

Games Gibbard–Satterthwaite Manipulators Play

Edith Elkind, Umberto Grandi, Francesca Rossi and Arkadii Slinko

Abstract

The Gibbard-Satterthwaite theorem implies the ubiquity of manipulators—the voters who could change the election outcome in their favor by unilaterally modifying their vote. In this paper, we ask what happens if a given profile admits several such voters. We model the strategic interactions among these voters, whom we call *Gibbard-Satterthwaite manipulators*, as a normal-form game. We classify the two-by-two games that can arise in this setting for some simple voting rules, and study the complexity of determining whether a given manipulative vote weakly dominates truth-telling, as well as existence of Nash equilibria.

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 3 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. This is problematic enough when such a manipulator is unique, as he then has a disproportional influence on the election outcome. However, the problem may be further exacerbated by the presence of multiple voters each of whom could get a more desirable outcome by lying—if he was the only one to do so. Indeed, if several such voters—we will call them Gibbard–Satterthwaite manipulators, or GS-manipulators—attempt to manipulate the election simultaneously in an uncoordinated fashion, the outcome may differ not just from the outcome under truthful voting, but also from the outcome that any of the GS-manipulators was trying to implement, due to complex interference among the different manipulative votes.

The goal of our paper is to develop a basic understanding of how multiple GS-manipulators can help or hurt each other when acting simultaneously, under some simple voting rules. We model the interaction among the GS-manipulators as a normal-form game, where the players are the GS-manipulators in a given profile, each player’s set of actions consists of his truthful vote and (a subset of) his manipulative votes, and players’ preferences over action profiles are determined by who wins in the resulting elections (see Section 4 for formal definitions). We call such games the *GS-games*. An important feature of these games, which distinguishes them from voting games that have been considered in prior literature (see Section 2), is that their set of players consists of the GS-manipulators only, and in their choice of actions the players are limited to strategies that would constitute a successful manipulation in the original profile. Thus, we aim at exploring the implications of the fact that many Gibbard–Satterthwaite manipulators can co-exist in a given profile, rather than investigating the full spectrum of strategic behaviors that may occur in voting scenarios.

We are interested in the properties of the normal-form games that arise in our model under common voting rules. Since the simplest non-trivial example of our framework involves two GS-manipulators, with one manipulative action each, we first focus on 2-by-2 games, and ask whether any such game can be represented as a GS-game. To answer this

question we develop a simple 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 GS-games. We then show that, while all 6 games can be obtained as GS-games under the Borda rule, for Plurality rule this is not the case.

We then move on to games with more than two players. Perhaps the most basic question one can ask in such games is whether a given manipulative vote is “safe”, i.e., whether the outcome obtained by submitting this vote is always at least as good as the outcome of truthful voting, no matter what actions other GS-manipulators choose. A vote with this property would be very attractive to a GS-manipulator, as he could then ignore the actions of other GS-manipulators. This question was originally investigated by Slinko and White (2008, 2013), in a setting where one considers a subset of GS-manipulators who all have identical preferences; it has been shown that finding a safe manipulation in this setting is easy for k -approval (Hazon and Elkind, 2010) and Borda (Ianovski et al., 2011). We show that in our model this problem is efficiently solvable for Plurality, and, with a mild restriction on players' strategy sets, for 2-approval; however, for 4-approval it becomes coNP-hard. We conjecture that the problem of safe manipulation is hard also for the case of 3-approval.

To further refine our understanding of GS-games, 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 natural conditions implying the existence of Nash equilibria for k -approval with $k = 2, 3$. However, we show that these conditions fail to ensure the existence of Nash equilibria for 4-approval.

In the remainder of the paper we omitted or sketched some proofs in the interest of space. Full versions of the proofs are available from the authors on request.

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., (Moulin, 1979; Myerson and Weber, 1993; Dhillon and Lockwood, 2004; Feddersen et al., 1990). The algorithmic aspects of such models have recently received some attention as well (Desmedt and Elkind, 2010; Xia and Conitzer, 2010; Thompson et al., 2013; Obraztsova et al., 2013). A closely related topic is voting dynamics, where players change their votes one by one in response to the current outcome (Meir et al., 2010; Lev and Rosenschein, 2012; Reijngoud and Endriss, 2012; Reyhaneh and Wilson, 2012). However, to the best of our knowledge, in all of these papers the set of players consists of all voters, i.e., a player is allowed to vote non-truthfully even if he would be unable or unwilling to manipulate the election on his own. Restricting the set of players to GS-manipulators in the original profile alters the problem substantially; for instance, it rules out “bad” Nash equilibria where all players vote for the same undesirable candidate.

Our work can be seen as an extension of the model of safe strategic voting proposed by Slinko and White (2008, 2013). However, unlike us, Slinko and White focus on a subset of GS-manipulators who (a) all have identical preferences and (b) choose between truth-telling and using a specific manipulative vote, and on the existence of weakly dominant non-truthful votes in this setting (such votes are called *safe strategic votes*). In contrast, we allow manipulators to have diverse preferences and to use strategy sets that contain more than one non-truthful vote, and the questions we are interested in include, but are not limited to, the existence of weakly dominant strategies.

3 Preliminaries

We consider n -voter elections over a candidate set $C = \{c_1, \dots, c_m\}$. 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 i ranks c_1 above c_2 ; if this is the case, we say that voter i *prefers* c_1 to 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 . Let $X = \{x_1, \dots, x_\ell\}$ and $Y = \{y_1, \dots, y_\ell\}$ be two disjoint sets of candidates. Then $v[X; Y]$ denotes the vote obtained by swapping x_j with y_j for $j = 1, \dots, \ell$ in the individual preference ordering v . If the sets X and Y are singletons, i.e., $X = \{x\}$, $Y = \{y\}$, we omit the curly braces, and simply write $v[x; y]$.

A *voting rule* is a mapping \mathcal{R} that, given a profile V , outputs a candidate $\mathcal{R}(V) \in C$. We consider the following voting rules in this paper.

- k -approval, $1 \leq k \leq m - 1$: under this rule, each candidate receives one point from each voter that ranks her in top k positions; 1-approval is also known as Plurality.
- Borda: under this rule, each candidate gets $m - j$ points from each voter that ranks her in position j .

Under both of these rules, the *score* of each candidate is the total number of points she receives, and the winner is the candidate with the highest score, with ties broken according to a fixed order $>$ over C . We denote the k -approval score of a candidate c in a profile V by $\text{sc}_k(c, V)$; c 's Borda score in V is denoted by $\text{sc}_B(c, V)$. We will sometimes denote the k -approval rule by k -App.

We say that two votes v and v' 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 . It is easy to see that v and v' are equivalent with respect to k -approval if and only if $\text{top}_k(v) = \text{top}_k(v')$, and v and v' are equivalent with respect to Borda if and only if $v = v'$.

4 The Model

Our goal is to investigate strategic situations that arise when one or more voters could improve the outcome of the election from their perspective by unilaterally changing their votes. We will now define such situations formally.

Definition 1. *We say that a voter i is a Gibbard–Satterthwaite manipulator, or a GS-manipulator, in a profile $V = (v_1, \dots, v_n)$ with respect to a voting rule \mathcal{R} if there exists a vote $v'_i \neq v_i$ such that i prefers $\mathcal{R}(V_{-i}, v'_i)$ to $\mathcal{R}(V)$. We denote the set of all GS-manipulators in a profile V with respect to a voting rule \mathcal{R} by $N(V, \mathcal{R})$. 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, moreover, 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 $\mathcal{R}(V_{-i}, v''_i)$.*

We say that i manipulates *in favor of* p by submitting a vote $v'_i \neq v_i$ if p is the winner in $\mathcal{R}(V_{-i}, v'_i)$. Note that a voter can manipulate in favor of several different candidates; however, in the definition of a GS-manipulation we require the voter to focus on his most preferred candidate among the ones he can make the election winner.

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., tuples of the form (a_1, \dots, a_n) , where $a_i \in A_i$ for all $i \in N$ (while one could define normal-form games 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).

For each preference profile V and each voting rule \mathcal{R} , we consider a family of normal-form games defined as follows. For each game in this family, the set of players N is the set of all GS-manipulators in V under \mathcal{R} . For each player i , his set of actions A_i consists of his truthful vote and a subset of his GS-manipulations; different choices of these subsets correspond to different games in the family. Finally, the preference relation of player i is determined by the outcome of \mathcal{R} on the preference profile that corresponds to a given action profile. Specifically, given an action profile $V^* = (v_i^*)_{i \in N}$, let $V[V^*] = (v'_1, \dots, v'_n)$ 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 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 $\mathcal{R}(V[V^{**}])$. We refer to any game of this form as a *GS-game* for a profile V and the voting rule \mathcal{R} , and 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., $(A_i)_{i \in N(V, \mathcal{R})}$. Thus, in what follows, we write $G = (V, \mathcal{R}, (A_i)_{i \in N(V, \mathcal{R})})$; 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 $V[V^*]$. We denote the set of all GS-profiles in a game G by $\mathcal{GSP}(G)$.

We emphasize that we allow the players to limit themselves to subsets of their GS-manipulations, rather than considering a single game where each player's set of actions consists of his truthful vote and *all* of his GS-manipulations. There are several reasons for that. First, the space of all GS-manipulations for a given voter can be very large, and a player may be unable 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). Second, the player may use a specific algorithm (e.g., that of Bartholdi et al., 1989) to find his GS-manipulation; in this case, his set of actions would consist of his truthful vote and the output of this algorithm. Third, the player may choose to ignore manipulations that are (weakly) dominated by other manipulations. Finally, a player may prefer not to change his vote beyond what is necessary to make his target candidate the election winner. One possible reason for this behavior is to keep his vote as close to his true preferences as possible (see Obraztsova and Elkind, 2012), or the fear of unintended consequences in the complex environment of the game.

5 2-by-2 GS 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. To address this question, 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) show 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 (a bidirectional edge

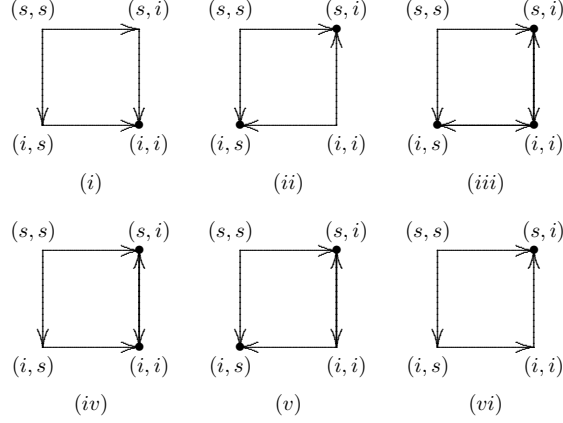


Figure 1: Diagrams for 2-by-2 GS-games

indicates indifference).

Now, let us focus on GS-games with 2 players and 2 actions per player. For each player, let s denote his truthful 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 can be represented by one of the 6 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 from Figure 1.

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

Example 3. Consider the GS-game for the preference profile (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; we can assume that their GS-manipulations are $v_1^* = v_1[a; b]$ and $v_2^* = v_2[b; a]$, 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 is *realizable* by a voting rule \mathcal{R} if there exists a preference profile V and a 2-by-2 game $G \in \mathcal{GS}(V, \mathcal{R})$ such that D is the diagram for G . Our next goal is to understand which diagrams are realizable by common voting rules.

Theorem 4. *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 GS-manipulators in V . Suppose that the Plurality winner in V is w , voter 1 manipulates in favor of a , and voter 2 manipulates in favor of b . Since 1 and 2 are GS-manipulators, we have $w \neq \text{top}(v_1)$ and $w \neq \text{top}(v_2)$: otherwise the two voters would not have an incentive to manipulate. We can therefore 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 $V^1 = (V_{-1}, v_1^*)$, $V^2 = (V_{-2}, v_2^*)$, $V^{12} = (V_{-2}^1, v_2^*)$.

The winner in V^{12} can be w , a , or b , and the case where it is w corresponds to diagram (ii). Further, if $a = b$, then a is the winner at V^1 , V^2 , and V^{12} , corresponding to diagram (iii). Thus, suppose the winner at V^{12} be 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 argued that diagrams (i) and (vi) are not realizable by Plurality. Example 3 shows how to realize diagram (ii). We now construct examples for the remaining three cases. Diagram (iii) can be realized in profile $V = (cabd, dabc, bacd)$ with ties broken alphabetically. The winner at V is b , and both the first and second players can manipulate in favor of a , which is therefore the winner at all manipulated profiles V^1 , V^2 , and V^{12} .

Consider now the profile $V = (dabce, ebacd, cabde)$ with tie-breaking order $b > a > c > d > e$. Candidate c is the winner in V , v_1 manipulates in favor of a , and v_2 in favor of b . If both players manipulate, the result is still b by the tie-breaking order. Since v_2 prefers b to a this example realizes diagram (iv). Diagram (v) can be obtained in profile $V = (dabc, abdc, cabd)$, and use $b > c > a > d$ as a tie-breaking rule. Candidate c is the winner in the initial profile, candidate a is winning at V^1 , and candidate b is winning at both V^2 and V^{12} , however this time v_2 prefers a to b and hence regrets her choice of manipulating. \square

In contrast, for Borda all six diagrams are realizable.

Theorem 5. *Diagrams (i)–(vi) are all realizable by Borda.*

Proof. We provide examples for each diagram in Figure 1:

Diagram (i). Let $V = (dbcae, dbcae, acebd, acebd)$, and assume that ties are broken according to $a > b > c > d > e$. The Borda winner in V is a . The first two voters are the GS-manipulators, and for each of them the vote $v^* = bdcea$ is a GS-manipulation in favor of c . Further, in (v^*, v^*, v_3, v_4) the Borda winner is b .

Diagram (ii). Let $V = (abc, abc, bca, cba)$, with $a > b > c$ as the tie-breaking order. The Borda winner at V is b , and each of the first two voters can manipulate in favor of a by submitting $v^* = acb$. However, in profile (v^*, v^*, v_3, v_4) the Borda winner is c , and both v_1 and v_2 prefer a to c .

Diagram (iii). Let $V = (cbad, dbac, abcd, abcd)$ with ties broken according to $a > b > c > d$. The Borda winner is a , and both v_1 and v_2 can manipulate in favor of b with strategies $v_1^* = cbad$ and $v_2^* = dbac$. If both manipulate, the winner is still b .

Diagram (iv). Let $V = (acdbe, bcdae)$, with $a > b > c > d > e$ as the tie-breaking order. The Borda winner is c . Let $v_1^* = adceb$ be v_1 's manipulation strategy in favor of a , and let $v_2^* = becda$ be v_2 's manipulation strategy in favor of b . The winner in profile (v_1^*, v_2^*) is a , hence realizing diagram (iv).

Diagram (v). Let $V = (dacbe, abcde, cbead, cebad)$ with ties broken according to $a > b > c > d > e$. The Borda winner in V is c . Voter v_1 has a manipulation strategy $v_1^* = adbec$ in favor of a , and voter v_2 has a manipulation strategy $v_2^* = badec$ in favor of b (note that v_2 cannot make her favorite candidate a win). The Borda winner in (v_1^*, v_2^*, v_3, v_4) is still b , and v_2 prefers a to b . Thus, this example realizes diagram (v).

Diagram (vi). Let $V = (abcd, cbda)$, with $a > b > c > d$ as the tie-breaking rule. The Borda winner is b . Let $v_1^* = adbc$ and $v_2^* = cdab$ be two manipulation strategies in favor of, respectively, a and c . The winner of (v_1^*, v_2^*) is d . Voter v_2 prefers d to a and does not regret manipulating, while v_1 prefers c to d , hence realizing diagram (vi). \square

6 Weak Dominance

In this section, we consider the complexity of checking whether a given GS-manipulation is always at least as good as truthful voting. This problem is captured by the game-theoretic notion of *weak dominance*.

Definition 6. Given a GS-game $G = (V, \mathcal{R}, (A_i)_{i \in N(V, \mathcal{R})})$ and two strategies $v_i^*, v_i^{**} \in A_i$, we say that v_i^* weakly dominates v_i^{**} if for every profile $V' \in \mathcal{GSP}(G)$ either $\mathcal{R}(V'_{-i}, v_i^*) = \mathcal{R}(V'_{-i}, v_i^{**})$ or voter i prefers $\mathcal{R}(V'_{-i}, v_i^*)$ to $\mathcal{R}(V'_{-i}, v_i^{**})$. We say that v_i^* is weakly dominant if v_i^* weakly dominates v_i^{**} for every $v_i^{**} \in A_i$.

In what follows, we will mostly be interested in determining whether a given GS-manipulation v_i^* weakly dominates the truthful vote v_i . A related problem is whether v_i^* is weakly dominant. However, in the context of GS-games the former problem is more relevant than the latter. Indeed, the main issue faced by a GS-manipulator is whether to manipulate or not, and if a certain vote can always ensure an outcome that is at least as good as that of truthful voting, this is a very strong incentive to use it, even if another non-truthful vote may be better in some situations. Note that our approach is consistent with the one taken in the study of safe manipulation (Slinko and White, 2008, 2013), where different manipulations are not compared to each other either.

In the rest of the paper, we limit our attention to k -approval rules. Thus, we will now state some definitions and observations that apply to this class of rules.

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 $X \subseteq \text{top}_k(v_i)$, $Y \subseteq C \setminus \text{top}_k(v_i)$. Consider a GS-manipulation $v_i[X; Y]$ of voter i in V under k -approval; we say that $v_i[X; Y]$ is *minimal* if for every other GS-manipulation v'_i of voter i it holds that $v'_i = v_i[X'; Y']$, where $|X'| \geq |X|$. That is, a GS-manipulation is minimal if it performs as few swaps as possible.

Consider a profile V such that the k -approval winner in V is w , and $\text{sc}_k(w, V) = t$. Let

$$\begin{aligned} S_1(V, k) &= \{c \in C \mid \text{sc}_k(c, V) = t, w > c\}, \\ S_2(V, k) &= \{c \in C \mid \text{sc}_k(c, V) = t - 1, c > w\}, \end{aligned}$$

and set $S(V, k) = S_1(V, k) \cup S_2(V, k)$. We say that a candidate c is *k -competitive* in V if $c \in S(V, k)$. Since every manipulative vote increases the score of each candidate by at most 1 and decreases the score of the winner by at most 1, it follows that under k -approval every GS-manipulation in V is in favor of some candidate $p \in S(V, k)$. In particular, if $S(V, k)$ is empty, there are no GS-manipulators. Further, we say that a candidate $c \in S(V, k)$ is *k -plausible* in V if there is a voter i such that $c \notin \text{top}_k(v_i)$ and $c \succ_i a$ for all $a \in S(V, k) \setminus \{\text{top}_k(v_i), c\}$; we denote the set of all k -plausible candidates in V by $S^+(V, k)$. Given a candidate set X , let $w(X, V, k)$ be the candidate with the highest k -approval score in V , with ties broken according to $>$. Set $p^*(V, k) = w(S(V, k), V, k)$, and if $S^+ \neq \emptyset$, set $p^+(V, k) = w(S^+(V, k), V, k)$. We omit V and k from the notation when they are clear from the context.

The following proposition characterizes the possible effects of a manipulative vote under k -approval.

Proposition 7. Let w be the k -approval winner in a profile V , let $v_i^* = v_i[X; Y]$, where $X \subseteq \text{top}_k(v_i)$, $Y \subseteq C \setminus \text{top}_k(v_i)$, let $V' = (V_{-i}, v_i^*)$, and let w' be the k -approval winner in V' . Then we have $w \in X$ or $w' \in Y$ or $w = w'$.

Proof. Suppose that $w \notin X$, $w' \notin Y$, but $w \neq w'$. Then $\text{sc}_k(w, V') \geq \text{sc}_k(w, V)$ and $\text{sc}_k(w', V') \leq \text{sc}_k(w', V)$. Since w beats w' in V , it follows that w also beats w' in V' , a contradiction. \square

With these definitions and results in hand, we can proceed with our analysis.

6.1 Plurality

Consider a GS-manipulator i in a profile V under Plurality, and let v_i^* be his GS-manipulation. Note that i does not rank the current winner first, since otherwise he would

not be a GS-manipulator. Let p_i be i 's most preferred candidate in $S \setminus \{\text{top}(v)\}$. It is clear that v_i^* is equivalent to $v_i[\text{top}(v_i); p_i]$ and p_i is the winner in (V_{-i}, v_i^*) , i.e., all GS-manipulations of voter i in V are equivalent to each other. Hence, there is essentially a unique GS-game that corresponds to V , namely, the one where $A_i = \{v_i, v_i[\text{top}(v_i); p_i]\}$ for each player i ; we will denote this game by $G_1^*(V)$. We will now characterize GS-manipulations that are weakly dominant in this game (as each player i has two strategies, his GS-manipulation v_i^* is weakly dominant if and only if it weakly dominates v_i).

Theorem 8. *Let v_i^* be the GS-manipulation of voter i in $G_1^*(V)$, and let $c = \text{top}(v_i)$. Then v_i^* is not weakly dominant in $G_1^*(V)$ if and only if $c \in S^+(V, 1)$.*

Proof. Let w be the Plurality winner in V . Suppose that there exists a GS-manipulator j such that $c \succ_j a$ for all $a \in S \setminus \{\text{top}(v_j), c\}$, and consider the profile V' where j submits his GS-manipulation (i.e., swaps $\text{top}(v_j)$ and c), whereas every other GS-manipulator votes truthfully. Clearly, c is the winner in V' , and moreover $\text{sc}_1(c, V') = \text{sc}_1(c, V) + 1$, and $\text{sc}_1(a, V') \leq \text{sc}_1(a, V)$ for all $a \in C \setminus \{c\}$. Suppose now that i changes his vote to v_i^* ; denote the resulting profile by V'' . We have $\text{sc}_1(c, V'') = \text{sc}_1(c, V)$ and $\text{sc}_1(w, V'') = \text{sc}_1(w, V)$, so c loses to w in V'' . But $c = \text{top}(v_i)$, hence v_i^* is not weakly dominant.

Conversely, suppose that $c \notin S^+$. Then for any profile $V' = (v'_1, \dots, v'_n) \in \mathcal{GSP}(G_1^*(V))$ such that $v'_i = v_i$ we have $\text{sc}_1(c, V') \leq \text{sc}_1(c, V)$. On the other hand, since no GS-manipulator ranks w first, we have $\text{sc}_1(w, V') = \text{sc}_1(w, V)$. Thus, c is not the Plurality winner in V' . Let w' be the Plurality winner in V' ; note that $w' \in \{w\} \cup S$. Let $V'' = (V'_{-i}, v_i^*)$. Suppose that $v_i^* = v[c; p]$. By Proposition 7 the Plurality winner in V'' is either p or w' . Since i weakly prefers p to every other candidate in $\{w\} \cup S$, it follows that v_i^* is weakly dominant. \square

Theorem 8 illustrates that every GS-game for Plurality is essentially a coordination game: the GS-manipulators have to coordinate in order to ensure that their efforts do not cancel out. We will now see that for k -approval with $k > 1$ the situation is more complicated.

6.2 2-Approval: An Algorithm

In this section, we focus on the 2-approval rule. Just as for Plurality, we first consider a GS-manipulator i in a profile V , and describe the GS-manipulations that are available to him. Let w be the 2-approval winner in V with score t . Again, since i is a GS-manipulator, he does not rank w first.

If i ranks w second, the only way for him to improve the outcome is to vote so that $p = \text{top}(v_i)$ wins. This is possible if and only if (1) $p = p^*$, and (2) there exists a candidate $c \neq p, w$ such that (2a) $\text{sc}_2(c) \leq t - 3$, or (2b) $\text{sc}_2(c) = t - 2$ and either $p \in S_1$ or $p \in S_2$, $p > c$, or (2c) $\text{sc}_2(c) = t - 1$, $p \in S_1$, and $p > c$. Note that condition (1) implies that if there are several GS-manipulators who rank w second, they all rank the same candidate p^* first, and have the same set of GS-manipulations. Moreover, every GS-manipulation for voters of this type is equivalent to swapping w with some candidate c that satisfies one of the conditions (2a)–(2c). We will therefore refer to such voters as *demoters*.

Finally, suppose that $w \notin \text{top}_2(v_i)$. Then i cannot lower the score of w by changing his vote. However, he can raise the scores of some candidates in $C \setminus \text{top}_2(v_i)$ by moving these candidates into top two positions; note that i can do that for two candidates simultaneously. Therefore, we refer to GS-manipulators who do not rank w in top two positions as *promoters*. Clearly, a promoter's vote is a GS-manipulation if and only if one of the promoted candidates is i 's most preferred candidate in $S \setminus \text{top}_2(v_i)$ (let us denote this candidate by p) and the other promoted candidate p' is such that $\text{sc}_2(p', V) < \text{sc}_2(p, V)$ or $\text{sc}_2(p', V) = \text{sc}_2(p, V)$ and $p > p'$ (i.e., promoting p' does not prevent p from becoming a winner).

Given a profile G , let $G_2^*(V)$ be the GS-game where each player's set of actions consists of his truthful vote and all his minimal manipulations. The argument above shows that

for demoters the minimality assumption imposes no additional constraints, whereas for promoters it excludes manipulations where two candidates are swapped into top two positions simultaneously. We can now state the main result of this section.

Theorem 9. *There is a polynomial-time algorithm for checking whether a given GS-manipulation v'_i of voter i weakly dominates his truthful vote v_i in $G_2^*(V)$.*

Proof sketch. Let w be the winner in V , and let $t = \text{sc}_2(w, V)$. Consider a GS-manipulator i and his GS-manipulation v'_i . Let q be the number of demoters in V .

Suppose first that i is a demoter, and let $v_i[w; c]$ be his GS-manipulation. Then it can be shown that $v_i[w; c]$ weakly dominates v_i if and only if

- (i) $S_1 \neq \emptyset$, $p^* > c$, and $\text{sc}_2(c, V) + q \leq t$, or
- (ii) $S_1 \neq \emptyset$, $c > p^*$, and $\text{sc}_2(c, V) + q \leq t - 1$, or
- (iii) $S_2 = \emptyset$, $p^* > c$, and $\text{sc}_2(c, V) + q \leq t - 1$, or
- (iv) $S_2 = \emptyset$, $c > p^*$, and $\text{sc}_2(c, V) + q \leq t - 2$.

For promoters, the analysis is considerably more complicated; we omit it due to space constraints. \square

6.3 4-Approval: A Hardness Proof

We will now show that for 4-approval there is a way to select the players' action sets so that in the resulting game the problem of checking whether a given GS-manipulation weakly dominates truthtelling is coNP-complete.

We reduce from the classic NP-complete problem EXACT COVER BY 3-SETS (X3C). An instance of this problem is given by a ground set $\Gamma = \{g_1, \dots, g_{3\nu}\}$ and a collection $\Sigma = \{\sigma_1, \dots, \sigma_\mu\}$ of 3-element subsets of Γ . It is a “yes”-instance if there is a subcollection $\Sigma' \subseteq \Sigma$ with $|\Sigma'| = \nu$ such that $\cup_{\sigma \in \Sigma'} \sigma = \Gamma$, and a “no”-instance otherwise.

Theorem 10. *The problem of deciding whether a given strategy v'_i weakly dominates truthtelling in a GS-game $(V, 4\text{-App}, (A_i)_{i \in N(V, 4\text{-App})})$ is coNP-complete.*

Proof sketch. It is immediate that this problem is in coNP. To prove coNP-hardness, consider an instance $I = (\Gamma, \Sigma)$ of X3C with $|\Gamma| = 3\nu$, $|\Sigma| = \mu$. We can assume without loss of generality that $\sigma_1 = \{g_1, g_2, g_3\}$ and no other set in Σ contains g_1 ; if this is not the case, we can modify I by adding three new elements and a single set containing them. We will now construct an instance of our problem. In what follows, when writing $X \succ Y$ in the description of an order \succ , we mean that all elements of X are ranked above all elements of Y , but the order of elements within X and within Y can be arbitrary.

There is a set of candidates $C' = \{c_1, \dots, c_{3\nu}\}$ that correspond to elements of Γ , three special candidates w, p , and c , and a set of dummy candidates $D = \bigcup_{i=0}^{\mu} D_i \cup D'$, where $|D_i| = 4$ for $i = 0, \dots, \mu$. We define the order \succ by setting $w > c > p > c_1 > \dots > c_{3\nu} > D$.

For each set $\sigma_i \in \Sigma$ we construct a vote v_i of the form

$$D_i \succ \sigma_i \succ c \succ \dots,$$

where candidates in σ_i are ordered according to \succ , and set

$$v_0 = D_0 \succ p \succ c_1 \succ c \succ C' \setminus \{c_1\} \succ \dots$$

We also construct additional voters so that in the resulting profile V we have $\text{sc}_4(w, V) = \text{sc}_4(p, V) = \nu + 1$, $\text{sc}_4(c_i, V) = \nu + 1$ for all $i = 1, \dots, 3\nu$, $\text{sc}_4(c, V) = 1$, the 4-approval

score of each candidate in D is at most 1, and the only GS-manipulators in V are voters $0, 1, \dots, \mu$.

Note that w is the 4-approval winner in V , and $S = C' \cup \{p\}$. We now define a GS-game for this profile by constructing the players' sets of actions as follows. Let $D_0 = \{d_1, d_2, d_3, d_4\}$. Then $v_0^* = v_0[\{d_1, d_2\}; \{p, c\}]$ is a GS-manipulation for voter 0, which makes p the winner with $\nu + 2$ points. Similarly, for each $i = 1, \dots, \mu$ the vote $v_i^* = v_i[D_i; \sigma_i \cup \{c\}]$ is a GS-manipulation which makes i 's top candidate in σ_i the winner with $\nu + 2$ points (note that i orders σ_i in the same way $>$ does, so tie-breaking favors i 's most preferred candidate in σ_i). We set $A_i = \{v_i, v_i^*\}$ for $i = 0, \dots, \mu$. This completes the description of our game. Clearly, we can construct the profile V and the players' sets of strategies in polynomial time given I . It can be shown that v_0^* weakly dominates v_0 if and only if we started with a “no”-instance of X3C; we omit the proof. \square

7 Nash Equilibrium

In this section, we study the existence of Nash equilibria in GS-games for k -approval with $k = 1, 2, 3, 4$. In what follows, when considering a GS-game, we assume that its set of GS-manipulators is not empty, since otherwise a Nash equilibrium exists trivially.

We will first show that for Plurality, a Nash equilibrium always exists.

Theorem 11. *For any profile V the game $G_1^*(V)$ has a Nash equilibrium.*

Proof. Fix a profile V . If the set of GS-manipulators in V is not empty, we have $S^+ \neq \emptyset$, so p^+ is well-defined. Let i be some voter whose GS-manipulation is in favor of p^+ , and let v_i^* be his GS-manipulation. It is easy to check that (V_{-i}, v_i^*) is a Nash equilibrium with winner p^+ . \square

A similar argument proves the existence of Nash equilibria for 2-approval, as long as every manipulator's action set contains at least one minimal GS-manipulation.

Theorem 12. *Any game $G = (V, 2\text{-App}, (A_i)_{i \in N(V, 2\text{-App})})$, where for each $i \in N(V, 2\text{-App})$ the set A_i contains some minimal manipulation of voter i , has a Nash equilibrium.*

Proof. We can assume that $N(V, 2\text{-App}) \neq \emptyset$. Assume first that the set of promoters in G is not empty. Let \bar{p} be the plausible candidate with the highest 3-approval score or highest tie-breaking rank such that there exists a promoter that can manipulate in its favor. Let therefore i be a promoter in favor of \bar{p} , and let v_i^* be her minimal GS-manipulation strategy. $V^* = (V_{-i}, v_i^*)$ is a Nash equilibrium with winner \bar{p} . By construction of \bar{p} , no promoter can change the outcome of V^* by deviating from her sincere strategy. Moreover, demoters can only decrease the score of w , and thus cannot change the outcome at V^* . Otherwise, all GS-manipulators are demoters. Then a profile where a demoter, say j , submits some GS-manipulation and everyone else votes truthfully is a Nash equilibrium with winner $\text{top}(v_j)$. \square

The assumption of minimality in Theorem 12 can be relaxed: for instance, we can allow a promoter to swap a second candidate into the top 2 positions, as long as he prefers this candidate to the one he is trying to make the winner (we omit the formal statement of this result and the proof). However, it is not clear if it can be eliminated altogether.

A similar result holds for 3-approval, though under stronger assumptions. A manipulation strategy $v_i[X; Y]$ is called *greedy* if it is minimal and if all other minimal strategies $v_i[X; Y']$ are such that i 's least preferred candidate in Y' is less preferred than i 's least preferred candidate in Y . Clearly, there is a unique greedy manipulation for each manipulator i .

Theorem 13. *Any game $G = (V, 3\text{-App}, (A_i)_{i \in N(V, 3\text{-App})})$, where for each player $i \in N(V, 3\text{-App})$ the set A_i consists of i 's truthful vote and a greedy GS-manipulation, has a Nash equilibrium.*

Proof. We can assume that the set of GS-manipulators $N(V, 3\text{-App})$ is non-empty, and let $w = 3\text{-App}(V)$ with $\text{sc}_3(w, V) = t$. As in the case of 2-approval (see Section 6.2), we can partition the set of GS-manipulators in a set of *promoters*, i.e., the set of $j \in N(V, 3\text{-App})$ such that $w \notin \text{top}_3(v_j)$, and a set of *demoters*, i.e., with $w \in \text{top}_3(v_j)$.

Assume first that the set of promoters is non-empty, and let \bar{p} be the candidate with the highest 3-approval score or highest tie-breaking rank such that there exists a promoter that can manipulate in its favor. Let therefore i be a promoter in favor of \bar{p} , and let v_i^* be her 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 favor 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. 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. Let therefore $v_j^* = v_j[w, \bar{p}; p, a]$ where a is an irrelevant candidate. Observe that by minimality of v_j^* we must have that either $\text{sc}_3(\bar{p}, V) > \text{sc}_3(p, V)$ or if the score is the same then \bar{p} is higher in the tie-breaking order than p . Let $V' = (V_{-j}^*, v_j^*)$. Observe that, by minimality again, $\text{sc}_3(V', \bar{p}) = \text{sc}_3(V, \bar{p})$, and moreover that $\text{sc}_3(V', p) = \text{sc}_3(V, p)$, this time because v_j^* is a demoter strategy. 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. We first show that there are at most two plausible candidates in S^+ . Assume for the sake of contradiction that p_1, p_2 and p_3 are three distinct alternatives in S^+ . Let p_3 be the alternatives with minimum 3-score or that sits lower in the tie-breaking order (and assume p_1 and p_2 are ordered in the same way). A demoter strategy in favor of alternative p_3 requires a GS-manipulator j to decrease the score of w , together with the score of p_1 and p_2 . However, with 3-approval this implies that $p_3 \notin \text{top}_3(v_j)$ in contradiction with the hypothesis that p_3 is preferred to w . Hence 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.

We first construct a NE in case $V_2 = \emptyset$. Let i be any voter in V_1 , and let v_i^* be her greedy manipulation strategy. It is easy to see that $V^* = (V_{-i}, v_i^*)$ is a NE, since all GS-manipulators, i.e., all demoters in V_1 , cannot change the winner in V^* – they can only increase the score of p_1 by minimal manipulation. Consider then the case in which V_1 and V_2 are not empty. For all pairs of candidates x, y different than w, p_1 , or p_2 , let $V_2^{x,y} = \{j \in V_2 \mid v_j[w, p_1; x, y] \in A_j\}$. Voters in V_1 who rank p_2 lower than x may have a countermanipulation strategy to voters in $V_2^{x,y}$, so we need to design a NE in which such deviations are not possible. Let therefore $V_1^x = \{j \in V_1 \mid v_j[w, x] \in A_j \text{ and } 3\text{-app}(V_{-\{i,j\}}, v_i^*, v_j^*) = x \text{ for some } j \in V_2^{x,-}\}$, i.e., V_1^x is the set of voters who have a countermanipulation move when x 's score is being raised to favor p_2 by some voters 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 NE: voters in V_1 cannot change the outcome, and voters in V_2 are satisfied with having p_2 the winner. Suppose then 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 $J = \{v_1, \dots, v_k\}$, and let $V^* = (V_{-J}, v_1^*, \dots, v_k^*)$ be the profile in which all GS-manipulators

in J play the manipulation strategy. 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 favor of either x or y , which by construction are less preferred by j than p_1 . This concludes the proof. \square

In contrast, for 4-approval the existence of Nash equilibria is no longer guaranteed, even if manipulation is restricted to greedy GS-manipulation strategies. We omit the proof in the interest of space.

Theorem 14. *There exists a game for 4-approval $G = (V, 4\text{-App}, (A_i)_{i \in N(V, 4\text{-App})})$, where for each player $i \in N(V, 4\text{-App})$ the set A_i consists of i 's truthful vote and a greedy GS-manipulation, such that G has no Nash equilibrium.*

8 Conclusions

We have initiated the study of games played by GS-manipulators. We have shown that for Plurality these games exhibit a fairly simple structure; however, for Borda and k -approval with $k > 1$ GS-games are quite complicated, and it may therefore be difficult for the players to coordinate their actions. Thus, these games may require extensive communication and a non-trivial computational effort, which may serve as a barrier against manipulation.

Many questions concerning GS-games remain open. The most immediate of them is to fully understand the role of minimality assumptions in our proofs. Further afield, it would be interesting to extend our study to other voting rules, and to identify reasonable restrictions on the manipulators' strategy spaces that lead to existence and uniqueness of Nash equilibria, and make it easy to compute manipulations that weakly dominate truth-telling.

Acknowledgements This research was initiated during the Institute for Mathematical Sciences (Singapore) program on Algorithmic Game Theory and Computational Social Choice. Edith Elkind and Arkadii Slinko gratefully acknowledge the financial support provided by IMS, as well as by the National Research Foundation (Singapore) grant NRF2009-08. The work of Umberto Grandi and Francesca Rossi was partially supported by the strategic project "KIDNEY - Incorporating patients' preferences in kidney transplant decision protocols" funded by the University of Padova.

References

- Y. Bachrach, N. Betzler, and P. Faliszewski. Probabilistic possible winner determination. In *AAAI'10*, pages 697–702, 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.
- Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *ACM EC'10*, pages 347–356, 2010.
- A. Dhillon and B. Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46:55–75, 2004.
- 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.

- 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.
- N. Hazon and E. Elkind. Complexity of safe strategic voting. In *SAGT'10*, pages 210–221, 2010.
- E. Ianovski, L. Yu, E. Elkind, and M. Wilson. The complexity of safe manipulation under scoring rules. In *IJCAI'11*, pages 246–251, 2011.
- O. Lev and J. S. Rosenschein. Convergence of iterative voting. In *AAMAS'12*, pages 611–618, 2012.
- R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria of plurality voting. In *AAAI'10*, pages 823–828, 2010.
- 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 *AAMAS'12*, pages 619–626, 2012.
- S. Obraztsova, E. Markakis, and D. R. M. Thompson. Plurality voting with truth-biased agents. In *SAGT'13*, 2013.
- A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *AAMAS'12*, pages 635–644, 2012.
- R. Reyhaneh and M. Wilson. Best reply dynamics for scoring rules. In *ECAI'12*, pages 672–677, 2012.
- 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.
- A. Slinko and S. White. Non-dictatorial social choice rules are safely manipulable? In *COMSOC'08*, pages 403–414, 2008.
- A. Slinko and S. White. Is it ever safe to vote strategically? *Social Choice and Welfare*, 2013. Online first.
- D. R. M. Thompson, O. Lev, K. Leyton-Brown, and J. S. Rosenschein. Empirical analysis of plurality election equilibria. In *AAMAS'13*, pages 391–398, 2013.
- L. Xia and V. Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *AAAI'10*, pages 805–810, 2010.

Edith Elkind
 Department of Computer Science
 University of Oxford
 United Kingdom
 Email: eelkind@gmail.com

Umberto Grandi
Department of Mathematics
University of Padova
Italy
Email: umberto.uni@gmail.com

Francesca Rossi
Department of Mathematics
University of Padova
Italy
Email: frossi@math.unipd.it

Arkadii Slinko
Department of Mathematics
University of Auckland
New Zealand
Email: a.slinko@auckland.ac.nz