented and Section 7 concludes the paper.

2 Related Work

There is a substantial body of research dating back to Farquharson (1969) that explores the consequences of modeling non-truthful voting as a strategic game; see, e.g., Fishburn (1978); Moulin (1979); Feddersen et al. (1990); Myerson and Weber (1993); De Sinopoli (2000); Dhillon and Lockwood (2004); Sertel and Sanver (2004); De Sinopoli et al. (2015); Desmedt and Elkind (2010); Obraztsova et al. (2013). The most popular framework so far has been the one introduced by Myerson and Weber (1993), which however stipulates that each voter has a utility for the election of each candidate (so it is not purely ordinal in nature). The main issue tackled by this line of research is the existence of numerous Pareto-dominated Nash equilibria, many of which involve seemingly irrational choices by the voters. The following methods were considered to resolve this problem: equilibria refinements (De Sinopoli, 2000), costly voting (Sinopoli and Iannantuoni, 2005), truth-biased voters (Obraztsova et al., 2013), trembling-hand equilibria (Obraztsova et al., 2016), generic utilities (De Sinopoli et al., 2015). However, in our opinion the problem remains not completely solved. With the notable exception of Obraztsova et al. (2015), in all of these papers the only rule under consideration is Plurality and the set of players consists of all voters, i.e., a player is allowed to vote non-truthfully even if he would be unable to manipulate the election on his own or countermanipulate. We believe that the main problem underlying the proliferation of irrational Nash equilibria is that voters are allowed to vote irrationally. Therefore, in our paper we propose to restrict the set of players to those voters that are GS-manipulators in the original profile, and their strategies to a subset of the manipulative actions of the original profile, altering the model significantly. For instance, we are able to rule out unnatural Nash equilibria such as situations in which all players vote for the same undesirable candidate.

One further approach is considered by Elkind et al. (2017), where they model voting by using an adaptation of the cognitive hierarchy model. In their work, it is assumed that the voters reason about potential actions of other voters assuming that their level of rationality is lower. They take non-strategic (sincere) voters as those belonging to level 0. The players of level 1 play a best response to the assumption that all other players belong to level 0, and the players of level 2 give their best response assuming that all other players belong to level 0 or level 1. We note that players of level 2 are already quite sophisticated. They can, for example, think of countermanipulating or they can strategically stay sincere when they can manipulate. The emphasis of the work of Elkind et al. (2017) is on the complexity of a level 2 voter deciding whether his manipulative strategy weakly dominates his sincere strategy.

The algorithmic aspects of voting games have recently received some attention as well (see, e.g., Elkind et al. (2016); Desmedt and Elkind (2010); Xia and Conitzer (2010); Elkind et al. (2015b)). Empirical analysis of Nash equilibria in plurality elections has been done in Thompson et al. (2013). Iterative voting is the topic considered in the literature that is the closest to our work. In this model players change their votes one by one in response to the current outcome (see, e.g., the recent survey by Meir (2017)). The main difference with our framework is conceptual: our purpose is the evaluation of a one-shot game resulting from Nature choosing a set of manipulators and a set of actions for each of them, while in iterative voting players can make several moves responding to various degrees of information provided on the current vote profile. More technically, our approach assumes a fixed strategy set for each player, defined by the initial truthful profile, while in iterative voting players decide on their next move depending on the profile that resulted after the previous moves, hence considering strategy sets that change over time. The common feature is that both approaches assume boundedly rational voters.

The graphical representation of best-responses in 2-by-2 games that we propose in Section 4 is reminescent of the literature on acyclicity in games (see, e.g., Kukushkin (2011); Apt and Simon (2015)). Meir et al. (2017) recently applied this terminology to the setting of iterative voting, contributing with further results to this literature.

3 Modelling Voting Manipulation Games

In this section we present basic definitions and distinctive features of voting manipulation games.

3.1 Basic Definitions

We consider elections over a candidate set $C = \{c_1, \dots, c_m\}$ in which n voters $1, 2, \dots, n$ participate. An election is defined by a *preference profile* $V = (v_1, \dots, v_n)$, where each v_i, $i = 1, \dots, n$, is a total order over C; we refer to v_i as the *vote*, or *preferences*, of voter i. For two candidates $c_1, c_2 \in C$ we write $c_1 \succ_i c_2$ if voter i ranks c_1 above c_2; if this is the case, we say that voter i *prefers* candidate c_1 to candidate c_2. For brevity we will sometimes write $ab \dots z$ to represent a vote v_i with $a \succ_i b \succ_i \dots \succ_i z$. We denote by $\text{top}(v_i)$ the top candidate in v_i. Also, we denote by $\text{top}_k(v_i)$ the set of top k candidates in v_i.

Given a preference profile $V = (v_1, \dots, v_n)$, we denote by (V_{-i}, v'_i) the preference profile obtained from V by replacing v_i with v'_i; for readability, we will sometimes omit the parentheses around (V_{-i}, v'_i) and write V_{-i}, v'_i. A (resolute) *voting rule* is a mapping $\mathcal{R}$ that, given a profile V, outputs a candidate $\mathcal{R}(V) \in C$ called the *winner* at V under $\mathcal{R}$. We say that two votes v and v' of voter i over the same candidate set C are *equivalent* with respect to a voting rule $\mathcal{R}$, if $\mathcal{R}(V_{-i}, v) = \mathcal{R}(V_{-i}, v')$ for every profile V.

In this paper we consider k-Approval voting rules. Under k-Approval, $1 \leq k \leq m-1$, each candidate receives one point from each voter who ranks her in top k positions. The candidate(s) with the highest score wins. The most well-known rule in this class is the Plurality rule (aka first-past-the-post), which corresponds to 1-Approval. Since any k-Approval voting rule is not resolute and two or more candidates can share the highest score, we complement it with a tie-breaking. In this paper we assume that ties are broken according to a fixed order $>$, usually alphabetic, over C.

We denote the k-Approval score of a candidate c in a profile V by $\mathrm{sc}_k(c, V)$. We will sometimes denote the k-Approval rule by k-App. It is easy to see that v and v' are equivalent with respect to k-Approval if and only if $\mathrm{top}_k(v) = \mathrm{top}_k(v')$. Let us also introduce the relation $c \sqsupset_V c'$ on candidates $c, c' \in C$, which means that at profile V either c has higher score than c' (the exact k-approval score used will be clear from the context), or they have equal scores but $c > c'$ in the tie-breaking order.

We conclude by introducing some useful notation. Let $X = (x_1, \ldots, x_\ell)$ and $Y = (y_1, \ldots, y_\ell)$ be two sequences over disjoint sets of candidates such that no candidate is repeated in any of them and v is a vote. Then $v[X; Y]$ denotes the vote obtained by swapping x_j with y_j for $j = 1, \ldots, \ell$ in the individual preference ordering v. We often denote sequences as $X = x_1 \ldots x_\ell$ and $Y = y_1 \ldots y_\ell$. Then we write $v[x_1 \ldots x_\ell; y_1 \ldots y_\ell]$ instead of $v[X; Y]$.

3.2 Manipulators and countermanipulators

Our goal is to model and investigate situations that arise in voting when one or more voters are strategic. We model such situations as normal form games that we call *voting manipulation games*. We will now define such situations and games formally. We start by defining the first and the main type of strategic voter.

Definition 1. *We say that a voter i is a* Gibbard–Satterthwaite manipulator*, or a* GS-manipulator*, at a profile $V = (v_1, \ldots, v_n)$ with respect to a voting rule $\mathcal{R}$, if there exists a vote $v'_i \neq v_i$ such that i strictly prefers $\mathcal{R}(V_{-i}, v'_i)$ to $\mathcal{R}(V)$. If $\mathcal{R}(V_{-i}, v'_i) = p$, we will also say that i manipulates* in favour of p.

Given a voting rule $\mathcal{R}$ and a profile V, we denote with $N(V, \mathcal{R})$ the set of GS-manipulators at V for $\mathcal{R}$. Next, we model how a strategic voter chooses between several strategic moves available to him.

Definition 2. *A vote v'_i is called a* GS-manipulation *of voter i if:*

- *i prefers $\mathcal{R}(V_{-i}, v'_i)$ to $\mathcal{R}(V)$, and,*
- *for every v''_i it holds that either $\mathcal{R}(V_{-i}, v'_i) = \mathcal{R}(V_{-i}, v''_i)$ or i prefers $\mathcal{R}(V_{-i}, v'_i)$ to $\mathcal{R}(V_{-i}, v''_i)$.*

Sometimes a voter can manipulate in favour of several different candidates; however, in Definition 2 we make the mild rationality assumption that a voter always focuses on his most preferred candidate among the ones that he can make election winners. We now define a second type of strategic voter that we will consider in this paper.

Definition 3. *Suppose voter i is a GS-manipulator at $V = (v_1, \ldots, v_n)$ with respect to a voting rule $\mathcal{R}$ and his manipulation is v'_i. We say that a voter j is a* countermanipulator *at a profile V against v'_i if:*

- *there exists a vote $v'_j \neq v_j$ such that j prefers $\mathcal{R}((V_{-i}, v'_i)_{-j}, v'_j)$ to $\mathcal{R}(V_{-i}, v'_i)$, and,*
- *for every v''_j it holds that either $\mathcal{R}((V_{-i}, v'_i)_{-j}, v'_j) = \mathcal{R}((V_{-i}, v'_i)_{-j}, v''_j)$ or j prefers $\mathcal{R}((V_{-i}, v'_i)_{-j}, v'_j)$ to $\mathcal{R}((V_{-i}, v'_i)_{-j}, v''_j)$.*

In this case vote v'_j is called a countermanipulation *of voter j against v'_i.*

The two above definitions introduce two basic types of strategic voters, and more sophisticated agents may of course be considered. For example, a voter may not be able to change the result of the election unilaterally (in large elections voters are seldom pivotal) but he may hope that there will be other like-minded voters who will also change their vote in a similar way and the desired change may come about as a result of combined efforts (Slinko and White, 2008, 2014). Alternatively, he may also countermanipulate against a coalition of manipulators. The higher his level of rationality, the more strategic motives the voter understands and the more complex game he faces as a result (for a more in-depth analysis see Elkind et al. (2017)).

3.3 Voting Manipulation Games

Recall that a *normal-form game* is defined by a set of *players* N, and, for each player $i \in N$, a set of *actions* A_i and a preference relation $\succeq_i$ defined on the space of *action profiles*, i.e., on tuples of the form $(a_1, \ldots, a_n)$, where $a_i \in A_i$ for all $i \in N$. Alternatively we can think that there is a function $f\colon A_1 \times \ldots \times A_n \to \mathcal{O}$, where $\mathcal{O}$ is the set of outcomes and relations $\succeq_i$ are defined on $\mathcal{O}$.[2] We denote the set of all strategic voters at a profile V with respect to a voting rule $\mathcal{R}$ by $N(V, \mathcal{R})$. After nature makes its move, this set is known to everybody. Moreover, in our case the preference relation of player i on the action profiles is determined by the outcome of $\mathcal{R}$ on those action profiles.

A *voting manipulation game* at a profile V is any game with the set of players the set of all GS-manipulators in V under $\mathcal{R}$, and such that for every player i his strategy set A_i, includes this voter's sincere vote. Given an action profile $V^* = (v^*_i)_{i \in N}$, let $V[V^*] = (v'_1, \ldots, v'_n)$ be the preference profile such that $v'_i = v_i$ for $i \notin N$ and $v'_i = v^*_i$ for $i \in N$. Then, given two action

[2]While normal-form games may be defined either in terms of utility functions or in terms of preference relations, the latter approach is more suitable for our setting, as we only have ordinal information about the voters' preferences.

profiles V^* and V^{**}, we write $V^* \succeq_i V^{**}$ if and only if $\mathcal{R}(V[V^*]) = \mathcal{R}(V[V^{**}])$ or player i prefers $\mathcal{R}(V[V^*])$ to $\mathcal{R}(V[V^{**}])$. In what follows, we denote this game $G = (V, \mathcal{R}, (A_i)_{i\in N(V,\mathcal{R})})$.

The class of voting manipulation games that is of primary interest to our analysis is the following: for each preference profile V and each voting rule $\mathcal{R}$, a *Gibbard-Sattertwhaite game*, or *GS-game* for short, is a normal-form game defined as having as the set of players N the set of all GS-manipulators in V under $\mathcal{R}$, where, for each player i, his set of actions A_i consists of his truthful vote and a subset (possibly empty) of his GS-manipulations. Different choices of the strategy sets correspond to different games in the family. We denote the set of all GS-games for V and $\mathcal{R}$ by $\mathcal{GS}(V, \mathcal{R})$. Note that all games in $\mathcal{GS}(V, \mathcal{R})$ have the same set of players, namely, $N(V, \mathcal{R})$, so an individual game in $\mathcal{GS}(V, \mathcal{R})$ is fully determined by the players' sets of actions, i.e., $(A_i)_{i\in N(V,\mathcal{R})}$. When V and $\mathcal{R}$ are clear from the context, we simply write $G = (A_i)_{i\in N}$. We refer to an action profile in a GS-game as a *GS-profile*; we will sometimes identify the GS-profile $V^* = (v_i^*)_{i\in N}$ with the preference profile $V[V^*]$.

In Section 4 will also consider another form of voting manipulation games, in which there are two strategic players, one of which is a GS-manipulator and another is a countermanipulator. A GS-manipulator has a GS-manipulation that changes the outcome of the election in his favour; and some other player, called countermanipulator, cannot manipulate but can perform a countermanipulation that neutralises to some extent the action of the manipulator (such as voter 4 in Example 1). We call such games *Manipulator-Countermanipulator games* or *MC-games* for short. We will only focus on 2-by-2 games of this type—i.e., one manipulator and one countermanipulator.

3.4 Discussion of the Model

We emphasise that in our definition of GS-games, rather than considering games where each player's set of actions consists of his truthful vote and *all* of his GS-manipulations, we allow the players to limit themselves to subsets of their GS-manipulations. There are several reasons for that. First, the introduction of strategy sets allows us to bring into a spotlight and to study in isolation various aspects of strategic manipulation e.g., interaction of GS-manipulators alone. Second, the space of all GS-manipulations for a given voter can be very large, and a player may be unable or unwilling to identify all such votes; indeed, even counting the number of GS-manipulations for a given voter is a non-trivial computational problem (Bachrach et al., 2010). Thus, the player may use a specific algorithm (e.g., the greedy algorithm proposed by Bartholdi et al., 1989 for the class of scoring rules) to find his GS-manipulation; in this case, his set of actions would consist of his truthful vote and the output of this algorithm. Also, the player may choose to ignore GS-manipulations that are (weakly) dominated by other GS-manipulations. Finally, a player may prefer not to change his vote beyond what is necessary to make his target candidate the election winner, either because he wants his vote to be as close to his true preferences as possible (see, e.g., Obraztsova and Elkind, 2012), or for fear of unintended consequences of such changes in the complex environment of the game.

4 2-by-2 Voting Manipulation Games

In this section, we investigate which 2-by-2 games (i.e., games with two players, and two actions per player) can be represented as GS-games or as MC-games. We show that even in such a restricted framework we can realise a surprising variety of games.

4.1 Representation of 2-player games

In this section we focus on the question of which 2-by-2 games can be realised as voting manipulation games of a particular kind. To do so, we need a suitable classification of 2-by-2 games. Note, first, that every such game corresponds to 4 action profiles, and is fully described by giving both players' preferences over these profiles. By considering all possible pairs of preference relations over domains of size 4, Fraser and Kilgour (1986) showed that there are 724 distinct 2-by-2 games. However, this classification is too fine-grained for our purposes. Thus, we propose a simplified approach that is based on the following two principles. First, we only compare action profiles that differ in exactly one component. Second, when comparing two profiles that differ in the i-th component ($i = 1, 2$), we only take into account the preferences of the i-th player. Thus, every 2-by-2 game can be represented by a diagram with 4 vertices and 4 directed edges, where an edge is directed from a less preferred profile to a more preferred profile and a bidirectional edge indicates indifference.

4.2 2-by-2 GS-games

Now, let us focus on GS-games with 2 players and 2 actions per player, one of which is their sincere vote. For each player, let s denote his sincere vote and let i denote his manipulative vote; thus, the vertices of our diagram are (s, s), (i, s), (s, i), and (i, i). For two edges of this diagram their direction is determined by the fact that i is a GS-manipulation: namely, both of the edges adjacent to (s, s) are directed away from (s, s). Thus, by renaming the players if necessary, any 2-by-2 GS-game for any voting rule can be represented by one of the six diagrams in Figure 1. Observe that an action profile in a 2-by-2 game is a Nash equilibrium if and only if the corresponding vertex in the diagram of the game has two incoming edges. The following proposition is immediate:

Proposition 1. *Every 2-by-2 GS-game has at least one Nash equilibrium.*

Proof. Since both insincere strategies are GS-manipulations we have arrows from (s, s) both to (i, s) and (s, i). For these not to be Nash equilibria both arrows from (i, s) to (i, i) and from

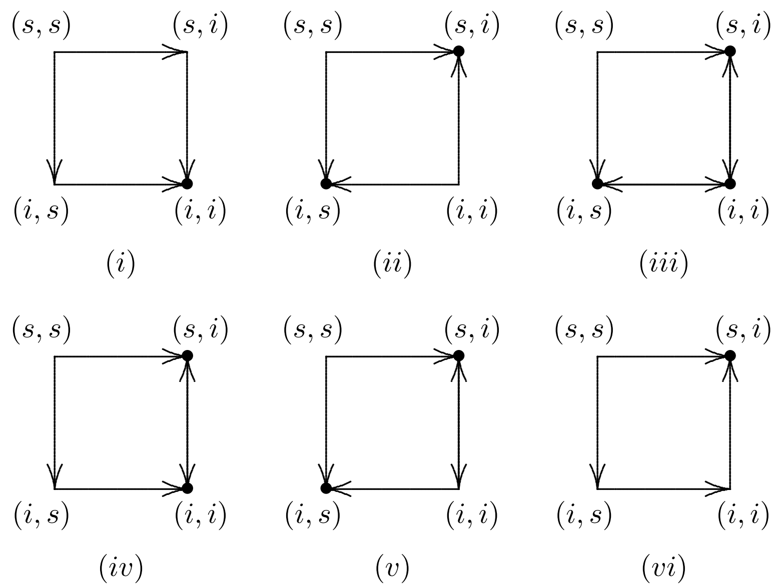

Figure 1: Diagrams for 2-by-2 GS-Games

(s, i) to (i, i) must be directed towards (i, i). Therefore, either (i, i) is a Nash equilibrium, or at least one of (i, s) and (s, i) is a Nash equilibrium. □

On the six diagrams in Figure 1 Nash equilibria are marked by black dots.

Example 2. *Consider the GS-game for the preference profile* $(v_1, v_2, v_3, v_4) = (abc, bac, cab, cba)$ *under the Plurality voting rule, with ties broken according to* $a > b > c$. *In this game players* 1 *and* 2 *are the GS-manipulators; their GS-manipulations are* $v_1^* = v_1[a; b]$ *and* $v_2^* = v_2[b; a]$, *which result in election of* b *and* a, *respectively. Note that, if both GS-manipulators vote insincerely,* c *remains the election winner. Thus, this game corresponds to diagram (ii) in Figure 1.*

We will say that a diagram D in Figure 1 is *realisable* as a diagram of a GS-game by a voting rule $\mathcal{R}$ on n-voter profiles if there exists a preference profile V consisting of at most n voters and a 2-by-2 game $G \in \mathcal{GS}(V, \mathcal{R})$ such that D is the diagram for G. We also say that D is *realisable* by a voting rule $\mathcal{R}$ if there exists a profile V (without restriction on the number of voters) such that D is the diagram for G.

Our next goal is to understand which diagrams are realizable by the voting rules we have chosen to concentrate on in this paper, i.e., k-Approval voting rules.

Theorem 1. *The only diagrams realizable by Plurality are (ii), (iii), (iv) and (v).*

Proof. Consider a profile V, and assume that voters 1 and 2 are the only GS-manipulators in V. Suppose that the Plurality winner in V is w, voter 1 manipulates in favour of a, and voter 2 manipulates in favour of b. Since 1 and 2 are GS-manipulators, we have $w \neq \text{top}(v_1)$, $w \neq \text{top}(v_2)$, and we can assume that the GS-manipulations of voters 1 and 2 are given by $v_1^* = v_1[\text{top}(v_1); a]$, $v_2^* = v_2[\text{top}(v_2); b]$. Let $V^1 = (V_{-1}, v_1^*)$, $V^2 = (V_{-2}, v_2^*)$, $V^{1,2} = (V_{-2}^1, v_2^*)$.

The winner in $V^{1,2}$ can be w, a, or b, and the case where the winner is w corresponds to diagram (ii). Further, if $a = b$, then a is the winner at V^1, V^2, and $V^{1,2}$; this corresponds to diagram (iii). Thus, suppose that the winner at $V^{1,2}$ is a or b and $a \neq b$. This means that one of the arrows adjacent to (i, i) must be bidirectional, ruling out the three diagrams (i), (ii) and (vi). We have thus proved that diagrams (i) and (vi) are not realizable by Plurality.

Example 2 shows how to realise diagram (ii). We now construct examples for the remaining three cases. Diagram (iii) can be realised in profile $V = (cawb, bawc, wabc)$ with ties broken according to $w > a > b > c$. The winner at V is w, and both the first and second players can manipulate in favour of a, which is therefore the winner at all profiles V^1, V^2, and $V^{1,2}$.

Consider now the profile $V = (dabwc, cbawd, wbacd)$ with tie-breaking order $w > a > b > c > d$. Candidate w is the winner in V, Voter v_1 manipulates in favour of a (horizontal arrow), and v_2 in favour of b. If both players manipulate, the result is still a by the tie-breaking order. Since v_1 prefers a to b this example realises diagram (iv).

Diagram (v) can be obtained on profile $V = (bacw, cbaw, wbac)$, using $a > w > b > c$ as tie-breaking rule. Candidate w is the winner at V, candidate a is winning at V^1 and $V^{1,2}$, and candidate b is winning at V^2, however this time v_1 prefers b to a and regrets his manipulation. □

The remaining diagrams are realizable by 2-Approval voting rule.

Theorem 2. *Diagrams (i) and (vi) are both realizable by 2-Approval voting rule.*

Proof. In both cases we will consider the alphabet tie-breaking, i.e., the tie-breaking order is $a > b > \ldots > y > z$.

Diagram (i). Let $V = (v_1, v_2, v_3) = (xywa\ldots, ztwb\ldots, cd\ldots)$. The 2-Approval winner at V is c. The first two voters are the GS-manipulators with manipulations $v_1^* = v_1[xy; aw]$ and $v_2^* = v_2[zt; bw]$ in favour of a and b, respectively. Further, at (v_1^*, v_2^*, v_3) the 2-Approval winner is w which both manipulators prefer over a and b.

Diagram (vi). Let $V = (v_1, v_2, v_3) = (xywa\ldots, ztbaw\ldots, cd\ldots)$. The 2-Approval winner at V is c. The first two voters are the GS-manipulators with manipulations $v_1^* = v_1[xy; aw]$ and $v_2^* = v_2[zt; bw]$ in favour of a and b, respectively. At (v_1^*, v_2^*, v_3) the 2-Approval winner is w, in which case voter 2 would regret manipulating as he'd rather prefer a than w. This realises diagram (vi). □

Observe that in the above proof, the second manipulator in the realization of diagram (vi) chose an illogical manipulation by promoting w and not promoting a. We will call such manipulations

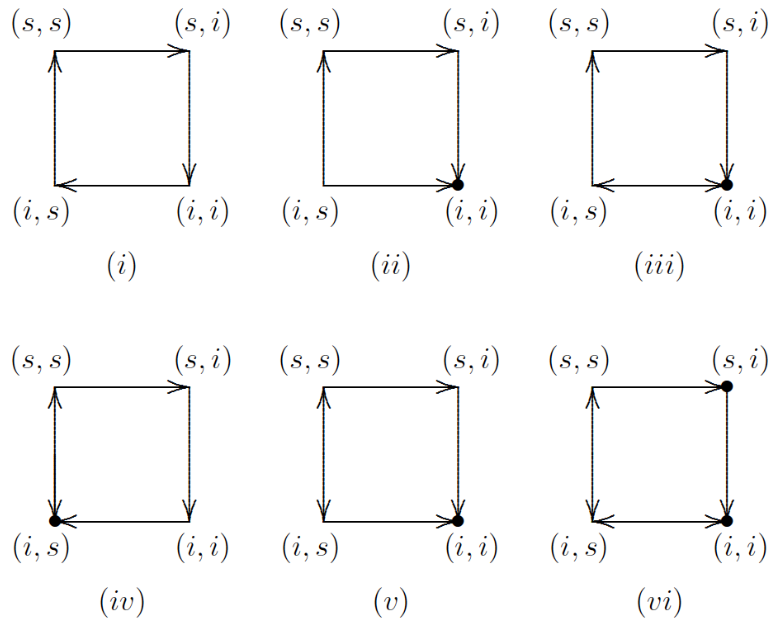

Figure 2: Diagrams for 2-by-2 Manipulator-Countermanipulator Games

unsound in Section 5. We note that for the Borda rule all six diagrams are realizable as well (see Elkind et al., 2015a, for a proof of this result).

4.3 2-by-2 MC-games

There are six possible game diagrams that can be achieved in a 2-by-2 MC-game, namely those (i)-(vi) presented in Figure 2. Observe that apart from the first one, all the diagrams in Figure 2 have at least one Nash equilibrium. As before, we denote by s the sincere action of the player and by i his insincere action.

Theorem 3. *Only the forms (i)-(vi) are possible to emerge as diagrams of* 2*-by-*2 *MC-games. All six diagrams are realizable under 2-Approval. Only (iii) is realizable for two voters under Plurality and (vi) can be realised with more than two voters under Plurality.*

Proof. We will use the relation $c \sqsupset_V c'$ on candidates $c, c' \in C$, with the interpretation that at profile V either c has higher Plurality score than c' or they have equal scores but $c > c'$.

Plurality: It can easily be seen that only one form is possible with two voters under Plurality. Suppose our two voters, 1 and 2, have favourite candidates a and w respectively, with w being higher on the tie-breaking order so w wins (note, if both have the same favourite neither can be a manipulator). Then voter 1 may manipulate in favour of a new candidate x which beats w on tie-break by submitting it first instead of a. The only way voter 2 can countermanipulate is promoting a new candidate y that beats x on tie-break. Then $y > x > w > a$. If voter 1 plays sincerely, but voter 2 countermanipulates, then y wins again but voter 2 would regret countermanipulating. Thus we get diagram (iii) and it is the only possible diagram realizable by Plurality on 2-voter profiles.

If there are more than two voters, then diagram (vi) is also realizable but no other. Again suppose a wins with s points at sincere profile V and that voter 1 has a manipulation in favour of some candidate x, where x originally either also had s points with $a > x$ or x had $s - 1$ points with $x > a$, obtaining a profile V_1 with x as the winner. Then voter 2, as the counter manipulator, can either promote a new candidate, y, to beat x, or promote a and make him winner again at profile V_2 depending on which is the best candidate in favour of whom voter 2 can countermanipulate. We show that the first option is not possible.

Indeed, if a was not voter 2's best candidate in favour of whom he can countermanipulate but y, then $V = (b \cdots x \cdots a \cdots, c \cdots y \cdots, U)$, where U is a subprofile, and a wins at V. Then x wins at $V_1 = (xb \cdots a \cdots, c \cdots y \cdots, U)$, and y wins at $V_2 = (xb \cdots a \cdots, yc \cdots, U)$. But then y will also win at $V_3 = (b \cdots x \cdots a \cdots, yc \cdots, U)$ since $y \sqsupset_{V_3} a$ (their positions are the same as in V_2) and $a \sqsupset_{V_3} b$ (their positions are the same as in V. But that would be a manipulation for voter 2 at V which is impossible.

If a was voter 2's best candidate in favour of whom he can countermanipulate, then $V = (b \cdots x \cdots, c \cdots a \cdots, U)$, x wins at $V_1 = (xb \cdots, c \cdots a \cdots, U)$, and after voter 2 promoted a she becomes unbeatable at $V_2 = (xb \cdots, ac \cdots, U)$ and also at $V_3 = (b \cdots x \cdots, ac \cdots, U)$. We therefore have diagram (vi).

2-Approval: In Table 2 we give examples of profiles for each diagram under 2-Approval voting with two voters (apart from (iii) which has already been realised with two voters under Plurality in Example 1). For all profiles the tie-breaking rule is $x > e > a > b > c > d > u$ and a wins under sincere votes. Voter 1 is the manipulator (column player) and voter 2 is the countermanipulator (row player). This proves the theorem. □

We note that apart from the diagram (i), all remaining diagrams in Figure 2 have at least one Nash equilibrium.

5 Nash Equilibria

One of the most important characteristics of any game is whether it has Nash equilibria (NE) in pure strategies or not. Various schools of thought provide different opinions on whether or not players will end up choosing equilibrium strategies. We do not wish to contribute to this debate, but we simply consider the existence of an NE as an important feature of the game, and one that will guide players in their analysis of a situation of strategic vote.

In Section 4, we showed that even for 2-by-2 games in the presence of countermanipulators we cannot expect a voting manipulation game to have a NE while for GS-games the existence of

Diagram	Voter	Preferences	Strategic votes
(i)	1	$c\,d \mid e\,x\,b\,a$	$v_1[e; d]$
	2	$a\,b \mid d\,x\,c\,e$	$v_2[ab; dx]$
(ii)	1	$c\,d \mid e\,u\,x\,b\,a$	$v_1[cd; eu]$
	2	$a\,b \mid u\,x\,d\,c\,e$	$v_2[ab; ux]$
(iv)	1	$b\,a \mid u\,c\,d$	$v_1[a; u]$
	2	$c\,d \mid u\,a\,b$	$v_2[d; u]$
(v)	1	$c\,d \mid u\,e\,x\,b\,a$	$v_1[cd; eu]$
	2	$a\,b \mid u\,x\,d\,c\,e$	$v_2[b; u]$
(vi)	1	$b\,a \mid u\,c\,d$	$v_1[a; u]$
	2	$c\,d \mid a\,u\,b$	$v_2[d; a]$

Table 2: Realizations of diagrams (i),(ii),(iv)-(vi) under 2-Approval voting rule.

a NE is guaranteed. However, when more than two strategic players are manipulators, things are not so obvious. In this section, we study the existence of Nash equilibria in arbitrary GS-games for k-Approval with $k = 1, 2, 3, 4$, introducing further conditions on the players' strategy sets when necessary.

5.1 Voting Manipulation Games under Plurality

We will first show that for Plurality, a Nash equilibrium always exists for every voting manipulation game, even in presence of countermanipulators.

Theorem 4. *Any voting manipulation game under Plurality has a Nash equilibrium in pure strategies.*

Proof. Fix a profile V and let w be the Plurality winner at V with score t. Let $S \subset C$ be the set of candidates such that either their score is t and they are beaten by w on the tie-break or their score is $t-1$ and they are higher than w on the tie-breaking order. Let $S^+ \subset S$ be the set of candidates in favour of whom there is a GS-manipulation. Let $p \in S^+$ be the candidate such that $p \sqsupset_V q$ for all $q \in S^+$ and suppose voter i can manipulate in favour of p with $v_i^* = v_i[\text{top}(v_i); p]$. Suppose, first, that there is another voter k who can manipulate in favour of p with $v_k^* = v_k[\text{top}(v_k); p]$. Then it is easy to check that $((V_{-i}, v_i^*)_{-k}, v_k^*)$ is a Nash equilibrium with winner p. If this second manipulator in favour of p does not exist, no GS-manipulator (existing at V) can change the result at (V_{-i}, v_i^*) by his actions due to the choice of p. However a countermanipulator can possibly change the result countermanipulating in favour of $q \in S \setminus S^+$ (exactly such situation has occurred in Example 1). Note that his top preference is not p thus he cannot make the score of p lower. Let $q \in S \setminus S^+$ be maximal with respect to $\sqsupset_V$ for which some voter, say voter j with v_j^*, can countermanipulate. Then $((V_{-i}, v_i^*)_{-j}, v_j^*)$ is a Nash equilibria. □

5.2 Manipulation Strategies for k-Approval

We first prove a lemma on manipulation strategies, to get more insight into GS-games for k-Approval with $k \geq 2$. Note, first, that under k-Approval any GS-manipulation of voter i is equivalent to a vote of the form $v_i[X;Y]$, where $X \subseteq \text{top}_k(v_i)$ and $Y \subseteq C \setminus \text{top}_k(v_i)$. There are two types of GS-manipulators for k-Approval voting: those who rank the current winner w in top k positions, and those who do not. We give a more precise description in the following lemma.

Lemma 1. *Let the voting rule be k-Approval for $k \geq 1$. Let V be a profile and w be the winner at V. Let also x be an alternative, other than w. Then any manipulation in favour of x at V falls under one and only one of the following two categories:*

Type 1 *A voter i increases the score of x by 1 without decreasing the score of w. In this case both w and x are not approved by the manipulator with $x \succ_i w$, the manipulator moves x to the first k positions leaving w not approved. We refer to voter i as a* promoter *of x.*

Type 2 *A voter i reduces the score of w and possibly the scores of some other alternatives by 1 without increasing the score of x. In this case both w and x are approved by the manipulator with $x \succ_i w$, the manipulator removes w from his top k positions leaving x there. We refer to voter i as a* demoter *of w.*

Clearly, for $k = 1$ only manipulations of type 1 are possible.

Proof. If voter i can increase the score of an alternative x in favour of which he manipulates, then it was not approved by this voter originally. He, however, must rank it higher than w (otherwise this is not a GS-manipulation). Thus, w was not approved either and voter i cannot decrease its score, falling under the first category which we named Type 1 manipulations. The second possibility is that the voter cannot increase the score of the alternative x in favour of which he manipulates since it is already approved by him. This means that he is left with reducing the scores of some of its competitors, including the current winner w. This means w is also in the approval set with $x \succ_i w$, corresponding to our definition of Type 2 manipulations. This type of manipulation cannot happen for $k = 1$ since both w and x must be in the approval set. □

5.3 2-Approval GS-Games with Unrestricted Manipulations

We begin by exhibiting a seemingly discouraging example which shows that for 2-Approval, there are GS-games with no Nash equilibria, albeit allowing for counterintuitive manipulation strategies.

Example 3. *Consider three players, v_1, v_2, v_3, with the following preferences under the 2-Approval voting rule. The tie-breaking rule is $n > a > m > x > b > c > d > e > f$. Let*

$$V = (v_1, v_2, v_3) = (ab|cdefnmx,\ cd|nmxbaef,\ ef|bmxnacd).$$

Under this profile a wins. Voter v_1 is not a manipulator as his most preferred candidate, a, wins. Both v_2 and v_3 can manipulate in favour of n and b, respectively. We define the strategy sets $A_2 = \{s_2, i_2, i'_2\}$ and $A_3 = \{s_3, i_3, i'_3\}$, where s_2, s_3 are the sincere strategies of v_2 and v_3 respectively. Here i_2 and i'_2 are manipulations of voter 2 where he swaps c and d with n and m, and n and x, respectively, i.e.,

$$i_2 = v_2[cd; nm], \qquad \textit{and} \qquad i'_2 = v_2[cd; nx].$$

Similarly, i_3 and i'_3 are manipulations of voter 3 where he swaps e and f with b and m and b and x, respectively, submitting

$$i_3 = v_3[ef; bm], \qquad \textit{and} \qquad i'_3 = v_3[ef; bx].$$

To see that this game has no Nash equilibrium we have to consider all 9 strategy profiles that can be realised. These are summarised in Table 3 where we give the winner in the given profile, which voter has an incentive to change their strategy and the change in strategy that would be favourable for that voter. As can be seen, in every profile it is favourable for at least one of the manipulators to change their strategy, hence no Nash equilibrium exists.

Profile	Winner	Voter	Change in Strategy
$(s_2,\ s_3)$	a	3	s_3 to i_3
$(s_2,\ i_3)$	b	2	s_2 to i_2
$(s_2,\ i'_3)$	b	2	s_2 to i'_2
$(i_2,\ s_3)$	n	3	s_3 to i'_3
$(i_2,\ i_3)$	m	3	i_3 to i'_3
$(i_2,\ i'_3)$	b	2	i_2 to i'_2
$(i'_2,\ s_3)$	n	3	s_3 to i_3
$(i'_2,\ i_3)$	b	2	i'_2 to i_2
$(i'_2,\ i'_3)$	x	3	i'_3 to i_3

Table 3: Profitable deviations at each of the 9 strategy profiles in Example 3

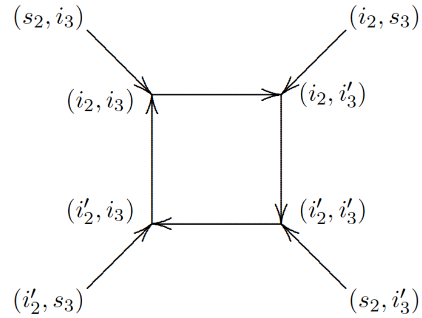

Figure 3: Diagram for the game in Example 3

A graphical representation of the game is in Figure 3, in line with the figures used in Section 4. As we can see, instead of a NE, we have an attractor in the form of a 4-cycle.

Example 3 exhibits a GS-game without Nash equilibrium with a minimal number of voters. For two voters, as we will see below, every game has a Nash equilibrium. To show this we first need to prove the following lemma.

Lemma 2. *In any GS-game under the 2-Approval voting rule all Type 2 manipulators must have the same two top preferences and in the same order. In particular, they manipulate in favour of the same candidate.*

Proof. First note that all Type 2 manipulators approve the winner w of the sincere profile V but place her in the second position. Also, if a voter is a Type 2 manipulator, its other approved candidate must either be on the same score as w and be lower in tie-breaking order or be on the score one less than w and be higher in the tie-breaking order. Suppose there are two Type 2 manipulators with different top preferences, call them v_1 and v_2 with $a_1 = \text{top}(v_1) \neq \text{top}(v_2) = a_2$, i.e., $v_1 = a_1 w \ldots$ and $v_2 = a_2 w \ldots$. When v_1 manipulates, demoting w and promoting some other candidate, a_1 wins. This means, in particular, $a_1 \sqsupset_V a_2$. But similarly, we get $a_2 \sqsupset_V a_1$. This is a contradiction, showing that all Type 2 manipulators must have the same top two preferences (and in the same order). □

We are now ready to prove the following:

Theorem 5. *Every 2-Approval GS-game with 2 voters has a Nash equilibrium in pure strategies.*

Proof. If only one voter is a manipulator then there is a trivial Nash equilibrium (the manipulator plays insincerely). Assume then that both players v_1 and v_2 are manipulators. Using the terminology of Lemma 1, let us first analyse the case of when both voters are of Type 1 or both of Type 2. First, both voters cannot be manipulators of Type 1, as this assumes that the winner in the sincere profile V is neither in $\text{top}_2(v_1)$ nor in $\text{top}_2(v_2)$, against the assumption that these are the only two voters. Second, if both are Type 2 manipulators, then their manipulations are in favour of the same candidate, creating a Nash equilibrium when either plays insincerely.

Let us now treat the case of v_1 a Type 2 manipulator and v_2 of Type 1. In this case, v_1's second best candidate wins at the sincere profile and their top candidate comes second. As v_2 is a Type 1 manipulator, if he manipulates in favour of candidate x then x must beat after manipulation both of v_1's top two candidates. Therefore, v_1's manipulation will not change the outcome in case v_2 manipulates, giving a Nash equilibrium when just v_2 votes insincerely. $\square$

One of the notable characteristics of the Example 3 is that the manipulators have more than one manipulation in their strategy sets (in this case they have two each). We conjecture that if manipulators are restricted to have one manipulative action each (which is not an unreasonable assumption since considering more than one manipulative action is hard to expect from most voters) then a Nash equilibrium exists in any such voting manipulation game. We already showed this claim for 2-by-2 GS-games in Section 4.

5.4 2-Approval GS-Games under Sound Manipulations

We saw in Example 3 that if voters choose an 'irrational' set of strategies, then NE may not exist. Indeed, both voters in the example could easily realise that strategies $i_2'' = v_2^*[d;n]$ and $i_3'' = v_3^*[f;b]$ are safer strategies for the situation they are facing. For both players, promoting not only the candidate in favour of whom they are manipulating, but also and additional candidate m, might be a risky strategy, as m can inadvertently win (and she does in Example 3). Let us now state an important assumption that will be used in the following sections.

Definition 4. *Under k-Approval a Type 1 GS-manipulation $v_i[Y,X]$ of voter i in favour of $x \in X$, where $Y \subseteq top_k(v_i)$ and $X \subseteq C \setminus top_k(v_i)$, is called* sound *if all the alternatives which are moved up together with x are preferred to x, i.e., for all $x' \in X$, different from x, we have that $x' >_i x$.*

The reason for this definition is as follows. A voter is manipulating in favour of x which is the highest alternative which he can get by manipulating alone. However it may be possible that two or more voters manipulating in unison may secure a better outcome. Hence a voter may wish, just in case, to increase the score of an alternative different from x. However, increasing the scores of alternatives which are below x may end up in promoting them as the new winners. We therefore formulate the following assumption:

Assumption 1 (Soundness assumption)**.** *Let V be a profile and G be a GS-game for V under k-Approval voting rule. If voter i has a manipulation strategy in A_i in favour of an alternative x, then A_i also contains a sound manipulation strategy in favour of x.*

We are now ready to prove the following:

Theorem 6. *Any GS-game for 2-Approval voting whose strategy sets satisfy the soundness assumption has a Nash equilibrium.*

Proof. Let V be a profile and w be a winner for 2-Approval voting at V. Then w has a maximal score, say s. Let $Q \subset C$ be the set of candidates in favour of whom a GS-manipulation at V exists. Then all candidates from Q have scores s or $s-1$. Suppose that manipulations of Type 1 exist at V, and that q is the highest candidate from Q relative to $\sqsupset_V$ in favour of whom manipulation of Type 1 exists. Due to the soundness assumption, every voter who can manipulate in favour of q has a sound manipulation strategy in favour of her. Suppose one of them, say voter i, has the manipulation $v_i^* = v_i[y; q]$, in which he moves up only q. Then we claim that $V^1 = (V_{-i}, v_i^*)$ is a Nash equilibrium. At V^1 we have $q \sqsupset_{V^1} w \sqsupset_{V^1} \mathrm{top}(v_\ell)$ for any voter ℓ. Indeed, no manipulation of Type 2 will change the result at V^1. Such a manipulation for voter ℓ has the form $v_\ell^* = v_\ell[w; a]$, where a loses to $\mathrm{top}(v_\ell)$ in $(V_{-\ell}, v_\ell^*)$ (in favour of whom he manipulates), and, hence in V^1. And $\mathrm{top}(v_\ell)$ loses to q in V^1. Any manipulation of Type 1 in favour of $p \in Q$ will not change the result either since $q \sqsupset_V p$, hence p cannot overtake q in V^1.

Suppose now that any manipulator of Type 1 in favour of q in his sound manipulation also promotes another candidate whom he ranks higher than q, e.g., his sound manipulation is $v_i^* = v_i[xy; qr]$ with $r \succ_i q$. Then, as we showed before, no other manipulation can lead to winning of $p \in Q$ different from q. However, it may lead to winning of r in case a differentvoter j manipulates with $v_j^* = v_j[zt; pr]$ with $r \succ_j p$. In such a case $r \sqsupset_{V^2} q$, where $V^2 = (V^1_{-j}, v_j^*)$. Then V^2 may be a Nash equilibrium since both voters i and j will be satisfied. The only case when it is not a Nash equilibrium is when another voter, say k, can also manipulate in favour of q with $v_k^* = v_k[uv; qs]$, where $s \neq r$. Then someone else, say voter m, may be in position to vote $v_m^* = v_m[ef; pr]$ making r again the winner. After the first two manipulations, only q and r are in a position to become the winner (all other candidates will have not enough points). Therefore, all those voters who manipulated would not want to revert to their sincere votes, and eventually a Nash equilibrium will be reached.

It remains to consider the case when only Type 2 manipulators exist. In this case, by Lemma 2, they all manipulate in favour of the same candidate and, if one of them manipulates, all others do not have incentives to change their strategy. □

Moreover, as Example 3 shows, the soundness assumption is a necessary condition for the existence of Nash equilibria in 2-Approval GS-games.

5.5 3-Approval

To secure a Nash equilibrium for 2-Approval GS-games we had to state the soundness assumption. A similar result holds for 3-Approval, although under the stronger rationality assumption of restricting the set of manipulation strategies to those that take a minimal number of changes.

Definition 5. *A manipulation of Type 1 in favour of x under k-Approval is* minimal *if x is the only alternative which is moved up and to the kth position, while the alternative which formerly occupied kth position becomes not approved. A manipulation of Type 2 is* minimal *if the smallest number, say ℓ, of alternatives are moved down while the best ℓ not previously approved alternatives are moved up.*

In particular, a Type 1 minimal manipulation it is always sound. Observe that in a GS-game there may be several minimal manipulation available to the same player. We now formulate the following assumption:

Assumption 2 (Minimality assumption (MA))**.** *Let V be a profile and G be a GS-game for V. Then G satisfies the minimality assumption if for every $j = 1, 2, \ldots, n$ all GS-manipulations in the strategy set A_j are minimal.*

And prove the following lemma:

Lemma 3. *In any GS-game for 3-Approval voting, the set of candidates in favour of whom Type 2 manipulators can manipulate has cardinality at most 2.*

Proof. Suppose that the set of Type 2 manipulators at profile V is not empty and let Q be the set of candidates in favour of whom they can manipulate. Let w be the winner at V and let $q \in Q$ be such that $q \sqsupset_V q'$ for all $q' \in Q$. Let $p \in Q$ be the candidate such that $p \sqsupset_V q'$ for all $q' \in Q$ such that $q' \neq q$.

We claim that $Q \subseteq \{q, p\}$. Indeed, for any third alternative $r \in C$ to win as a result of Type 2 manipulation, the respective voter must demote w, q, p, which is impossible since r by Lemma 1 must be among the approved candidates. □

We are now ready to prove the following:

Theorem 7. *Any GS-game for 3-Approvalwhose strategy sets satisfy the minimality assumption has a Nash equilibrium.*

Proof. We can assume that the set of GS-manipulators $N(V, \text{3-App})$ is non-empty. Let w be a 3-Approval winner at the sincere profile V with the score $\text{sc}_3(w, V) = t$. By Lemma 1, we can partition the set of GS-manipulators into a set of *promoters*, i.e., the set of $j \in N(V, \text{3-App})$ such that $w \notin \text{top}_3(v_j)$, and a set of *demoters*, i.e., such voters v_ℓ for whom $w \in \text{top}_3(v_\ell)$.

Assume first that the set of promoters is non-empty. Let Q be the set of candidates in favour of whom promoters can manipulate, and let $\bar{p} \in Q$ such that $\bar{p} \sqsupset_V q$ for all $q \neq \bar{p}$. Let therefore i be a promoter in favour of $\bar{p}$, and let v_i^* be his GS-manipulation strategy. We now show that $V^* = (V_{-i}, v_i^*)$ is a NE. Observe that by minimality of v_i^* and by definition of $\bar{p}$ no other promoter can change the outcome of V^*. We can therefore focus on the set of demoters. Let j be a demoter and let v_j^* be its manipulation strategy in favour of candidate p.

By minimality assumption, v_j^* either removes only w from $\text{top}_3(v_j)$, or removes w together with a second candidate (and the third candidate in $\text{top}_3(v_j)$ by Lemma 1 must be then p). While the first case would not be a profitable deviation at V^* since the result of the election does not change, we need more attention in the second case since $\bar{p}$ could be that second demoted candidate by voter j.

The fact that voter j had to demote $\bar{p}$, due to the minimality assumption, means that $\bar{p} \sqsupset_V p$. In this case we will also have $\bar{p} \sqsupset_{V^{**}} p$, where $V^{**} = (V_{-j}^*, v_j^*)$. Hence $\bar{p}$ wins against p in V^{**}, and v_j^* is not a profitable deviation for j at V^*.

We can now assume that the set of promoters is empty, and that therefore all GS-manipulators in $N(V, 3\text{-App})$ are demoters. By Lemma 3 there are at most two candidates in favour of whom manipulation is possible. Let us denote them p_1 and p_2 with $p_1 \sqsupset_V p_2$.

The set of demoters can be partitioned into a set V_1 of GS-manipulators for p_1, whose minimal strategy is to lower the current winner w only, and a set V_2 of GS-manipulators for p_2, whose minimal strategy is to lower both w and p_1. Note that voters in V_2 has p_2 as their top candidate. If $V_2 = \emptyset$, all manipulators manipulate in favour of p_1 and Nash equilibrium obviously exist. Consider then the case in which V_1 and V_2 are both non-empty. Note that once a voter from V_2 manipulates, p_1 can no longer win no matter how other voters vote. For all pairs of candidates x, y different than w, p_1, or p_2, let

$$V_2^{x,y} = \{j \in V_2 \mid v_j[w, p_1; x, y] \in A_j\}.$$

A voter $v \in V_1$ with manipulation $v^* = v[w, x]$ ranks p_2 lower than x and this manipulation may serve as a countermanipulation strategy to manipulations of voters in $V_2^{x,y}$, making, under certain circumstances x to win instead of p_2, so we need to design a strategy profile in which such situations do not occur. Let therefore

$$V_1^x = \{j \in V_1 \mid v_j[w; x] \in A_j \text{ and } 3\text{-app}(V_{-\{i,j\}}, v_i^*, v_j^*) = x \text{ for some } i \in V_2^{x,y}\},$$

i.e., V_1^x is the set of voters who have a countermanipulation move in favour of x when x's score is being raised by a manipulator in favour of p_2 by some voter in V_2.

If there exists $j \in V_2$ and x, y such that $j \in V_2^{x,y}$ but both V_1^x and V_1^y are empty, then it is easy to see that $(V_{-j}, v_j[w, p_1; x, y])$ is a Nash equilibrium: voters in V_1 cannot change the outcome, and voters in V_2 are satisfied with having p_2 the winner. Suppose then that this is not the case, i.e., for each pair of candidates x, y, either $V_2^{x,y}$ is empty, or one of V_1^x and V_1^y are not empty. Pick one voter from each non-empty V_1^x – they are all distinct since each voter belongs to at most one V_1^x, having a single manipulation strategy. Without loss of generality let them be $J = \{1, \ldots, k\}$, and let $V^* = (V_{-J}, v_1^*, \ldots, v_k^*)$ be the profile in which all GS-manipulators in J play their manipulation strategies $v_1^*, \ldots, v_k^*$. Since p_1 is the winner in V^*, all voters in V_1 do not have incentives to deviate. Voters in V_2 also do not have incentives to deviate. For if any $j \in V_2^{x,y}$ manipulate in V^* the result would change in favour of either x or y, which by construction are less preferred by j than p_1. This concludes the proof. $\square$

5.6 4-Approval

In contrast with the results obtained in the previous sections, for 4-Approval the existence of Nash equilibria is no longer guaranteed, even if GS-manipulations are restricted to minimal ones.

Theorem 8. *There exists a game $G = (V, 4\text{-}App, (A_i)_{i \in N(V,4\text{-}App)})$, where for each player $i \in N(V, 4\text{-}App)$ the set A_i consists of i's truthful vote and i's minimal GS-manipulation, such that G has no Nash equilibrium.*

Proof. Let $\{u_1, u_2, u_3, v_1, v_2, v_3\}$ be a set of voters using 4-Approval to choose one among eight candidates $\{w, d_1, d_2, d_3, d, e, c, x\}$. Let the tie-breaking rule be $w \succ d_1 \succ d_2 \succ d_3 \succ c \succ d \succ e \succ x$. Let V be the profile in Table 4, where the four approved candidates are those to the left of the vertical line in each individual preference:

Voters	Preferences
Voter u_1	$w\, d_1\, e\, c \,\vert\, d\, x\, d_2\, d_3$
Voter u_2	$c\, d_2\, d\, d_3 \,\vert\, w\, e\, x\, d_1$
Voter u_3	$d\, d_3\, c\, d_2 \,\vert\, w\, d_1\, e\, x$
Voter v_1	$d_1\, w\, d\, d_3 \,\vert\, c\, d_2\, e\, x$
Voter v_2	$d_2\, w\, d_1\, x \,\vert\, c\, d\, d_3\, e$
Voter v_3	$d_3\, w\, d_2\, d_1 \,\vert\, d\, x\, e\, c$

Table 4: A Profile for the 4-Approval GS-game with no Nash Equilibrium

The scores of alternatives are as follows: all of w, d_1, d_2, d_3 get 4 points, c and d get 3 points, e and x gets 1 point. The winner at V is therefore w. The first three candidates cannot manipulate: u_1 ranks the winner w on top; u_2 and u_3 rank w just below the approval line and hence there is no candidate they prefer to w that can be promoted. Thus the set of GS-manipulators $N(V, 4\text{-App}) = \{v_1, v_2, v_3\}$. We restrict the set of strategies of voter j to the sincere strategy s_j and the minimal manipulation strategy i_j, hence $A_j = \{s_j, i_j\}$, where

- $i_1 = v_1[w; c]$ making d_1 winner;
- $i_2 = v_2[wd_1; cd]$ making d_2 winner;
- $i_3 = v_3[wd_2d_1; dxe]$ making d_3 winner.

Profile	Winner	Deviation	New Winner
(s_1, s_2, s_3)	w	1 switches to i_1	d_1
(i_1, s_2, s_3)	d_1	3 switches to i_3	d_3
(s_1, i_2, s_3)	d_2	1 switches to i_1	c
(s_1, s_2, i_3)	d_3	2 switches to i_2	d
(i_1, s_2, i_3)	d_3	2 switches to i_2	c
(s_1, i_2, i_3)	d	3 switches to s_3	d_2
(i_1, i_2, s_3)	c	2 switches to s_2	d_1
(i_1, i_2, i_3)	c	1 switches to s_1	d

Table 5: Deviations from strategy profiles.

There are eight strategy profiles in this game, and in Table 5 we indicate for each strategy profile which of the candidates is the winner and which of the voters have an incentive to change the strategy. At every strategy profile there is at least one player that prefers the winner of a different profile to the current one. Hence, there is no Nash equilibrium. □

6 Discussion

In this section we discuss the results obtained in Sections 4 and 5, as well as some of the crucial assumptions behind our model of voting manipulation games.

6.1 The game-theoretic complexity of manipulation

The results presented in this paper should be interpreted as providing a partial taxonomy of the game-theoretic complexity of manipulation. Starting from the assumption that GS-manipulators do not act in isolation, we showed in this paper that strategic voters might encounter situations of varying complexity depending on the strategic actions that are available to them.

In Section 4 we showed that all 2-by-2 voting manipulation games have at least one Nash equilibrium, but that in one case the two manipulators find themselves playing a coordination game with two competing equilibria. We observe further that not all 2-by-2 games are solvable by reasoning about dominated strategies, introducing an already high level of game-theoretic complexity even for these simple games. When a countermanipulator is added to the game the situation might even be more complex, with examples in which the two strategic voters have to play a game that does not have any Nash equilibria.

In Section 5 we focus on the case of an arbitrary number of GS-manipulators and a simple measure for game-theoretic complexity, namely whether the voting manipulation games admit at least one Nash equilibria. While the answer is negative in arbitrary games, we identify intuitive conditions on the strategy sets of voters that guarantee the existence of a Nash equilibria for 2 and 3-Approval voting (respectively, the soundness and the minimality assumption). This is however not the case for 4-Approval (and we conjecture it is so for any $k > 3$), where we show that even the strongest restriction on strategy sets does not guarantee the existence of a Nash equilibria, suggesting that the game-theoretic complexity of those games might be a sufficient deterrent from strategic voting. For the case of 1-Approval (plurality) we show that there always exists a Nash equilibria, without having to resort to additional assumptions on the strategy sets.

6.2 Model assumptions

Our results are based on a number of important assumptions, which we detail and defend in this section.

Boundedly rational voters. Unlike Myerson and Weber (1993), and follow-up papers, in our approach voters are boundedly rational and cannot see beyond manipulations in Gibbard-Satterthwaite sense (and countermanipulations to those for the special case of MC-games). To contemplate a countermanipulation, voters must have a higher degree of rationality—at least level-2 in the cognitive hierarchy model of Camerer et al. (2004)—so most of the time we assume that the game is played by Gibbard-Satterthwaite manipulators alone. In particular, we assume that voters vote sincerely if they cannot change the result, a similar assumption to the truth-bias of Obraztsova et al. (2013). It is extremely hard to estimate how many strategic voters are present in a given election, but the percentage of those who actually manipulated is easier to estimate. For instance, Kawai and Watanabe (2013) estimate the number of such voters, called misaligned, in Japanese elections between 2,5% and 5,5%. Moreover, Benjamin et al. (2013) show that preference misrepresentation is related to cognitive skills, and Choi et al. (2014) demonstrate that decision-making ability in laboratory experiments correlates strongly with socio-economic status and wealth. Therefore, it is reasonable to assume, as we do in the definition of GS-games, that only a small fraction of voters in an election would act strategically when given an opportunity to do so.

Complete information. The Gibbard-Satterthwaite theorem is proved under the assumption that voters know the sincere preferences of other voters. We make the same assumption when we investigate the interaction of Gibbard-Satterthwaite manipulators in our work.

Randomised strategies. Our work is positioned within an ordinal model of individual preferences, as we do not assume that voters can identify the utilities that they will enjoy when each candidate wins. If we assumed that voters were allowed to include randomised strategies in their strategy sets, then outcomes would become lotteries and voters must be able to compare them. However, within the ordinal model of preferences a voter with preferences $a > b > c$ is unable, for example, to compare the lottery that gives 50% chance of a and 50% chance of c

with the sure thing lottery that gives him 100% chance of b.

Tie-breaking. Our results clearly depend on the way ties are broken, and the alphabetic tie-breaking is the most popular method for doing this. It violates neutrality but preserves anonymity which is usually considered as a more important property. Randomised tie-breaking is not applicable in our framework as it will also lead to having to compare lotteries. It would be interesting, however, to see if our main results survive under other methods of deterministic tie-breaking.

7 Conclusions and Further Research

In this paper we start from the observation that realistic models of manipulation in voting must take into consideration the bounded rationality of voters. We propose a model in which the set of strategic voters is restricted to a set of classical Gibbard-Sattertwhaite manipulators, who reason about the effects of the possible manipulations of the other players in the resulting game.

We investigate the game-theoretic complexity of such games, in particular the existence of Nash equilibria, enlarging the class of voting rules studied to k-Approval rules, while the literature has typically focused on the special case of 1-Approval or Plurality. We fully characterise the variety of games that can be encountered when two manipulators (or one manipulator and one countermanipulator) play a game with one manipulation strategy each. The presence of a countermanipulator may induce games without Nash equilibria even for the simplest games, thus, when generalising to an arbitrary number of players and strategies, we concentrate on games played by Gibbard-Satterthwaite manipulators only. We provide initial in-depth results on the structure of such games, showing that while for Plurality they exhibit a fairly simple structure, for k-Approval with $k > 1$ GS-games are quite complicated to analyze, and it may therefore be difficult for a strategic voter to decide on their actions.

Many technical questions concerning GS-games remain open. The most immediate of them is to fully understand the role of minimality assumption in the proof for 3-Approval. Further afield, it would be interesting to extend our study to other voting rules, most notably Borda, and to identify reasonable restrictions on the manipulators' strategy spaces that lead to the existence of Nash equilibria or that allows for an easy computation of weakly dominant manipulations.

In conclusion, our results shows that Gibbard-Satterthwaite manipulators are at times forced to play rather complicated games, e.g., some in which manipulating might not be a dominant strategy, or in which coordination with the other manipulators would be required to elect a more preferred candidate or, even worse, games where no Nash equilibria exists. Our findings put therefore into question the notion of manipulability of a vote profile, which points at a very interesting direction for future research.

References

K. R. Apt and S. Simon. A classification of weakly acyclic games. *Theory and Decision*, 78(4): 501–524, 2015.

Y. Bachrach, N. Betzler, and P. Faliszewski. Probabilistic possible winner determination. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI-2010)*, 2010.

J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989.

D. J. Benjamin, S. A. Brown, and J. M. Shapiro. Who is behavioral? cognitive ability and anomalous preferences. *Journal of the European Economic Association*, 11:1231–1255, 2013.

C. F. Camerer, T.-H. Ho, and J.-K. Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, pages 861–898, 2004.

S. Choi, S. Kariv, W. Muller, and D. Silverman. Who is (more) rational? *American Economic Review*, 104:1518–1550, 2014.

F. De Sinopoli. Sophisticated voting and equilibrium refinements under plurality rule. *Social Choice and Welfare*, 17(4):655–672, 2000.

F. De Sinopoli, G. Iannantuoni, and C. Pimienta. On stable outcomes of approval, plurality, and negative plurality games. *Social Choice and Welfare*, 44(4):889–909, 2015.

Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic Commerce (EC-2010)*, 2010.

A. Dhillon and B. Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46:55–75, 2004.

E. Elkind, U. Grandi, F. Rossi, and A. Slinko. Gibbard-Satterthwaite games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI-2015)*, 2015a.

E. Elkind, E. Markakis, S. Obraztsova, and P. Skowron. Equilibria of plurality voting: Lazy and truth-biased voters. In M. Hoefer, editor, *Proceedings of the 8th International Symposium on Algorithmic Game Theory (SAGT-2015)*, 2015b.

E. Elkind, E. Markakis, S. Obraztsova, and P. Skowron. Complexity of finding equilibria of plurality voting under structured preferences. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2016)*, 2016.

E. Elkind, U. Grandi, F. Rossi, and A. Slinko. Cognitive hierarchy and voting manipulation. *Working paper available on request*, 2017.

R. Farquharson. *Theory of Voting*. Yale University Press, 1969.

T. J. Feddersen, I. Sened, and S. G. Wright. Rational voting and candidate entry under plurality rule. *American Journal of Political Science*, 34(4):1005–1016, 1990.

P. C. Fishburn. A strategic analysis of nonranked voting systems. *SIAM Journal on Applied Mathematics*, 35(3):488–495, 1978.

N. Fraser and D. Kilgour. Non-strict ordinal 2x2 games: A comprehensive computer-assisted analysis of the 726 possibilities. *Theory and Decision*, 20(2):99–121, 1986.

A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.

K. Kawai and Y. Watanabe. Inferring strategic voting. *American Economic Review*, 103:624–662, 2013.

N. S. Kukushkin. Acyclicity of improvements in finite game forms. *International Journal of Game Theory*, 40(1):147–177, 2011.

R. Meir. Iterative voting. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 4, pages 69–86. AI Access, 2017.

R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Iterative voting and acyclic games. *Artificial Intelligence*, 252:100–122, 2017.

H. Moulin. Dominance solvable voting schemes. *Econometrica*, 47:1337–1351, 1979.

R. Myerson and R. Weber. A theory of voting equilibria. *American Political Science Review*, 87(1):102–114, 1993.

S. Obraztsova and E. Elkind. Optimal manipulation of voting rules. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2012)*, 2012.

S. Obraztsova, E. Markakis, and D. R. M. Thompson. Plurality voting with truth-biased agents. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory*, 2013.

S. Obraztsova, O. Lev, E. Markakis, Z. Rabinovich, and J. S. Rosenschein. Beyond plurality: Truth-bias in binary scoring rules. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT)*, 2015.

S. Obraztsova, Z. Rabinovich, E. Elkind, M. Polukarov, and N. R. Jennings. Trembling hand equilibria of plurality voting. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-2015)*, 2016.

M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10 (2):187–217, 1975.

M. R. Sertel and M. R. Sanver. Strong equilibrium outcomes of voting games are the generalized Condorcet winners. *Social Choice and Welfare*, 22(2):331–347, 2004.

F. D. Sinopoli and G. Iannantuoni. On the generic strategic stability of Nash equilibria if voting is costly. *Economic Theory*, 25(2):477–486, 2005.

A. Slinko and S. White. Non-dictatorial social choice rules are safely manipulable? In *Proceedings of the 2nd International Workshop on Computational Social Choice (COMSOC-2008)*, 2008.

A. Slinko and S. White. Is it ever safe to vote strategically? *Social Choice and Welfare*, 43: 403–427, 2014.

D. R. M. Thompson, O. Lev, K. Leyton-Brown, and J. S. Rosenschein. Empirical analysis of plurality election equilibria. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2013)*, 2013.

L. Xia and V. Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI-2010)*, 2010.