

Cognitive Hierarchy and Voting Manipulation in k -Approval Voting*

Edith Elkind¹, Umberto Grandi^{†2}, Francesca Rossi³, and Arkadii Slinko⁴

¹*University of Oxford, United Kingdom*

²*University of Toulouse, France*

³*IBM T.J. Watson Research Lab, New York, USA*

⁴*The University of Auckland, New Zealand*

Abstract

By the Gibbard–Satterthwaite theorem, every reasonable voting rule for three or more alternatives is susceptible to manipulation: there exist elections where one or more voters can change the election outcome in their favour by unilaterally modifying their vote. When a given election admits several such voters, strategic voting becomes a game among potential manipulators: a manipulative vote that leads to a better outcome when other voters are truthful may lead to disastrous results when other voters choose to manipulate as well. We consider this situation from the perspective of a boundedly rational voter, using an appropriately adapted cognitive hierarchy framework to model voters' limitations. We investigate the complexity of algorithmic questions that such a voter faces when deciding on whether to manipulate. We focus on k -approval voting rules, with $k \geq 1$. We provide polynomial-time algorithms for $k = 1, 2$ and hardness results for $k \geq 4$ (NP and co-NP), supporting the claim that strategic voting, albeit ubiquitous in collective decision making, is computationally hard if the manipulators try to reason about each others' actions.

JEL classification: D72.

Keywords: strategic voting, bounded rationality, computational complexity

1 Introduction

Imagine that you and your friends are choosing a restaurant to go to for dinner. Everybody is asked to name their two most preferred cuisines, and the cuisine named most frequently will

*Some of the results in this paper were presented at the 24th International Joint Conference on Artificial Intelligence (IJCAI) (Elkind et al., 2015a). We are grateful to the IJCAI-15 and the MSS reviewers, as well as to the audiences of the Workshop on Iterative Voting and Voting Games in Padova in 2014, the 5th International Workshop on Computational Social Choice (COMSOC-2014) in Pittsburgh, and to the participants of seminar talks in Toulouse in 2014 and in Dagstuhl in 2015 for their useful feedback.

[†]Corresponding author. Email addresses: elkind@cs.ox.ac.uk (E. Elkind), umberto.grandi@irit.fr (U. Grandi), Francesca.Rossi2@ibm.com (F. Rossi), a.slinko@auckland.ac.nz (A. Slinko).

be selected (this voting rule is known as 2-approval). Your favourite cuisine is Japanese and your second most preferred cuisine is Indian. Indian is quite popular among your friends, and you know that if you name it among your favourite two cuisines, it will be selected. On the other hand, you also know that only a few of your friends like Chinese food. Will you vote for Japanese and Chinese to give Japanese cuisine a chance?

This example illustrates that group decision-making is a complex process that involves agents reasoning about other agents' preferences. Individual decision-makers would like to influence the final decision in a way that is beneficial to them, and hence they may be strategic in communicating their individual preferences. Indeed, it is essentially impossible to eliminate strategic behaviour by a clever choice of the voting rule: the groundbreaking result of Gibbard (1973) and Satterthwaite (1975) states that, under any onto and non-dictatorial social choice rule, there exist situations where a voter can achieve a better outcome by casting a strategic vote rather than the sincere one, provided that everyone else votes sincerely; in what follows, we will refer to voters that can benefit from voting strategically when others remain sincere as *Gibbard–Satterthwaite manipulators* or simply *GS-manipulators*.

The Gibbard–Satterthwaite theorem alerts us that strategic behaviour of voters cannot be ignored, but it does not tell us under which circumstances it actually happens. Of course, if there is just a single GS-manipulator at a given profile, and he¹ is fully aware of other voters' preferences, it is rational for him to manipulate. However, even in this case this voter may prefer to vote truthfully, simply because he may assign a high value to announcing his true preferences; we call such voters *ideological*, or *non-strategic*. Moreover, if there are two or more GS-manipulators, it is no longer easy for them to make up their mind in favour of manipulation: while the Gibbard–Satterthwaite theorem tells us that each of these voters would benefit from voting strategically assuming that all other voters remain truthful, it does not offer any predictions when several voters may be able to manipulate simultaneously.

voter 1	voter 2	voter 3	voter 4
<i>b</i>	<i>b</i>	<i>a</i>	<i>c</i>
<i>c</i>	<i>c</i>	<i>c</i>	<i>b</i>
<i>a</i>	<i>a</i>	<i>b</i>	<i>a</i>

Table 1: A preference profile. In each column, the alternatives are ranked from the most preferred (top) to the least preferred (bottom).

Example 1. Suppose four people are to choose among three alternatives a , b , and c , by means of 2-approval, with ties broken according to the order $a > b > c$. That is, each voter specifies two approved alternatives, and the winning alternative is selected among the alternatives that received the largest number of approvals by means of the tie-breaking rule. Let the profile of sincere preferences be as in Table 1. There are two voters who prefer b to c to a , one voter who prefers a to c to b , and one voter who prefers c to b to a . If everyone votes sincerely, then c gets 4 approvals, b gets 3 approvals and a gets 1 point, so c is elected. We observe that voters 1 and 2 are Gibbard–Satterthwaite manipulators: each of them can make b the winner by approving $\{a, b\}$, *ceteris paribus*.² Let us consider this game from the first voter's perspective, assuming

¹We use 'he' to refer to voters and 'she' to refer to candidates.

²In the formal model that will be used throughout the paper, voters report their preference rankings; under

that he is strategic; let A_i denote the strategy set of voter i , $i = 1, 2, 3, 4$. The strategy set of voter 1 can then be assumed to be $A_1 = \{\{b, c\}, \{b, a\}\}$. Voter 1 has a good reason to believe that voters 3 and 4 will vote sincerely, as voter 3 cannot achieve an outcome that he would prefer to the current outcome and voter 4 is fully satisfied.

Case 1. If voter 1 believes that voter 2 is ideological, then he is analysing the game where $A_2 = \{\{b, c\}\}$, $A_3 = \{\{a, c\}\}$ and $A_4 = \{\{c, b\}\}$. In this case he just votes $\{b, a\}$ and expects b to become the winner.

Case 2. Suppose now that voter 1 believes that voter 2 is also strategic. Now voter 1 has to analyse the game with $A_1 = A_2 = \{\{b, c\}, \{b, a\}\}$, $A_3 = \{\{a, c\}\}$ and $A_4 = \{\{c, b\}\}$. If either one of the strategic players—voter 1 or voter 2—manipulates and another stays sincere, b will be the winner. However, if they both manipulate, their least preferred alternative a will become the winner. Thus, in this case voter 1’s manipulative strategy does not dominate his sincere vote, and if voter 1 is risk-averse, he should refrain from manipulating.

The overarching goal of this paper is to investigate the complexity of the decision that a strategic participant of a voting game must make in the presence of other strategic voters. To obtain a realistic model, we assume that a voter is a boundedly rational agent, most notably in its representation of other agents’ strategic abilities, and we suggest a model for such a behaviour. We also assume that voters have bounded computational abilities, and use the toolbox of computational social choice (Brandt et al., 2015) to study the computational complexity of the problems that strategic voters face. Thus, our work merges two views of bounded rationality: the strategic one, which is inspired by the game theory and economics literature, and the algorithmic one, which is more common in the computer science literature and goes back to Bartholdi et al. (1989) and Bartholdi and Orlin (1991)). As this is the first study to combine these two views of bounded rationality in the context of voting, it is natural to start by investigating simple voting rules. Therefore, throughout the paper we focus on the family of k -approval voting rules for $k \geq 1$.

1.1 Voting as a Strategic Game

A prevalent approach up to date has been to view voting as a strategic game among the voters, and use various game-theoretic solution concepts to predict the outcomes (see Section 1.3 for a survey of related work). The most common such concept is Nash equilibrium, which is defined as a combination of strategies, one for each player, such that each player’s strategy is a best response to other players’ strategies. In these terms, the Gibbard–Satterthwaite theorem says that under every reasonable voting rule there are situations where truthful voting is not a Nash equilibrium. For example, the game analysed in Example 1 (Case 2) has two Nash equilibria: in the first one, voter 1 manipulates and voter 2 remains truthful, and in the second one the roles are reversed. If we increase the strategy sets, allowing non-manipulators to act, we obtain further Nash equilibria, but truthful voting is never one of them.

The principle that players can always be expected to choose equilibrium strategies is not universally applicable. Specifically, if players have enough experience with the game they are playing (or with similar games), both theory and experimental results suggest that players are often able to learn equilibrium strategies (Fudenberg and Levine, 1998). However, it is also well-known since the early work of Shapley (1964) that learning dynamics may fail to converge to an

2-approval, voters can indicate that they approve $\{a, b\}$ by reporting $a \succ b \succ c$ or $b \succ a \succ c$. For readability, in this example we use approval ballots.

equilibrium. Moreover, in many applications—and voting is one of them—players’ interactions have only imperfect precedents, or none at all, making learning hard if not impossible.³ If equilibrium is justified in such applications, it must be via strategic thinking of players rather than learning. In reality, even in simple games the required reasoning is often too complex for such a justification of equilibrium to be behaviourally plausible (Harsanyi and Selten, 1988; Brandenburger, 1992). This is fully applicable to voting, where such reasoning, beyond very simple profiles, is impossible because of the number of voters involved.

In fact, a number of recent experimental and empirical studies suggest that players’ responses in games often deviate systematically from equilibrium strategies, and are better explained by the structural non-equilibrium level- k (Nagel, 1995; Stahl and Wilson, 1994) or cognitive hierarchy (CH) models (Camerer et al., 2004); see also a survey by Crawford et al. (2013). In a level- k model players anchor their beliefs in a non-strategic initial assessment of others’ likely responses to the game. Non-strategic players are said to be level-0 players. Level-1 players believe that all other players are at level 0, and they give their best response on the basis of this belief. Thus, in the context of voting, a level-1 player has an incentive to vote non-truthfully if and only if he is a GS-manipulator. Level-2 players assume that all other players belong to level 1, and, more generally, players at level k give their best response assuming that all other players are at level $k - 1$. The cognitive hierarchy model is similar, but with an essential difference: in this model players of level k respond to a mixture of types from level 0 to level $k - 1$. It is frequently assumed that other players’ levels are drawn from a Poisson distribution. Some further approaches based on similar ideas are surveyed by Wright and Leyton-Brown (2010). Although none of the game-theoretic models of boundedly rational voters seem to be perfect, we believe that specifics of voting, and, in particular, the heterogeneity of types of voters in real electorates, make the cognitive hierarchy framework more appropriate than others for our purposes and we adapt it appropriately.

This, however, is not a trivial task. First, an important feature that distinguishes voting from many other applications of both level- k and CH models is the role of level-0 players. Level-0 (non-strategic) players are typically assumed to choose their strategy at random while, in contrast, in voting situations it is natural to associate level-0 players with ideological or non-strategic voters, who have a significant presence in real elections. For instance, in the famous Florida vote (2000), where Bush won over Gore by just 537 votes, 97,488 Nader supporters voted for Nader—even though in such a close election strategic voters should have voted either for Gore or for Bush (and an overwhelming majority of Nader supporters preferred Gore to Bush) because of a real possibility of becoming pivotal. However, in the level- k analysis voters of level 2 assume that all other voters have level 1, i.e., level- k models cannot be used to accommodate ideological voters. We therefore focus on the cognitive hierarchy approach. However, as a vast majority of voters in elections can be assumed to be non-strategic, or level-0 players, we do not assume that types are drawn from a Poisson distribution. Moreover, we limit ourselves to considering the first three levels of the hierarchy (i.e., level-0, level-1, and level-2 players), as it seems plausible that very few voters are capable of higher-level reasoning (for evidence supporting this assumption we refer to the survey by Crawford et al. (2013)).

Another difficulty in adapting the cognitive hierarchy model to our setting is that in the standard model of social choice, voters’ preferences over alternatives are ordinal rather than cardinal (see, e.g., the already-cited classical work of Gibbard (1973) and Satterthwaite (1975)),

³We note in passing the setting of repeated elections or iterative voting, in which the information available to voters at each iteration step can be used as a basis for learning—see, e.g., the survey by Meir (2017).

which implies that, in general, voters at level 2 or higher do not have a straightforward definition of best response. We therefore resort to the concept of *weak dominance*, assuming that voters choose strategies that are not weakly dominated according to their beliefs. We present our formal definitions and further discussion in Section 3.

1.2 Main Results

Our first results provide a classification of level-1 strategies under k -approval; this is a necessary step to understand level-2 strategies in this setting (Section 4). We then switch our attention to higher-level strategies, and, in particular, to the complexity of computing them. For k -approval with $k = 1$ (i.e., the classic plurality rule) we describe an efficient algorithm that decides whether a given strategy weakly dominates another strategy; as a corollary of this result, we conclude that under the plurality rule level-2 strategies can be efficiently computed and efficiently recognised (Section 5). We obtain a similar result for 2-approval under an additional *minimality* assumption (Section 6). Briefly, this assumption means that the level-2 player expects all level-1 players to manipulate by making as few changes to their votes as possible. For larger values of k , finding level-2 strategies becomes computationally challenging: we show that this problem is NP-hard for k -approval with $k \geq 4$ (Section 7). As the problem of finding a level-1 strategy under k -approval is computationally easy for any value of $k \geq 1$ (this follows immediately by combining our characterisation of level-1 strategies with the classic results of Bartholdi et al. (1989)), this demonstrates that higher levels of voters' sophistication come with a price tag in terms of algorithmic complexity.

1.3 Related Work

There is a substantial body of research in social choice theory and in political science that models non-truthful voting as a strategic interaction, with a strong focus on the plurality rule; this line of work dates back to Farquharson (1969) and includes important contributions by Cain (1978), Feddersen et al. (1990) and Cox (1997), to name a few.

More recently, voting games and their equilibria have also received a considerable amount of attention from researchers in artificial intelligence, with a variety of approaches used to eliminate counterintuitive Nash equilibria. Meir (2018) provides an excellent survey of this line of work; we discuss a few specific papers below. Some authors assume that voters have a slight preference for abstaining or for voting truthfully when they are not pivotal (Battaglini, 2005; Dutta and Sen, 2012; Desmedt and Elkind, 2010; Thompson et al., 2013; Obraztsova et al., 2013; Elkind et al., 2015b; Obraztsova et al., 2015a). Other works consider refinements of Nash equilibrium, such as subgame-perfect Nash equilibrium (Desmedt and Elkind, 2010; Xia and Conitzer, 2010), strong equilibrium (Messner and Polborn, 2007; Sertel and Sanver, 2004) or trembling-hand equilibrium (Obraztsova et al., 2016), or model the reasoning of voters who have incomplete or imperfect information about each others' preferences (Myerson and Weber, 1993; Myatt, 2007; Meir et al., 2014). Dominance-based solution concepts have been investigated as well (Moulin, 1979; Dhillon and Lockwood, 2004; Buenrostro et al., 2013; Dellis, 2010; Meir et al., 2014), also from a computational perspective (Conitzer et al., 2011). All the aforementioned papers do not impose restrictions on the voters' reasoning ability, de facto assuming that they are fully rational.

Similarly to our work, Grandi et al. (2019) also consider boundedly rational voters. However, their work focuses on strategic interactions among Gibbard–Satterthwaite manipulators, and

studies conditions that ensure existence of pure strategy Nash equilibria in the resulting games. In contrast, in this paper we go further, and formally define the degree of voters’ rationality by using the cognitive hierarchy approach.

Level- k models and the cognitive hierarchy framework have been long used to model a variety of strategic interactions; see the survey of Crawford et al. (2013). Nevertheless, to the best of our knowledge, ours is the first attempt to apply these ideas in the context of voting.

A topic closely related to voting games is voting dynamics, where players change their votes one by one in response to the current outcome (Reijngoud and Endriss, 2012; Reyhaneh and Wilson, 2012; Obraztsova et al., 2015b; Endriss et al., 2016; Lev and Rosenschein, 2016; Koolyk et al., 2017; Meir et al., 2017); see also a recent survey by Meir (2017). However, this line of work assumes that voters are myopic, i.e., choose actions that are best responses to the current selections of the other voters, without trying to model the other voters’ reasoning.

Our work can also be seen as an extension of the model of safe strategic voting proposed by Slinko and White (2014). 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 a weakly dominant strategic vote in this setting (such votes are called *safe strategic votes*). In contrast, in our model the decision-maker takes into account that manipulators may have diverse preferences and have strategy sets that contain more than one strategic vote. It is therefore not surprising that computing safe strategic votes is easier than finding level-2 strategies: Hazon and Elkind (2010) show that safe strategic votes with respect to k -approval can be computed efficiently for every $k \geq 1$, whereas we obtain hardness results for $k \geq 4$.

One of our contributions is a classification of manipulative votes under k -approval with lexicographic tie-breaking. Peters et al. (2012) propose a similar classification for several approval-based voting rules. However, they view k -approval as a non-resolute voting rule, and therefore their results do not apply in our setting.

1.4 Paper Outline

The paper is organised as follows. We introduce the basic terminology and definitions in Section 2. Section 3 presents the adaptation of the cognitive hierarchy framework to the setting of voting games. We then focus on the study of k -approval. Section 4 describes the structure of level-1 strategies under k -approval. In Section 5 we provide an efficient algorithm for identifying level-2 strategies with respect to the plurality rule. Section 6 contains our results for 2-approval, and in Section 7 we present our hardness results for k -approval with $k \geq 4$. Section 8 summarises our results and suggests directions for future work.

2 Preliminary Definitions

In this section we introduce the relevant notation and terminology concerning preference aggregation and normal-form games. Throughout the paper, we write $[n]$ to denote the set $\{1, 2, \dots, n\}$.

2.1 Votes, Strategic Votes and Voting Rules

We consider n -voter elections over a candidate set $C = \{c_1, \dots, c_m\}$; in what follows we use the terms *candidates* and *alternatives* interchangeably. Let $\mathcal{L}(C)$ denote the set of all linear orders over C . The preferences of a particular society are defined by a *sincere preference profile*

$V = (v_1, \dots, v_n)$, where each v_i , $i \in [n]$, is a linear order over C ; we refer to v_i as the *sincere 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 , and 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 such that $a \succ_i b \succ_i \dots \succ_i z$. We denote the top candidate in v_i by $\text{top}(v_i)$. Also, we denote the set of top k candidates in v_i by $\text{top}_k(v_i)$; note that $\text{top}_1(v_i) = \{\text{top}(v_i)\}$ and $a \succ_i b$ for all $a \in \text{top}_k(v_i)$ and $b \in C \setminus \text{top}_k(v_i)$.

Voters sometimes prefer not to report their sincere preferences, but to vote strategically. Given a sincere preference profile $V = (v_1, \dots, v_n)$ and a linear order $v'_i \in \mathcal{L}(C)$, 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 . We will often use this notation when voter i submits a strategic vote v'_i instead of his sincere vote v_i .

A (resolute) *voting rule* \mathcal{R} is a mapping that, given a profile V (sincere or strategic), outputs a candidate $\mathcal{R}(V) \in C$, which we call the *winner of the election defined by V* , or simply the *winner at V* . In this paper we focus on the family of voting rules known as *k-approval*. Under *k-approval*, $k \in [m-1]$, each candidate receives one point from each voter who ranks her in top k positions in his vote; the *k-approval score* of a candidate c , denoted by $\text{sc}_k(c, V)$, is the total number of points that she receives. The winner is chosen among the candidates with the highest score according to a fixed tie-breaking linear order $>$ on the set of candidates C : specifically, the winner is the highest-ranked candidate with respect to this order among the candidates with the highest score. The 1-approval voting rule is widely used and known as the *plurality rule*. We will denote the *k-approval rule* (with tie-breaking based on a fixed linear order $>$) by \mathcal{R}_k . We say that a candidate x *beats* a candidate y at V with respect to \mathcal{R}_k and the tie-breaking order $>$ if $\text{sc}_k(x, V) > \text{sc}_k(y, V)$ or $\text{sc}_k(x, V) = \text{sc}_k(y, V)$ and $x > y$.

Definition 1. Consider a sincere profile $V = (v_1, \dots, v_n)$, a voter i , and a voting rule \mathcal{R} . We say that a linear order v'_i is a *manipulative vote of voter i at V with respect to \mathcal{R}* if $\mathcal{R}(V_{-i}, v'_i) \succ_i \mathcal{R}(V)$. We say that i *manipulates* in favour of candidate c by submitting a vote v'_i if c is the winner at $\mathcal{R}(V_{-i}, v'_i)$. A voter i is a *Gibbard–Satterthwaite manipulator*, or a *GS-manipulator*, at V with respect to \mathcal{R} if the set of his manipulative votes at V with respect to \mathcal{R} is not empty. We denote the set of all *GS-manipulators at V with respect to \mathcal{R}* by $N(V, \mathcal{R})$.

Note that a voter may be able to manipulate in favour of several different candidates. Let $F_i = \{c \in C \mid \mathcal{R}(V_{-i}, v'_i) = c \text{ for some } v'_i \in \mathcal{L}(C)\}$; we say that the candidates in F_i are *feasible for i at V with respect to \mathcal{R}* . Note that $F_i \neq \emptyset$ for all $i \in [n]$, as this set contains the \mathcal{R} -winner at V under truthful voting.

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 voter $i \in [n]$ and every profile V_{-i} of other voters' preferences. It is easy to see that v and v' are equivalent with respect to *k-approval* if and only if $\text{top}_k(v) = \text{top}_k(v')$.

2.2 Normal-form Games

A *normal-form game* $(N, (A_i)_{i \in N}, (\succeq_i)_{i \in N})$ is defined by a set of *players* N , and, for each $i \in N$, a set of *strategies* A_i and a preference relation \succeq_i defined on the space of *strategy profiles*,⁴ i.e., tuples of the form $s = (s_1, \dots, s_n)$, where $s_i \in A_i$ for all $i \in N$. For each pair of strategy profiles

⁴While one usually defines normal-form games in terms of utility functions, defining them in terms of preference relations is more appropriate for our setting, as preference profiles only provide ordinal information about the voters' preferences.

s, t and a player $i \in N$, we write $s \succ_i t$ if $s \succeq_i t$ and $t \not\succeq_i s$. A normal-form game is viewed as a game of complete and perfect information, which means that all players are fully aware of the structure of the game they are playing.

Given a strategy profile $s = (s_1, \dots, s_n)$ and a strategy $s'_i \in A_i$, we denote by (s_{-i}, s'_i) the strategy profile $(s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$ that is obtained from s by replacing s_i with s'_i . We say that a strategy $s_i \in A_i$ *weakly dominates* another strategy $s'_i \in A_i$ if for every strategy profile s_{-i} of other players we have $(s_{-i}, s_i) \succeq_i (s_{-i}, s'_i)$ and there exists a profile s_{-i} of other players' strategies such that $(s_{-i}, s_i) \succ_i (s_{-i}, s'_i)$.

3 The Model

Our goal is to analyse voting as a strategic game and consider it from the perspective of the cognitive hierarchy model. As we reason about voters' strategic behaviour, we consider games where players are voters, their strategies are ballots they can submit, and their preferences over strategy profiles are determined by election outcomes under a given voting rule. Thus, a voting game is a triple $G = (V, \mathcal{R}, (A_i)_{i \in [n]})$, where $V = (v_1, \dots, v_n)$ is the list of voters' preferences, \mathcal{R} a voting rule, and for each $i \in [n]$ the set A_i consists of all actions available to voter i .

3.1 A Cognitive Hierarchy Framework for Voting Games

When the cognitive hierarchy (CH) framework is used in the context of normal-form games, the players at level 0 are typically assumed to choose their action at random. Indeed, in general normal-form games a player who is unable to deliberate about other players' actions usually has no reason to prefer one strategy over another. In contrast, in the context of voting, there is an obvious focal strategy, namely, truthful voting.

At the next level of the hierarchy are level-1 voters. These voters believe that all other voters are sincere (i.e., are at level 0), and choose their vote so as to get the best outcome they consider possible under this assumption. That is, a level-1 voter votes so as to make his most preferred feasible candidate the election winner. If voter i is a level-1 voter and the set of feasible outcomes F_i is a singleton, which means that voter i cannot change the outcome at all no matter which vote he submits, we assume that i votes truthfully. This assumption reflects the special role of the sincere vote in the context of voting.

Definition 2. *We say that a vote v'_i of a voter i is a level-1 strategy at profile V with respect to \mathcal{R} if $\mathcal{R}(V_{-i}, v'_i) \succ_i c$ for all $c \in F_i \setminus \{\mathcal{R}(V_{-i}, v'_i)\}$.*

We emphasise that in our adaptation of the CH framework a level-1 voter that is not a Gibbard–Satterthwaite manipulator sees no reason to vote non-truthfully, since, given his belief that all other voters will vote truthfully, he does not expect to be able to change the election outcome according to his tastes.

We are now ready to discuss level-2 voters. These voters believe that all other voters are at levels 0 or 1 of the cognitive hierarchy. Specifically, a level-2 voter i identifies, for every other voter j , a set of votes A_j that he believes j may use. For each $j \in [n] \setminus \{i\}$, this set contains j 's truthful vote v_j , i.e., i entertains the possibility that j is an ideological voter. Further, if j is not a GS-manipulator, then $A_j = \{v_j\}$. On the other hand, if j is a GS-manipulator, the set A_j consists of v_j and a (possibly empty) subset of j 's level-1 strategies; this subset represents

i 's beliefs about j 's ability and desire to manipulate.⁵ In particular, if i believes that j is an ideological voter, this subset is empty.

Importantly, we do not assume that $A_j \setminus \{v_j\}$ is either empty or consists of all level-1 strategies of voter j . Indeed, the set of all level-1 strategies for a particular voter can be very large, and a level-2 voter may be unable or unwilling to identify all such votes. For example, a level-2 voter may know or believe that the other voters use a specific algorithm (e.g., that of Bartholdi et al. (1989)) to find their level-1 strategies. In this case, the set of strategies for each voter j would consist of the truthful vote v_j and the output of the respective algorithm. Further, a level-2 voter may expect that level-1 voters would not choose manipulations that are (weakly) dominated by other manipulations. Finally, voters may prefer not to change their vote beyond what is necessary to make their target candidate the election winner, either because they want their vote to be as close to the true preference order as possible (see the work of Obraztsova and Elkind (2012)), or for a fear of unintended consequences of such changes in the complex environment of the game; this, again, may rule out some of the level-1 strategies. Thus, a preference profile together with a voting rule define not just a single game, but a family of games, which differ in sets of actions available to GS-manipulators.

3.2 Gibbard–Satterthwaite Games

We will now complete the model, by defining games that enable us to reason about the decisions faced by a level-2 player. For convenience, we assume that voter 1 is a level-2 player and describe a normal-form game that captures his perspective of the strategic interaction, i.e., his beliefs about the game he is playing.

Fix a voting rule \mathcal{R} , let V be a sincere profile over a set of candidates C , let $N = N(V, \mathcal{R})$ be the set of GS-manipulators at V with respect to \mathcal{R} , and set $N_1 = N \cup \{1\}$. Since any level-1 player who is not a GS-manipulator must vote sincerely, and voter 1 believes that all other voters are level-0 or level-1 players, in his view of the game only the GS-manipulators can deviate from their true preferences; that is, he only needs to reason about the actions of the voters in $N_1 \setminus \{1\}$.

We consider a family of normal-form games defined as follows. In each game the set of players is N_1 , i.e., voter 1 is a player irrespective of whether he is actually a GS-manipulator. For each player $i \in N_1 \setminus \{1\}$, i 's strategy set A_i consists of his truthful vote and a (possibly empty) subset of his level-1 strategies; for voter 1 we have $A_1 = \mathcal{L}(C)$, i.e., voter 1 can submit an arbitrary ballot. It remains to describe the voters' preferences over strategy profiles. For a strategy profile $V^* = (v_i^*)_{i \in N_1}$, where $v_i^* \in A_i$ for all $i \in N_1$, let $V[V^*] = (v'_1, \dots, v'_n)$ be the preference profile such that $v'_i = v_i$ for $i \notin N_1$ and $v'_i = v_i^*$ for $i \in N_1$. Then, given two strategy profiles V^* and V^{**} and a voter $i \in N_1$, we write $V^* \succeq_i V^{**}$ if and only if i prefers $\mathcal{R}(V[V^*])$ to $\mathcal{R}(V[V^{**}])$ or $\mathcal{R}(V[V^*]) = \mathcal{R}(V[V^{**}])$. We refer to any such game as a *GS-game*.

We denote the set of all GS-games for V and \mathcal{R} by $\mathcal{GS}(V, \mathcal{R})$. Note that an individual game in $\mathcal{GS}(V, \mathcal{R})$ is fully determined by the GS-manipulators' sets of strategies, i.e., $(A_i)_{i \in N(V, \mathcal{R})}$ (voter 1's set of strategies is always the same, namely, $\mathcal{L}(C)$). 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 a strategy profile in a GS-game as a *GS-profile*, and we will sometimes identify the GS-profile $V^* = (v_i^*)_{i \in N_1}$ with the preference profile $V[V^*]$. We denote the set of all GS-profiles in a game G by $\mathcal{GSP}(G)$.

⁵With this approach, we depart from the recent literature of unawareness in game theory (Rêgo and Halpern, 2014, 2012), which studies interactive situations where a player might not have complete knowledge of the possible actions available to other players.

We will now argue that games in $\mathcal{GS}(V, \mathcal{R})$ reflect the perspective of voter 1 when he is at the second level of the cognitive hierarchy. Fix a game $G \in \mathcal{GS}(V, \mathcal{R})$. Note first that, since voter 1 believes that all other voters belong to levels 0 and 1 of the cognitive hierarchy, he expects all voters who are not GS-manipulators to vote truthfully, i.e., he does not need to reason about their strategies at all. This justifies having $N_1 = N \cup \{1\}$ as the set of players. On the other hand, consider a voter $i \in N_1 \setminus \{1\}$. voter 1 considers it possible that i is a level-0 voter, who votes truthfully. voter 1 also entertains the possibility that i is a level-1 voter, in which case i 's vote has to be a level-1 strategy; as argued above, voter 1 may also rule out some of i 's level-1 strategies. Consequently, the set A_i , which, by definition, contains v_i , consists of all strategies that voter 1 considers possible for i . Thus, voter 1's view of other voters' actions is captured by G .

3.3 Level-2 Strategies

We are now ready to discuss level-2 strategies. In the game-theoretic literature, it is typical to assume that a level-2 player is endowed with probabilistic beliefs about other players' types as well as a utility function describing his payoffs under all possible strategy profiles. Under these conditions, it makes sense to define player 1's level-2 strategies as those that maximise his expected payoff with respect to his beliefs. However, in the absence of numerical information, as in the case of voting games, we cannot reason about expected payoffs. Nevertheless, we can compare different strategies pointwise, and remove strategies that are weakly dominated by other strategies. On the other hand, if a strategy v is not weakly dominated, a level-2 player may hold beliefs that make him favour v , so no such strategy can be removed from consideration without making additional assumptions about the behaviour of the players in $N(V, \mathcal{R})$. This reasoning motivates the following definition of a level-2 strategy.

Definition 3. *Given a GS-game $G = (V, \mathcal{R}, (A_i)_{i \in N(V, \mathcal{R})})$, we say that a strategy $v \in A_1$ of player 1 is a level-2 strategy if no other strategy of player 1 weakly dominates v .*

We note that being weakly undominated is not a very demanding property: a strategy can be weakly undominated even if it fares badly in many scenarios, as illustrated by the following example.

Example 2. Consider the 4-voter profile over $\{a, b, c, d\}$ given in Table 2. Suppose that the voting rule is the plurality rule and the tie-breaking rule is $a > b > c > d$. As always, we assume that voter 1 is the level-2 player. Voters 2, 3, and 4 are GS-manipulators; their most preferred manipulations are, respectively, in favour of d , b , and c . Consider the GS-game where $A_2 = \{bdac, dbac\}$, $A_3 = \{cbad, bcad\}$, $A_4 = \{dcab, cdab\}$. In this game every vote that does not rank d first is a level-2 strategy for the first voter. Indeed, a vote that ranks a first is optimal when all other voters submit their sincere votes; a vote that ranks b first is optimal when voters 2 and 3 stay sincere, but voter 4 votes for c ; and a vote that ranks c first is optimal when voter 2 votes for d , but voters 3 and 4 stay sincere. Note, in particular, that, by changing his vote from $abcd$ (his sincere vote) to $cabd$, voter 1 changes the outcome from a (his top choice) to c (his third choice) when other voters vote truthfully; however, this behaviour is rational if voter 1 expects voters 3 and 4 (but not voter 2) to vote sincerely.

Example 2 illustrates that level-2 strategies are not 'safe': there can be circumstances where a level-2 strategy results in a worse outcome than sincere voting. Now, a cautious level-2 player may prefer to stick to his sincere vote unless he can find a manipulative vote which leads to an

voter 1	voter 2	voter 3	voter 4
a	b	c	d
b	d	b	c
c	a	a	a
d	c	d	b

Table 2: A profile where voter 1 has three distinct level-2 strategies under the plurality rule.

outcome that is at least as desirable as the outcome under truthful voting, for any combination of actions of other players that he considers possible. The following definition, which is motivated by the concept of *safe strategic voting* (Slinko and White, 2014), describes the set of strategies that even a very cautious level-2 player would prefer to sincere voting.

Definition 4. *Given a GS-game $G = (V, \mathcal{R}, (A_i)_{i \in N(V, \mathcal{R})})$, we say that a strategy $v \in \mathcal{L}(C)$ of voter 1 is an improving strategy if v weakly dominates voter 1's sincere strategy v_1 .*

We note that a level-2 strategy may fail to be an improving strategy, and, conversely, an improving strategy is not necessarily a level-2 strategy. For instance, in Example 1 approving a and b is a level-2 strategy, but not an improving strategy, and none of the level-2 strategies in Example 2 are improving. However, it is easy to see that if a player has an improving strategy, he also has an improving strategy that is a level-2 strategy. Moreover, an improving strategy exists if and only if sincere voting is not a level-2 strategy.

A natural goal for a level-2 player would be to identify strategies that weakly dominate all other (non-equivalent) strategies. However, while strategies with this property are highly desirable, from the perspective of a strategic voter it is more important to find out whether his truthful strategy is weakly dominated. Indeed, the main issue faced by a strategic voter is whether to manipulate at all, and if a certain vote can always ensure an outcome that is at least as good, and sometimes better, as that guaranteed by his truthful vote, this is a very strong incentive to use it, even if another non-truthful vote may be better in some situations. This issue is illustrated by Example 3 below, which describes a profile where a player has two incomparable improving strategies.

Example 3. Let the profile of sincere preferences be as in Table 3, and assume that the voting rule is the plurality rule and the tie-breaking order is given by $w > d > c > b > a$. The winner at the sincere profile is w . All level-1 strategies of voter 2 are equivalent to $cbdwa$, whereas all level-1 strategies of voter 3 are equivalent to $dcbwa$; voters 4 and 5 are not GS-manipulators. Consider the GS-game where for $i \in \{2, 3\}$ the set of strategies of player i consists of his truthful vote and all of his level-1 strategies. voter 1, who is a level-2 player, can manipulate either in favour of b or in favour of d , by ranking the respective candidate first. Indeed, for voter 1 both $badwc$ and $dabwc$ weakly dominate truth-telling. However, neither of these strategies weakly dominates the other: $badwc$ is preferable if no other player uses a level-1 strategy, whereas $dabwc$ is preferable if voter 2 uses his level-1 strategy, but voter 3 votes sincerely.

We note that a level-2 player may find it useful to act as a counter-manipulator (Pattanaik, 1976; Grandi et al., 2019), i.e., to submit a vote that is not a successful manipulation with respect to the truthful profile, but neutralises potential strategic actions of another voter. The following example illustrates one such situation.

voter 1	voter 2	voter 3	voter 4	voter 5
a	b	c	d	w
b	c	d	w	a
d	d	b	c	b
w	w	w	a	c
c	a	a	b	d

Table 3: voter 1 has two incomparable improving strategies.

Example 4. Let the profile of sincere preferences be as in Table 4, and assume that the voting rule is the plurality rule and the tie-breaking order is given by $a > b > c$. Under truthful voting a wins, so voter 6 is the only GS-manipulator: if he changes his vote to bca then b wins, and he prefers b to a . Therefore, for voter 1 voting acb is preferable to voting truthfully: this insincere vote has no impact if voter 6 votes truthfully, but prevents b from becoming a winner when voter 6 submits a manipulative vote. Thus, in this example acb is an improving strategy, and truthful voting is not a level-2 strategy as it is weakly dominated by voting acb . In contrast, acb is a level-2 strategy, as no other strategy weakly dominates it.

voter 1	voter 2	voter 3	voter 4	voter 5	voter 6
c	a	a	b	b	c
a	b	b	a	a	b
b	c	c	c	c	a

Table 4: Countermanipulation under the plurality rule.

4 Characterising Level-1 Strategies under k -Approval

The goal of this section is to understand and classify level-1 strategies under the k -approval voting rule; this will help us reason about level-2 strategies in subsequent sections. In what follows, we fix a linear order $>$ used for tie-breaking. We start with a simple, but useful lemma.

Lemma 1. *Fix $k \geq 1$. Consider a sincere profile V over C , let w be the k -approval winner at V , and let x be an alternative in $C \setminus \{w\}$. Then any manipulative vote by voter i in favour of x at V falls under one of the following two categories:*

Type 1 *voter i increases the score of x by 1 without decreasing the score of w . In this case $w, x \notin \text{top}_k(v_i)$, $x \succ_i w$, and the manipulative vote v'_i satisfies $x \in \text{top}_k(v'_i)$, $w \notin \text{top}_k(v'_i)$. In such cases voter i will be referred to as a promoter of x .*

Type 2 *voter i decreases the score of w (and possibly that of some other alternatives) by 1 without increasing the score of x . In this case $w, x \in \text{top}_k(v_i)$, $x \succ_i w$, and the manipulative vote v'_i satisfies $x \in \text{top}_k(v'_i)$, $w \notin \text{top}_k(v'_i)$. In such cases voter i will be referred to as a demoter of w . Manipulations of type 2 only exist for $k \geq 2$.*

Proof. Suppose that voter i manipulates in favour of x . If i can increase the score of x , then $x \notin \text{top}_k(v_i)$. However, i must rank x higher than w (otherwise, this would not be a manipulation).

Thus, $w \notin \text{top}_k(v_i)$ and therefore voter i cannot decrease w 's score. Moreover, if $w \in \text{top}_k(v'_i)$, then w would beat x under k -approval in (V_{-i}, v'_i) ; thus, $w \notin \text{top}_k(v'_i)$.

On the other hand, suppose that i cannot increase the score of x . This means that $x \in \text{top}_k(v_i)$ and hence i can only reduce the scores of some of x 's competitors including the current winner w . For this to be possible, it has to be the case that $w \in \text{top}_k(v_i)$ and $x \succ_i w$. Also, we have $w \notin \text{top}_k(v'_i)$, as otherwise w would beat x under k -approval in (V_{-i}, v'_i) . Finally, as $w \neq x$, we can only have $x, w \in \text{top}_k(v_i)$ if $k \geq 2$. \square

The classification in Lemma 1 justifies our terminology: a promoter promotes a new winner and a demoter demotes the old one. With the plurality rule, i.e., when $k = 1$, we only have promoters.

Let $X = \{x_1, \dots, x_\ell\}$ and $Y = \{y_1, \dots, y_\ell\}$ be two disjoint sets of candidates. Given a linear order v over C , we denote by $v[X; Y]$ the vote obtained by swapping x_j with y_j for $j \in [\ell]$. If the sets X and Y are singletons, i.e., $X = \{x\}$, $Y = \{y\}$, we omit the curly braces, and simply write $v[x; y]$. Clearly, under k -approval any manipulative vote of voter i is equivalent to a vote of the form $v_i[X; Y]$, where $X \subseteq \text{top}_k(v_i)$, $Y \subseteq C \setminus \text{top}_k(v_i)$. We can now state a corollary of Lemma 1, which characterises the possible effects of a manipulative vote under k -approval.

Corollary 2. *Let w be the k -approval winner at a sincere profile V , let $v_i^* = v_i[X; Y]$ be a manipulative vote, where $X \subseteq \text{top}_k(v_i)$ and $Y \subseteq C \setminus \text{top}_k(v_i)$. Let $V' = (V_{-i}, v_i^*)$, and let $w' \neq w$ be the k -approval winner at V' . Then either $w \in X$ or $w' \in Y$ but not both.*

Consider a manipulative vote $v_i[X; Y]$ of voter i at a sincere profile V under k -approval; we say that $v_i[X; Y]$ is *minimal* if for every other manipulative vote v'_i of voter i there is a vote $v_i[X'; Y']$ that is equivalent to v'_i and satisfies $|X'| \geq |X|$. That is, a manipulative vote is minimal if it performs as few swaps as possible. Arguably, minimal manipulative votes are the main tool that a rational voter would use, as they achieve the desired result in the most straightforward way possible.

We now introduce some useful notation. Fix a profile V . Let w be the k -approval winner at V , and let $t = \text{sc}_k(w, V)$. Set

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

and set $S(V, k) = S_1(V, k) \cup S_2(V, k)$.

The following proposition explains our attention to this set: only candidates from $S(V, k)$ can become k -approval winners as a result of a manipulation.

Proposition 1. *Suppose that some voter can manipulate in favour of a candidate $p \in C$ at a sincere profile V with respect to k -approval. Then $p \in S(V, k)$.*

Proof. Let w be the k -approval winner at V ; clearly, $w \neq p$. Suppose that voter i can manipulate in favour of p at V by submitting a vote v'_i ; let $V' = (V_{-i}, v'_i)$. Set $t = \text{sc}_k(w, V)$; then $\text{sc}_k(p, V) \leq t$. Note that if $\text{sc}_k(p, V) = t$, it has to be the case that $w > p$, since otherwise p would beat w at V . Thus, in this case $p \in S_1(V, k)$. Now, suppose that $\text{sc}_k(p, V) = t - 1$. By Corollary 2 we have either $\text{sc}_k(w, V') = \text{sc}_k(p, V') = t$ (if p was promoted) or $\text{sc}_k(w, V') = \text{sc}_k(p, V') = t - 1$ (if w was demoted). In both cases we have to have $p > w$, as otherwise w would beat p at V' . Therefore, in this case $p \in S_2(V, k)$. Finally, note that it cannot be the case that $\text{sc}_k(p, V) \leq t - 2$, since in this case by Corollary 2 we would have either $\text{sc}_k(w, V') \geq t - 1$, $\text{sc}_k(p, V') \leq t - 2$ or $\text{sc}_k(w, V') = t$, $\text{sc}_k(p, V') \leq t - 1$, i.e., w would beat p at V' . \square

Suppose that $S(V, k) \neq \emptyset$. If $S_1(V, k) \neq \emptyset$, then by $p^*(V, k)$ we denote the top-ranked candidate in $S_1(V, k)$ with respect to $>$; otherwise, we denote by $p^*(V, k)$ the top-ranked candidate in $S_2(V, k)$ with respect to $>$. Thus, $p^*(V, k)$ beats all candidates other than w at V , and would become a winner if it were to gain one point or if w were to lose one point. We omit V and k from the notation when they are clear from the context.

We are now ready to embark on the computational complexity analysis of level-2 strategies under k -approval, for various values of k .

5 The Plurality Rule

Recall that the plurality rule is \mathcal{R}_1 , i.e., k -approval with $k = 1$. For this rule we only have manipulators of type 1 (recall the definition in Section 4), and all manipulative votes of voter i in favour of candidate c are equivalent: in all such votes c is ranked first. The main result of this section shows that the problem of deciding whether a given strategy of voter 1 weakly dominates another strategy can be decided in polynomial-time. Note that, since under the plurality rule there are only m votes that are pairwise non-equivalent, this means that we can check if a given strategy is a level-2 strategy or an improving strategy, or find a level-2 strategy or an improving strategy (if it exists) in polynomial time; we formalise this intuition in Corollary 5 at the end of this section.

Fix a sincere preference profile V over a candidate set C , let $N = N(V, \mathcal{R}_1)$, and consider a GS-game $(V, \mathcal{R}_1, (A_i)_{i \in N})$. Let w be the plurality winner at V . As argued above, for each $i \in N \setminus \{1\}$ the set A_i consists of v_i and possibly a number of pairwise equivalent manipulative votes; without loss of generality, we can remove all but one manipulative vote, so that $|A_i| \leq 2$ for all $i \in N \setminus \{1\}$. We will now explain how, given two votes v'_1 and v''_1 , voter 1 can efficiently decide if one of these votes weakly dominates the other.

We will first describe a subroutine that will be used by our polynomial-time algorithm.

Lemma 3. *There is a polynomial-time procedure*

$$\text{Alg} = \text{Alg}(G, r, r', x, y, C^{[1]}, C^{[0]}, C^{[-1]}, C^{[-2]})$$

that, given a GS-game $G = (V, \mathcal{R}_1, (A_i)_{i \in N(V, \mathcal{R}_1)})$ with $|V| = n$, two integers $r, r' \in \{0, \dots, n\}$, two distinct candidates $x, y \in C$, and a partition of candidates in $C \setminus \{x, y\}$ into $C^{[1]}$, $C^{[0]}$, $C^{[-1]}$ and $C^{[-2]}$, decides whether there is a strategy profile V^* in G such that

- $\text{sc}_1(x, V[V^*]_{-1}) = r$,
- $\text{sc}_1(y, V[V^*]_{-1}) = r'$, and
- for each $c \in C \setminus \{x, y\}$ and each $\ell \in \{1, 0, -1, -2\}$ if $c \in C^{[\ell]}$ then $\text{sc}_1(c, V[V^*]_{-1}) \leq r + \ell$.

Proof. We proceed by reducing our problem to an instance of network flow with capacities and lower bounds, as follows. We construct a source, a sink, a node for each voter $i \in [n] \setminus \{1\}$ and a node for each candidate in C . There is an arc from the source to each voter node; the capacity and the lower bound of this arc are set to 1, i.e., it is required to carry one unit of flow. Also, there is an arc with capacity 1 and lower bound 0 from voter i to candidate c if $i \in N(V, \mathcal{R}_1) \setminus \{1\}$ and $c = \text{top}(v)$ for some $v \in A_i$ or if $i \in [n] \setminus (N(V, \mathcal{R}_1) \cup \{1\})$ and $c = \text{top}(v_i)$. Finally, there is an arc from each candidate c to the sink. The capacity of this arc is set to $r + \ell$ if $c \in C^{[\ell]}$ for some $\ell \in \{1, 0, -1, -2\}$; the lower bounds for these arcs are 0. For x , both the capacity and the

lower bound of the arc to the sink are set to r , and for y they are both set to r' . We note that some of the capacities may be negative, in which case there is no valid flow. It is immediate that an integer flow that satisfies all constraints corresponds to a strategy profile in G where all candidates have the required scores; it remains to observe that the existence of a valid integer flow can be decided in polynomial time. \square

We are now ready to describe our algorithm.

Theorem 4. *Given a GS-game $G = (V, \mathcal{R}_1, (A_i)_{i \in N(V, \mathcal{R}_1)})$ and two strategies $v'_1, v''_1 \in \mathcal{L}(C)$ of voter 1 we can decide in polynomial time whether v'_1 weakly dominates v''_1 .*

Proof. We will design a polynomial-time procedure that, given two strategies u, v of voter 1, decides if there exists a profile V_{-1}^* of other players' strategies such that $\mathcal{R}_1(V[V_{-1}^*, u]) \succ_1 \mathcal{R}_1(V[V_{-1}^*, v])$; by definition, v'_1 weakly dominates v''_1 if this procedure returns 'yes' for $u = v'_1, v = v''_1$ and 'no' for $u = v''_1, v = v'_1$.

Let $a = \text{top}(u), b = \text{top}(v)$. We can assume without loss of generality that $a \neq b$, since otherwise u and v are equivalent with respect to plurality. Consider an arbitrary profile V_{-1}^* of other players' strategies, and let $V^u = V[V_{-1}^*, u], V^v = V[V_{-1}^*, v], w^u = \mathcal{R}_1(V^u), w^v = \mathcal{R}_1(V^v)$. We note that $w^u \neq a$ implies $w^v \neq a$: if w^u beats a at V^u , this is also the case at V^v . Similarly, if $w^v \neq b$ then also $w^u \neq b$. Now, suppose that $w^u \neq a$ and $w^v \neq b$. We claim that in this case $w^u = w^v$. Indeed, suppose for the sake of contradiction that $w^u \neq w^v$. As $w^u \neq a, w^v \neq b$, the argument above shows that $\{w^u, w^v\} \cap \{a, b\} = \emptyset$. Thus, both w^u and w^v have the same score under the plurality rule at V^u and V^v ; as w^u beats w^v at V^u , this must also be the case at V^v , a contradiction.

Note that $\mathcal{R}_1(V[V_{-1}^*, u]) \succ_1 \mathcal{R}_1(V[V_{-1}^*, v])$ if and only if $w^u \succ_1 w^v$. By the argument in the previous paragraph, this can happen in one of the following three cases: (i) $w^u = a, w^v = b$ and $a \succ_1 b$; (ii) $w^u = a, w^v = w$ for some $w \neq b, a \succ_1 w$; (iii) $w^u = w, w^v = b$ for some $w \neq a, w \succ_1 b$. (We note that we can merge case (i) into case (ii) or case (iii); we choose not to do so for the sake of clarity of presentation.) We will now explain how to check if there exists a profile V_{-1}^* that corresponds to any of these three situations.

Case (i): $w^u = a, w^v = b$.

Suppose first that $a > b$. Then a desired profile V_{-1}^* exists if and only if there is some value $t \in [n]$ such that $\text{sc}_1(a, V^u) = t$ and

- (a) $\text{sc}_1(b, V^u) = t, \text{sc}_1(c, V^u) \leq t$ for all $c \in C \setminus \{a, b\}$ with $a > c, \text{sc}(c, V^u) \leq t - 1$ for all $c \in C \setminus \{a, b\}$ with $c > a$, or
- (b) $\text{sc}_1(b, V^u) = t - 1, \text{sc}_1(c, V^u) \leq t$ for all $c \in C \setminus \{a, b\}$ with $b > c$, and $\text{sc}(c, V^u) \leq t - 1$ for all $c \in C \setminus \{a, b\}$ with $c > b$.

Note that $\text{sc}_1(a, V_{-1}^u) = \text{sc}_1(a, V^u) - 1$ and $\text{sc}_1(c, V_{-1}^u) = \text{sc}_1(c, V^u)$ for $c \in C \setminus \{a\}$. Thus, to check if condition (a) is satisfied for some $t \in [n]$, we set $C^{[1]} = \{c \in C \setminus \{a, b\} \mid a > c\}, C^{[0]} = \{c \in C \setminus \{a, b\} \mid c > a\}, C^{[-1]} = C^{[-2]} = \emptyset$ and call

$$\text{Alg}(G, t - 1, t, a, b, C^{[1]}, C^{[0]}, C^{[-1]}, C^{[-2]}).$$

Similarly, to determine whether condition (b) is satisfied for some $t \in [n]$, we set $C^{[1]} = \{c \in C \setminus \{a, b\} \mid b > c\}, C^{[0]} = \{c \in C \setminus \{a, b\} \mid c > b\}, C^{[-1]} = C^{[-2]} = \emptyset$ and call

$$\text{Alg}(G, t - 1, t - 1, a, b, C^{[1]}, C^{[0]}, C^{[-1]}, C^{[-2]}).$$

The answer is ‘yes’ if one of these calls returns ‘yes’ for some $t \in [n]$.

For the case $b > a$ the analysis is similar. In this case, we need to decide whether there exists a value of $t \in [n]$ such that $\text{sc}_1(a, V^u) = t$ and

- (a) $\text{sc}_1(b, V^u) = t - 1$, $\text{sc}_1(c, V^u) \leq t$ for all $c \in C \setminus \{a, b\}$ with $a > c$, and $\text{sc}(c, V^u) \leq t - 1$ for all $c \in C \setminus \{a, b\}$ with $c > a$, or
- (b) $\text{sc}_1(b, V^u) = t - 2$, $\text{sc}_1(c, V^u) \leq t - 1$ for all $c \in C \setminus \{a, b\}$ with $b > c$, and $\text{sc}(c, V^u) \leq t - 2$ for all $c \in C \setminus \{a, b\}$ with $c > b$.

Again, this can be decided by calling the procedure *Alg* with appropriate parameters; we omit the details.

Case (ii): $w^u = a$, $w^v = w$ for some w with $a \succ_1 w$. In this case, we go over all candidates $w \in C \setminus \{a, b\}$ with $a \succ_1 w$ and all values of $t \in [n]$ and call *Alg* with appropriate parameters.

Specifically, if $a > w$, we start by setting $r = t - 1$, $r' = t$, and

$$C^{[1]} = \{c \in C \setminus \{a, w, b\} \mid w > c\}, C^{[0]} = \{c \in C \setminus \{a, w, b\} \mid c > w\}, C^{[-1]} = C^{[-2]} = \emptyset.$$

We then place b in $C^{[0]}$ if $w > b$ and in $C^{[-1]}$ otherwise; our treatment of b reflects the fact that she gets an extra point at V^v .

If $w > a$ we start by setting $r = t - 1$, $r' = t - 1$, and

$$C^{[1]} = \emptyset, C^{[0]} = \{c \in C \setminus \{a, w, b\} \mid w > c\}, C^{[-1]} = \{c \in C \setminus \{a, w, b\} \mid c > w\}, C^{[-2]} = \emptyset.$$

We then place b in $C^{[-1]}$ if $w > b$ and in $C^{[-2]}$ otherwise.

Finally, we call

$$\text{Alg}(G, r, r', a, w, C^{[1]}, C^{[0]}, C^{[-1]}, C^{[-2]}).$$

The answer is ‘yes’ if one of these calls returns ‘yes’ for some $t \in [n]$ and some w with $a \succ_1 w$.

Case (iii): $w^u = w$, $w^v = b$ for some w with $w \succ_1 b$. The analysis is similar to the previous case; we omit the details. \square

Theorem 4 immediately implies that natural questions concerning level-2 strategies and improving strategies are computationally easy.

Corollary 5. *Given a GS-game $G = (V, \mathcal{R}_1, (A_i)_{i \in N(V, \mathcal{R}_1)})$ and a strategy $v'_1 \in \mathcal{L}(C)$ of voter 1 we can decide in polynomial time whether v'_1 is a level-2 strategy or an improving strategy. Moreover, we can decide in polynomial time whether voter 1 has a level-2 strategy or an improving strategy in G .*

Proof. Let $a = \text{top}(v'_1)$. To decide whether v'_1 is an improving strategy, we use the algorithm described in the proof of Theorem 4 to check whether v'_1 weakly dominates v_1 . Similarly, to decide whether v'_1 is a level-2 strategy, for each $c \in C \setminus \{a\}$ we construct a vote v^c with $\text{top}(v^c) = c$ and check whether v^c weakly dominates v'_1 using the algorithm from the proof of Theorem 4. As every strategy of voter 1 is equivalent either to v'_1 or to one of the votes we constructed, v'_1 is a level-2 strategy if and only if it is not weakly dominated by any of the votes v^c , $c \in C \setminus \{a\}$.

Similarly, to decide whether voter 1 has a level-2 strategy (respectively, an improving strategy), we consider all of his m pairwise non-equivalent strategies, and check if any of them is a level-2 strategy (respectively, an improving strategy), as described in the previous paragraph. \square

Observe that, as the number of non-equivalent strategies available to a voter under the plurality rule is m , the problem of deciding whether voter 1 has an improving or a manipulative strategy has the same computational complexity as the problem of finding it, and the proposed algorithms can be used for the latter task as well.

6 2-Approval

In this section, we study the computational complexity of identifying level-2 strategies and improving strategies in GS-games under 2-approval. We show that if the level-2 player believes that level-1 players can only contemplate minimal manipulations, he can efficiently compute his level-2 strategies as well as his improving strategies. As argued in Section 4, minimality is a reasonable assumption, as level-1 players have no reason to use complex strategies when simple strategies can do the job.

Specifically, we prove that, under the minimality assumption, given two strategies v' and v'' , the level-2 player can decide in polynomial time whether one of these strategies weakly dominates the other; just as in the case of the plurality rule, this implies that he can check in polynomial time whether a given strategy is a level-2 (respectively, improving) strategy or identify all of his level-2 (respectively, improving) strategies.

The following observations play a crucial role in our analysis.

Proposition 2. *Consider a GS-game $G = (V, \mathcal{R}_2, (A_i)_{i \in N(V, \mathcal{R}_2)})$. Let w be the 2-approval winner at V . Then for each player $i \in N(V, \mathcal{R}_2) \setminus \{1\}$ such that $w \in \text{top}_2(v_i)$ it holds that $\text{top}(v_i) \neq w$ and the candidate $\text{top}(v_i)$ is ranked in the top two positions in every vote $v \in A_i$.*

Proposition 2 concerns voters who are demoters, and follows immediately from Lemma 1; note also that it does not depend on the minimality assumption.

Proposition 3. *Consider a GS-game $G = (V, \mathcal{R}_2, (A_i)_{i \in N(V, \mathcal{R}_2)})$. Let w be the 2-approval winner at V . Consider a player $i \in N(V, \mathcal{R}_2) \setminus \{1\}$ such that $w \notin \text{top}_2(v_i)$ and the set A_i consists of i 's truthful vote and a subset of i 's minimal manipulations. Let $\text{top}_2(v) = \{a, a'\}$. Then there is a candidate $c \in C \setminus \{a, a'\}$ such that for each $v \in A_i$ we have $\text{top}_2(v) \in \{\{a, a'\}, \{a, c\}, \{a', c\}\}$.*

Proof. Player i cannot lower the score of w by changing his vote, but he can raise the scores of some candidates in $C \setminus \text{top}_2(v_i)$ by moving these candidates into the top two positions. In general, i can do that for two candidates simultaneously; however, the minimality assumption implies that i only moves one candidate into the top two positions. Thus, i is a promoter (see Section 2). For a vote v'_1 to be a level-1 strategy the promoted candidate has to be i 's most preferred candidate in $S(V, 2) \setminus \text{top}_2(v_i)$ (let us denote this candidate by p). Thus, in this case voter i has three options: (1) to vote truthfully, (2) to swap p with the candidate that he ranks first or (3) to swap p with the the candidate he ranks second. This completes the proof. \square

Propositions 2 and 3 enable us to establish an analogue of Lemma 3 for 2-approval under the minimality assumption.

Lemma 6. *There is a polynomial-time procedure*

$$\text{Alg}' = \text{Alg}'(G, r, r', x, y, C^{[1]}, C^{[0]}, C^{[-1]}, C^{[-2]})$$

that, given a GS-game $G = (V, \mathcal{R}_2, (A_i)_{i \in N(V, \mathcal{R}_2)})$ with $|V| = n$, where for each $i \in N \setminus \{1\}$ the set A_i consists of i 's truthful vote and a subset of i 's minimal manipulations, two integers

$r, r' \in \{0, \dots, n\}$, two distinct candidates $x, y \in C$, and a partition of candidates in $C \setminus \{x, y\}$ into $C^{[1]}$, $C^{[0]}$, $C^{[-1]}$ and $C^{[-2]}$, decides whether there is a strategy profile V^* in G such that

- $\text{sc}_2(x, V[V^*]_{-1}) = r$,
- $\text{sc}_2(y, V[V^*]_{-1}) = r'$, and
- for each $\ell \in \{1, 0, -1, -2\}$ and each $c \in C^{[\ell]}$ it holds that $\text{sc}_2(c, V[V^*]_{-1}) \leq r + \ell$.

Proof. We transform an instance of our problem into an instance of the network flow problem using essentially the same construction as in the proof of Lemma 3. Let w be the 2-approval winner at V . If $S(V, 2) \neq \emptyset$, set $p^* = p^*(V, 2)$. We construct a flow network as follows. The set of nodes consists of a source, a sink, one node for each voter in $[n] \setminus \{1\}$, and one node for each candidate $c \in C$. For each $i \in [n] \setminus \{1\}$, the capacity and the lower bound of the arc from the source to node i are equal to 2, and the capacities and lower bounds of the arcs from candidates to the sink are defined as in the proof of Lemma 3. It remains to describe the arcs connecting voters and candidates.

If $i \notin N$, we add an arc from i to c for each $c \in \text{top}_2(v_i)$; the capacity and the lower bound of these arcs are 1, encoding the fact that i has to vote for his top two candidates.

Now, consider a voter $i \in N \setminus \{1\}$ who is a demoter; if such a voter exists, we have $S(V, 2) \neq \emptyset$ and hence p^* is well-defined. By Proposition 2 we have $\text{top}_2(v_i) = \{p^*, w\}$ and $A_i = \{v_i[w; c] \mid c \in C_i\}$ for some $C_i \subset C \setminus \{p^*, w\}$. Then we introduce an arc from i to p^* whose capacity and lower bound are both set to 1, and arcs with capacity 1 and lower bound 0 from i to each $c \in C_i \cup \{w\}$.

Finally, consider a voter $i \in N \setminus \{1\}$ who is a promoter; let $\text{top}_2(v_i) = \{a, a'\}$ and let p be i 's most preferred candidate in $S(V, 2) \setminus \{a, a'\}$. If A_i contains both a vote v' with $\text{top}_2(v') = \{a, p\}$ and a vote v'' with $\text{top}_2(v'') = \{a', p\}$, then by Proposition 3 it suffices to add arcs with capacity 1 and lower bound 0 that go from i to a, a' , and p . If we have $\text{top}_2(v) \in \{\{a, a'\}, \{a, p\}\}$ for each $v \in A_i$, we add an arc with capacity 1 and lower bound 1 from i to a and arcs with capacity 1 and lower bound 0 from i to a' and p . Similarly, if we have $\text{top}_2(v) \in \{\{a, a'\}, \{a', p\}\}$ for each $v \in A_i$, we add an arc with capacity 1 and lower bound 1 from i to a' and arcs with capacity 1 and lower bound 0 from i to a and p .

It is clear from the construction that a valid integer flow in this network corresponds to a strategy profile V^* with the desired properties. \square

We are now ready to prove the main result of this section.

Theorem 7. *Given a GS-game $G = (V, \mathcal{R}_2, (A_i)_{i \in N(V, \mathcal{R}_2)})$, where for each $i \in N \setminus \{1\}$ the set A_i consists of i 's truthful vote and a subset of i 's minimal manipulations, and two strategies $v'_1, v''_1 \in \mathcal{L}(C)$ of voter 1, we can decide in polynomial time whether v'_1 weakly dominates v''_1 .*

Proof. Just as in the proof of Theorem 4, it suffices to design a polynomial-time procedure that, given two strategies u, v of voter 1, decides if there exists a profile V_{-1}^* of other players' strategies such that $\mathcal{R}_2(V[V_{-1}^*, u]) \succ_1 \mathcal{R}_2(V[V_{-1}^*, v])$. Let $\text{top}_2(u) = \{a, a'\}$, $\text{top}_2(v) = \{b, b'\}$. We can assume that $\{a, a'\} \neq \{b, b'\}$, and we will focus on the case where $\{a, a'\} \cap \{b, b'\} = \emptyset$; the case where $\{a, a'\} \cap \{b, b'\}$ is a singleton is similar (and simpler).

We use the same notation as in the proof of Theorem 4: given a profile V_{-1}^* of other players' strategies, we let $V^u = V[V_{-1}^*, u]$, $V^v = V[V_{-1}^*, v]$, $w^u = \mathcal{R}_2(V^u)$, $w^v = \mathcal{R}_2(V^v)$. Our goal then is to decide if there exists a profile V_{-1}^* such that $w^u \succ_1 w^v$. To this end, we go over all

values of $t \in [n]$ and all candidates $w, w' \in C$ with $w \succ_1 w'$, and ask if there is a profile V_{-1}^* such that w wins at $V[V_{-1}^*, u]$ with t points, whereas w' wins at $V[V_{-1}^*, v]$. As in the proof of Theorem 4, for each triple (t, w, w') we have to consider a number of possibilities, depending on whether $w \in \{a, a'\}$, $w' \in \{b, b'\}$ as well as on the relative positions of w, w', a, a', b , and b' with respect to the tie-breaking order. The analysis is straightforward; to illustrate the main points, we consider two representative cases.

$w = a, w' = b, a > a' > b > b'$ In this case, a wins with t points at V^u if and only if $\text{sc}_2(a', V_{-1}^u) \leq t - 1$ and for each $c \in C \setminus \{a, a'\}$ we have $\text{sc}_2(c, V_{-1}^u) \leq t$ if $a > c$ and $\text{sc}_2(c, V_{-1}^u) \leq t - 1$ if $c > a$. Suppose that these conditions are satisfied. Then b can win at V^v with $t + 1$ or t points. The former case is possible if and only if $\text{sc}_2(b, V_{-1}^v) = t$. The latter case is possible if and only if $\text{sc}_2(b, V_{-1}^v) = t - 1$, $\text{sc}_2(b', V_{-1}^v) \leq t - 1$, and for each $c \in C \setminus \{a, b, b'\}$ we have $\text{sc}_2(c, V_{-1}^v) \leq t$ if $b > c$ and $\text{sc}_2(c, V_{-1}^v) \leq t - 1$ if $c > b$.

Thus, to decide whether this situation is possible, we have to call Alg' twice. For our first call, we set $C^{[1]} = \{c \in C \setminus \{a', b\} \mid a > c\}$, $C^{[0]} = \{c \in C \setminus \{a'\} \mid c > a\} \cup \{a'\}$, $C^{[-1]} = C^{[-2]} = \emptyset$ and call

$$\text{Alg}'(G, a, b, t - 1, t, C^{[1]}, C^{[0]}, C^{[-1]}, C^{[-2]}).$$

For our second call, we set $C^{[1]} = \{c \in C \setminus \{b'\} \mid b > c\}$, $C^{[0]} = \{c \in C \setminus \{a\} \mid c > b\} \cup \{b'\}$, $C^{[-1]} = C^{[-2]} = \emptyset$ and call

$$\text{Alg}'(G, a, b, t - 1, t - 1, C^{[1]}, C^{[0]}, C^{[-1]}, C^{[-2]}).$$

$w \notin \{a, a', b, b'\}, w' = b, b' > b > a' > w > a$ If w wins at V^u with t points, this means that $\text{sc}_2(w, V_{-1}^u) = t$, $\text{sc}_2(a', V_{-1}^u) \leq t - 2$, $\text{sc}_2(a, V_{-1}^u) \leq t - 1$, $\text{sc}_2(c, V_{-1}^u) \leq t$ for all $c \in C \setminus \{w, a, a'\}$ with $w > c$, and $\text{sc}_2(c, V_{-1}^u) \leq t - 1$ for all $c \in C \setminus \{w, a, a'\}$ with $c > w$. Suppose that these conditions are satisfied. As w still receives t points at V^v , this means that b wins at V^v if and only if $\text{sc}_2(b, V_{-1}^v) = t - 1$, $\text{sc}_2(b', V_{-1}^v) \leq t - 2$. Thus, we set $C^{[1]} = \emptyset$, $C^{[0]} = \{c \in C \setminus \{a\} \mid w > c\}$, $C^{[-1]} = \{c \in C \setminus \{a', b, b'\} \mid c > w\}$, $C^{[-2]} = \{a', b'\}$ and call

$$\text{Alg}'(G, w, b, t, t - 1, C^{[1]}, C^{[0]}, C^{[-1]}, C^{[-2]}). \quad \square$$

Just as for the plurality rule, we obtain the following corollary, which describes the complexity of finding and testing level-2 strategies and improving strategies under 2-approval.

Corollary 8. *Given a GS-game $G = (V, \mathcal{R}_2, (A_i)_{i \in N(V, \mathcal{R}_2)})$, where for each $i \in N$ the set A_i consists of i 's truthful vote and a subset of his minimal manipulations, and a strategy $v'_1 \in \mathcal{L}(C)$ of voter 1 we can decide in polynomial time whether v'_1 is a level-2 strategy and whether v'_1 is an improving strategy. Moreover, we can decide in polynomial time whether voter 1 has a level-2 strategy or an improving strategy in G .*

We remark that the minimality assumption plays an important role in our analysis. Indeed, in the absence of this assumption a promoter i may manipulate by swapping two different candidates into the top two positions. Let $v_i[\text{top}_2(v_i); \{p, c\}]$ be some such manipulation, where p is i 's most preferred candidate in $S(V, 2) \setminus \text{top}_2(v_i)$. If we try to model this possibility via a network flow construction, we would have to add edges from i to both p and c ; the lower bounds on these edges would have to be set to 0, to allow i to vote truthfully. However, there may then

be a flow that uses the edge (i, c) , but not (i, p) , which corresponds to a vote that promotes c , but not p ; such a vote is not a level-1 strategy.

Interestingly, a level-2 player may want to swap two candidates into the top two positions, even if he assumes that all level-1 players use minimal strategies. In fact, the following example shows that a strategy of this form may weakly dominate all other non-equivalent strategies.

Example 5. Let the profile of sincere preferences be as in Table 5, and assume that the voting rule is 2-approval and the tie-breaking order is given by $a > b > c > d > \dots$. Assume that voter 1 is the level-2 player. The winner under 2-approval is a with two points; candidates b , c , and d also have two points each. Voters 4 and 5 are GS-manipulators; voter 4 may manipulate by swapping c into the top two positions, and voter 5 may manipulate by swapping d into the top two positions.

Consider the GS-game where $N = \{1, 4, 5\}$, $A_4 = \{v_4, v_4[b; c]\}$, and $A_5 = \{v_5, v_5[b; d]\}$ (note that both A_4 and A_5 only contain a proper subset of the respective player's minimal manipulations; for instance, $v_4[u_1; c] \notin A_4$). We claim that $v_1[\{u_5, u_6\}; \{b, c\}]$ is a weakly dominant strategy for voter 1. Indeed, consider the four possible scenarios:

- Voters 4 and 5 are truthful. Then the best outcome that voter 1 can ensure is that b wins.
- Voter 4 is truthful, but voter 5 manipulates. Then the best outcome that voter 1 can ensure is that c wins.
- Voter 4 manipulates, but voter 5 is truthful. Then the best outcome that voter 1 can ensure is that c wins.
- Voters 4 and 5 both manipulate. Then the best outcome that voter 1 can ensure is that c wins.

Now, it is clear that only votes that rank b and c in the top two positions achieve all of these objectives simultaneously.

voter 1	voter 2	voter 3	voter 4	voter 5	voter 6	voter 7
u_5	a	a	b	b	c	c
u_6	d	d	u_1	u_2	u_3	u_4
b	c	d	a	a
c
d
...

Table 5: The strategy $bc\dots$ of voter 1 weakly dominates all non-equivalent strategies.

7 k -Approval for $k \geq 3$

Regrettably, our analysis of k -approval under the minimality assumption does not extend from $k = 2$ to $k = 3$. Specifically, the argument breaks down when we consider a potential demoter under 3-approval who can only help his top candidate by swapping his second and third candidate out of the top three positions. If he chooses to manipulate, he has to perform both of these

swaps at once; he can also remain truthful and not perform any swaps. It is not clear how to capture this all-or-nothing behaviour via network flows. We conjecture that finding a level-2 strategy under 3-approval is computationally hard, even under the minimality assumption. We will now prove a weaker result, showing that this problem is NP-hard for k -approval with $k \geq 4$ (and without the minimality assumption). Moreover, we will also show that it is coNP-hard to decide whether a given strategy is improving.

Theorem 9. *For every fixed $k \geq 4$, given a GS-game $G = (V, \mathcal{R}_k, (A_i)_{i \in N})$ and a strategy v of voter 1, it is NP-hard to decide whether v is a level-2 strategy, and it is coNP-hard to decide whether v is an improving strategy.*

Proof. Our hardness proof proceeds by a reduction 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.

We will first establish that our problems are hard for $k = 4$; towards the end of the proof, we will show how to extend our argument to $k > 4$.

Consider an instance $I^0 = (\Gamma^0, \Sigma^0)$ of X3C with $|\Gamma^0| = 3\nu'$. We will first modify this instance as follows. We add three new elements to Γ^0 and a set containing them to Σ^0 . We then add $\nu' + 2$ triples x_i, y_i, z_i , $i \in [\nu' + 2]$, of new elements to Γ^0 and for each such triple we add the set $S_i = \{x_i, y_i, z_i\}$ to Σ^0 . Finally, we add sets $S'_i = \{y_i, z_i, x_{i+1}\}$, $i \in [\nu' + 1]$, and $S'_{\nu'+2} = \{y_{\nu'+2}, z_{\nu'+2}, x_1\}$ to Σ^0 . We then renumber the elements of the ground set so that the elements added at the first step are numbered g_1, g_2, g_3 . We denote the resulting instance by (Γ, Σ) , and let $\nu = |\Gamma|/3$, $\mu = |\Sigma|$. Clearly, $I = (\Gamma, \Sigma)$ is a ‘yes’-instance of X3C if and only if $I^0 = (\Gamma^0, \Sigma^0)$ is. We let $\widehat{\Sigma} = \{S_i, S'_i \mid i \in [\nu' + 2]\}$; we have $\nu = 2\nu' + 3$, $|\widehat{\Sigma}| = 2(\nu' + 2)$.

We construct a GS-game as follows. We introduce a set of candidates $C' = \{c_1, \dots, c_{3\nu}\}$ that correspond to elements of Γ , three special candidates w, p, c , and, finally, a set of dummy candidates

$$D = \bigcup_{i=0}^{\mu} D_i \cup D_c \cup \bigcup_{j=1}^{\nu+1} E_j \cup \bigcup_{i=1}^{3\nu} \bigcup_{j=1}^{\nu+1} F_{i,j},$$

where $|D_i| = 4$ for $i = 0, \dots, \mu$, $|D_c| = 3$, and $|E_j| = 2$, $|F_{i,j}| = 3$ for $i \in [3\nu]$, $j \in [\nu + 1]$. Thus, the set of candidates is $C = \{w, p, c\} \cup C' \cup D$. We define the tie-breaking order $>$ on C by setting

$$w > c > p > c_1 > \dots > c_{3\nu} > D.$$

For each $j \in [\mu]$, we let $C_j = \{c_i \mid g_i \in \sigma_j\}$.

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 , while the elements within X and within Y are ranked according to the tie-breaking order $>$. The profile V consists of $2 + \mu + (3\nu + 1)(\nu + 1)$ votes defined as follows:

$$\begin{aligned} z_0 &= D_0 \succ p \succ c_1 \succ c \succ C' \setminus \{c_1\} \succ \dots, \\ z_i &= D_i \succ C_i \succ c \succ \dots, & i \in [\mu], \\ u &= c \succ D_c \succ w \succ \dots, \\ u_j &= w \succ p \succ E_j \succ \dots, & j \in [\nu + 1], \\ u_{i,j} &= c_i \succ F_{i,j} \succ w \succ \dots & i \in [3\nu], j \in [\nu + 1]. \end{aligned}$$

We have

$$\text{sc}_4(w, V) = \text{sc}_4(p, V) = \text{sc}_4(c_i, V) = \nu + 1 \quad \text{for all } i \in [3\nu],$$

$\text{sc}_4(c, V) = 1$, and $\text{sc}_4(d, V) \leq 1$ for each $d \in D$. Thus, w wins under 4-approval because of the tie-breaking rule.

We have $S(V, 4) = C' \cup \{p\}$. The set of GS-manipulators in this profile consists of the first $\mu + 1$ voters; we assume that the first voter (i.e., voter 0) is the level-2 voter. We now define a GS-game for this profile by constructing the players' sets of strategies as follows:

$$z'_i = z_i[D_i; C_i \cup \{c\}], \quad A_i = \{z_i, z'_i\} \quad \text{for all } i \in [\mu].$$

Observe that for each $i \in [\mu]$ the vote z'_i is a level-1 strategy for voter i , which makes i 's top candidate in C_i the winner with $\nu + 2$ points (note that voter i orders C_i in the same way as $>$ does, so tie-breaking favours i 's most preferred candidate in C_i). This completes the description of the game G .

Fix some $d, d' \in D_0$ and let

$$z'_0 = z_0[\{d, d'\}; \{p, c\}], \quad z''_0 = z_0[d; p].$$

Note that both z'_0 and z''_0 are level-1 strategies for voter 0, which make p the winner with $\nu + 2$ points. Clearly, we can construct the profile V and the players' sets of strategies in polynomial time given I .

We will now argue that z'_0 is an improving strategy if and only if $I = (\Gamma, \Sigma)$ is a 'no'-instance of X3C, and that z''_0 is a level-2 strategy if and only if $I = (\Gamma, \Sigma)$ is a 'yes'-instance of X3C.

As a preliminary observation, consider some strategy z of voter 0 such that $\text{top}_4(z)$ consists of p and three dummy candidates. By construction, for every profile of other players' strategies, z and z''_0 result in the same outcome. Moreover, if everyone except voter 0 votes truthfully, voter 0 strictly prefers z''_0 to every strategy \hat{z} with $\text{top}_4(\hat{z}) \subseteq D$. Thus, z''_0 can only be weakly dominated by a strategy that places at least one candidate from $C' \cup \{c, w\}$ in top 4 positions.

Suppose first that I is a 'yes'-instance of X3C. Fix a subcollection Σ' witnessing this, and consider a profile V' where the GS-manipulators that correspond to sets in Σ' vote strategically, whereas everyone else votes truthfully. We have $\text{sc}_4(p, V') = \text{sc}_4(c, V') = \text{sc}_4(w, V') = \nu + 1$, $\text{sc}_4(c_i, V') = \nu + 2$ for all $c_i \in C'$, so c_1 wins. However, if voter 0 changes his vote to z'_0 , the winner would be c , and voter 0 prefers c_1 to c , so voter 0 strictly prefers voting z_0 over voting z'_0 in this case, i.e., z'_0 is not an improving strategy.

Now, if voter 0 changes her vote to z''_0 instead, p becomes the election winner, which is the best feasible outcome from voter 0's perspective. The only way for voter 0 to achieve this outcome is to rank p and some dummy candidates in the top 4 positions; any vote \hat{z} with $\text{top}_4(\hat{z}) \cap (C' \cup \{c, w\}) \neq \emptyset$ is strictly worse for voter 0, and hence cannot weakly dominate z''_0 . As we have already observed that no strategy \hat{z} with $\text{top}_4(\hat{z}) \cap (C' \cup \{c, w\}) = \emptyset$ can weakly dominate z''_0 , it follows that if I is a 'yes'-instance of X3C then z''_0 is a level-2 strategy.

On the other hand, suppose that I is a 'no'-instance of X3C. Consider a strategy profile V^* in G , and let Σ'' be a subcollection of Σ that corresponds to players in $[\mu]$ who vote non-truthfully in V^* ; we know that Σ'' is not an exact cover of Γ . We will argue that voter 0 weakly prefers z'_0 to both z_0 and z''_0 for every choice of Σ'' , and there are choices of Σ'' for which this preference is strict.

If $\Sigma'' = \emptyset$, i.e., all voters in $[\mu]$ are truthful, then voter 0 benefits from changing his vote from z_0 to z'_0 , as this vote makes p the winner. Similarly, suppose that all sets in Σ'' are pairwise

disjoint (and hence $|\Sigma''| \leq \nu - 1$). Then candidate c gets at most ν points and the winner in $V[V_{-0}^*, z_0]$ is one of the candidates from C' (with $\nu + 2$ points). On the other hand, the winner in $V[V_{-0}^*, z'_0]$ is p (with $\nu + 2$ points), so voter 0 benefits from changing his vote to z'_0 . In both of these cases, z''_0 has the same effect as z'_0 .

Now, suppose that the sets in Σ'' are not pairwise disjoint. Let X be the set of elements that appear in the largest number of sets in Σ'' , and let g_ℓ be the element of X with the smallest index. Note that $g_\ell \neq g_1$, since we modified our instance of X3C so that g_1 only occurs in one set. The winner in $V[V_{-0}^*, z_0]$ is either c_ℓ or c , and the winner's score is at least $\nu + 3$. Suppose that voter 0 changes his vote from z_0 to z'_0 . If the winner in $V[V_{-0}^*, z_0]$ was c , this remains to be the case, and if the winner was c_ℓ then either c_ℓ remains the winner or c becomes the winner, and voter 0 prefers c to c_ℓ . Thus, in this case voting z'_0 is at least as good as voting z_0 , and voting z''_0 has the same effect as voting z_0 .

We conclude that whenever Σ'' is not an exact cover of Γ , voting z'_0 is at least as good as voting z_0 or z''_0 . It remains to establish that z'_0 is sometimes strictly better than either of these strategies. To this end, suppose that $\Sigma'' = \widehat{\Sigma}$. If voter 0 votes z'_0 , then the scores of the candidates covered by sets in $\widehat{\Sigma}$ are $\nu + 3$, the score of c is $1 + 2(\nu' + 2) + 1 = 2\nu' + 6 = \nu + 3$, and all other candidates have lower scores, so c wins. However, if voter 0 votes z_0 or z''_0 , c 's score is $\nu + 2$, and therefore the winner is a candidate in C' . Thus, in those circumstances, voter 0 strictly prefers z'_0 to both z_0 and z''_0 . Hence, if I is a 'no'-instance of X3C, z'_0 weakly dominates z_0 and z''_0 , and hence z''_0 is not a level-2 strategy. This completes the proof for $k = 4$.

For $k > 4$, we modify the construction by introducing $|V|$ additional groups of dummy candidates $H_1, \dots, H_{|V|}$ of size $k - 4$ each. We renumber the voters from 1 to $|V|$ and modify the preferences of the i -th voter, $i \in [\mu]$, by inserting the group H_i in positions $5, \dots, k$, and adding all other new dummy candidates at the bottom of his ranking. Then the k -approval scores of all candidates in C remain the same as in the original construction, and the k -approval score of each new dummy candidate is 1. The rest of the proof then goes through without change. \square

We note that the strategies of level-1 players in our hardness proof are not minimal; determining whether our hardness result remains true under the minimality assumption is an interesting research challenge.

Our complexity lower bounds are not tight: we do not know whether the computational problems we consider are in, respectively, NP and coNP. The following argument provides upper bounds on their complexity. Recall that the complexity class DP (difference polynomial-time) is composed by languages that are an intersection of an NP-language and a coNP-language (Papadimitriou and Yannakakis, 1984; Wechsung, 1985).

Proposition 4. *For every GS-game based on a polynomial-time voting rule (including k -approval) the problem of checking whether a given strategy is improving is in DP.*

Proof. For every n -player game $G = (N, (A_i)_{i \in N}, (\succeq_i)_{i \in N})$, where each relation \succeq_i is represented by a polynomial-time computable function of its arguments, and for every pair of strategies u, v of player 1, the problem of deciding whether u weakly dominates v belongs to the complexity class DP. Indeed, u weakly dominates v if and only if

- (a) for every profile P_{-1} of other players' strategies we have $(P_{-1}, u) \succeq_1 (P_{-1}, v)$ (which can be checked by a coNP algorithm), and

- (b) for some profile P_{-1} of other players' strategies we have $(P_{-1}, u) \succ_1 (P_{-1}, v)$ (which can be checked by an NP algorithm),

i.e., the language associated with our problem is an intersection of an NP-language and a coNP-language. \square

As a corollary, we obtain that for k -approval with a fixed value of k the problem of checking whether a given strategy is a level-2 strategy belongs to the Boolean hierarchy (see, e.g., Cai et al. (1988, 1989)). Indeed, there are only $\binom{m}{k} \leq m^k$ pairwise non-equivalent votes, and it suffices to check that none of these votes weakly dominates the given strategy.

8 Conclusions and Further Research

We have initiated the analysis of voting games from the perspective of the appropriately modified cognitive hierarchy framework. We have adopted a distribution-free approach that uses the concept of weak dominance in order to reason about players' actions. The resulting framework is mathematically rich, captures some interesting behaviours, and presents a number of algorithmic challenges, even for simple voting rules. To illustrate this, we focused on a well-known family of voting rules, namely, k -approval with $k \geq 1$, and investigated the complexity of finding level-2 strategies and improving strategies with respect to various rules in this family. This choice enabled us to start from one of the simplest and most used voting rules, i.e., the plurality rule, and then move upward in order to identify the boundary between tractable and intractable questions. For the plurality rule, i.e., for $k = 1$, we found that level-2 strategies and improving strategies are easy to find, and for $k \geq 4$ the associated problems are computationally hard, but for $k = 2, 3$ we do not have a full understanding of their computational complexity. We identify a natural assumption (namely, the minimality assumption), which is sufficient to obtain an efficient algorithm for $k = 2$; however, it is not clear if it remains useful for larger values of k . We list a few specific algorithmic questions that remain open:

- Is there a polynomial-time algorithm for computing level-2 strategies and improving strategies under 2-approval without the minimality assumption?
- Does Theorem 9 remain true under the minimality assumption?
- What can be said about 3-approval, with or without the minimality assumption?
- What can be said about other prominent voting rules, most importantly the Borda rule?

In our analysis, we have focused on level-1 and level-2 players. It would also be interesting to extend our formal definitions to level- ℓ players for $\ell \geq 3$ and to investigate the associated algorithmic issues. While it is intuitively clear that the view of the game for these players will be more complex, it appears that for the plurality rule our algorithm can be extended in a straightforward manner; however, it is not clear if this is also the case for 2-approval. Another interesting question, which can be analysed empirically, is whether truthful voting is likely to be a level-2 strategy, or, more broadly, how many votes in $\mathcal{L}(C)$ are level-2 strategies; again, this question can also be asked for level- ℓ strategies with $\ell \geq 2$.

A yet broader question, which can only be answered by combining empirical data and theoretical analysis, is whether the cognitive hierarchy approach provides a plausible description of strategic behaviour in voting. While our paper makes the first steps towards answering it, there is more to be done to obtain a full picture.

Acknowledgements This work was supported by an ERC Starting Grant ACCORD (GA 639945), and by Marsden Fund 3706352 of The Royal Society of New Zealand.

References

- J. J. Bartholdi and J. B. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.
- J. J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.
- M. Battaglini. Sequential voting with abstention. *Games and Economic Behavior*, 51:445–463, 2005.
- A. Brandenburger. Knowledge and equilibrium in games. *The Journal of Economic Perspectives*, 6(4):83–101, 1992.
- F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2015.
- L. Buenrostro, A. Dhillon, and P. Vida. Scoring rule voting games and dominance solvability. *Social Choice and Welfare*, 40(2):329–352, 2013.
- J. Cai, T. Gundermann, J. Hartmanis, L. A. Hemachandra, V. Sewelson, K. W. Wagner, and G. Wechsung. The Boolean hierarchy I: structural properties. *SIAM Journal on Computing*, 17(6):1232–1252, 1988.
- J. Cai, T. Gundermann, J. Hartmanis, L. A. Hemachandra, V. Sewelson, K. W. Wagner, and G. Wechsung. The Boolean hierarchy II: applications. *SIAM Journal on Computing*, 18(1): 95–111, 1989.
- B. Cain. Strategic voting in Britain. *American Journal of Political Science*, 22:639–655, 1978.
- C. F. Camerer, T.-H. Ho, and J.-K. Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898, 2004.
- V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*, 2011.
- G. Cox. *Making Votes Count*. Cambridge University Press, 1997.
- V. P. Crawford, M. A. Costa-Gomes, and N. Iriberry. Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1):5–62, 2013.
- A. Dellis. Weak undominance in scoring rule elections. *Mathematical Social Sciences*, 59(1): 110–119, 2010.
- Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, 2010.

- A. Dhillon and B. Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46:55–75, 2004.
- B. Dutta and A. Sen. Nash implementation with partially honest individuals. *Games and Economic Behavior*, 74(1):154–169, 2012.
- E. Elkind, U. Grandi, F. Rossi, and A. Slinko. Gibbard–Satterthwaite games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2015a.
- E. Elkind, S. Obraztsova, E. Markakis, and P. Skowron. Equilibria of plurality voting: Lazy and truth-biased voters. In *Proceedings of the 8th International Symposium on Algorithmic Game Theory (SAGT)*, 2015b.
- U. Endriss, S. Obraztsova, M. Polukarov, and J. S. Rosenschein. Strategic voting with incomplete information. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- 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.
- D. Fudenberg and D. K. Levine. *The theory of learning in games*, volume 2. MIT press, 1998.
- A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.
- U. Grandi, D. Hughes, F. Rossi, and A. Slinko. Gibbard-Satterthwaite games for k -approval voting rules. *Mathematical Social Sciences*, 99:24–35, 2019.
- J. C. Harsanyi and R. Selten. *A general theory of equilibrium selection in games*. MIT Press, 1988.
- N. Hazon and E. Elkind. Complexity of safe strategic voting. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*, 2010.
- A. Koolyk, T. Strangeway, O. Lev, and J. S. Rosenschein. Convergence and quality of iterative voting under non-scoring rules. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.
- O. Lev and J. S. Rosenschein. Convergence of iterative scoring rules. *Journal of Artificial Intelligence Research*, 57:573–591, 2016.
- R. Meir. Iterative voting. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access, 2017.
- R. Meir. *Strategic Voting*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2018.
- R. Meir, O. Lev, and J. S. Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of the 15th ACM Conference on Electronic Commerce (EC)*, 2014.

- R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Iterative voting and acyclic games. *Artificial Intelligence*, 252:100–122, 2017.
- M. Messner and M. K. Polborn. Strong and coalition-proof political equilibria under plurality and runoff rule. *International Journal of Game Theory*, 35(2):287–314, 2007.
- H. Moulin. Dominance solvable voting schemes. *Econometrica*, 47:1337–1351, 1979.
- D. P. Myatt. On the theory of strategic voting. *The Review of Economic Studies*, 74:255–281, 2007.
- R. Myerson and R. Weber. A theory of voting equilibria. *American Political Science Review*, 87(1):102–114, 1993.
- R. Nagel. Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5):1313–1326, 1995.
- S. Obraztsova and E. Elkind. Optimal manipulation of voting rules. In *Proceedings of 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 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 (SAGT)*, 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)*, 2015a.
- S. Obraztsova, E. Markakis, M. Polukarov, Z. Rabinovich, and N. R. Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015b.
- 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)*, 2016.
- C. H. Papadimitriou and M. Yannakakis. The complexity of facets (and some facets of complexity). *Journal of Computer and System Sciences*, 28(2):244–259, 1984.
- P. K. Pattanaik. Threats, counter-threats, and strategic voting. *Econometrica: Journal of the Econometric Society*, pages 91–103, 1976.
- H. Peters, S. Roy, and T. Storcken. On the manipulability of approval voting and related scoring rules. *Social Choice and Welfare*, 39(2-3):399–429, 2012.
- L. Rêgo and J. Halpern. Generalized solution concepts in games with possibly unaware players. *International Journal of Game Theory*, 41:131–155, 2012.
- L. Rêgo and J. Halpern. Extensive games with possibly unaware players. *Mathematical Social Sciences*, 70:42 – 58, 2014.

- A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.
- R. Reyhaneh and M. Wilson. Best reply dynamics for scoring rules. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI)*, 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.
- M. Sertel and R. Sanver. Strong equilibrium outcomes of voting games are the generalized condorcet winners. *Social Choice and Welfare*, 22(2):331–347, 2004.
- L. S. Shapley. Some topics in two-person games. *Advances in game theory, Annals of Mathematical Studies*, 52:1–29, 1964.
- A. Slinko and S. White. Is it ever safe to vote strategically? *Social Choice and Welfare*, 43:403–427, 2014.
- D. O. Stahl and P. W. Wilson. Experimental evidence on players’ models of other players. *Journal of economic behavior & organization*, 25(3):309–327, 1994.
- 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.
- G. Wechsung. On the Boolean closure of NP. In *Proceedings of the 4th International Conference on Fundamentals of Computation Theory (FCT)*, 1985.
- J. R. Wright and K. Leyton-Brown. Beyond equilibrium: Predicting human behavior in normal-form games. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 2010.
- L. Xia and V. Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 2010.