An NTU Cooperative Game Theoretic View of Manipulating Elections

Michael Zuckerman[1], Piotr Faliszewski[2], Vincent Conitzer[3], and Jeffrey S. Rosenschein[1]

[1] School of Computer Science and Engineering, The Hebrew Univ. of Jerusalem, Israel. {michez, jeff}@cs.huji.ac.il
[2] AGH University of Science and Technology, Kraków, Poland. faliszew@agh.edu.pl
[3] Department of Computer Science, Duke University Durham, NC 27708, USA. conitzer@cs.duke.edu

Abstract. Social choice theory and cooperative (coalitional) game theory have become important foundations for the design and analysis of multiagent systems. In this paper, we use cooperative game theory tools in order to explore the coalition formation process in the coalitional manipulation problem. Unlike earlier work on a cooperative-game-theoretic approach to the manipulation problem [2], we consider a model where utilities are not transferable. We investigate the issue of stability in coalitional manipulation voting games; we define two notions of the core in these domains, the α-core and the β-core. For each type of core, we investigate how hard it is to determine whether a given candidate is in the core. We prove that for both types of core, this determination is at least as hard as the coalitional manipulation problem. On the other hand, we show that for some voting rules, the α- and the β-core problems are no harder than the coalitional manipulation problem. We also show that some prominent voting rules, when applied to the truthful preferences of voters, may produce an outcome not in the core, even when the core is not empty.

1 Introduction

Voting constitutes a natural methodology for a group of agents to make a joint decision in spite of (possibly) conflicting preferences. Unfortunately, voting and elections are not a universal, perfect solution to preference aggregation problems. For example, the classic result of Gibbard and Satterthwaite [9, 13] says that in sufficiently general settings, any reasonable voting rule may lead to a situation where some voter(s) are better off by casting votes different from their true preferences (this is called *manipulation* or *strategic voting*). One of the most influential ideas regarding the computational study of elections, due to Bartholdi, Orlin, Tovey, and Trick [4, 3], was to study the computational complexity of computing manipulative votes. The rationale behind it was that if planning a manipulation were computationally hard, then in practice voters would not be able to vote strategically (see the surveys of Faliszewski, Hemaspaandra, and Hemaspaandra [7] and of Faliszewski and Procaccia [8] for a detailed overview of this approach and for two viewpoints regarding its applicability).

Formally, in the *coalitional* manipulation problem, introduced by Conitzer, Sandholm, and Lang [5], the voters are divided into two groups, the manipulators and the nonmanipulators. The votes of the nonmanipulators are assumed to be known, and the problem is to determine whether the manipulators can use their votes to achieve a given goal. The goal is either to ensure that some preferred candidate wins (the *constructive* variant) or to ensure that some despised candidate does not win (the *destructive* variant). Further, we obtain different flavors of the problem depending on whether the votes are weighted or not, which voting rule is used, etc. Some results on coalitional manipulation can be found in [17, 15, 14, 16]; see also surveys [8, 7].

Let us focus on unweighted constructive coalitional manipulation. If all the manipulators have identical preferences, this is exactly the problem that they want to be able to solve—they would first try to ensure their most preferred candidate's victory; then, if that were impossible, they would try the second best one, third best one, and so on, until they would either find a successful manipulation or determine that they cannot do better than electing the truthful winner.

Nonetheless, generally manipulators do *not* have identical preferences. In this case, the manipulators may still work together, but it is much less clear which candidate they should try to promote (even ignoring computational considerations). While they may all agree that they would prefer a different winner than the truthful one, deciding *which* candidate to support is a whole new game that they need to play among themselves. (To push our scenario to the limit, consider a situation where all the voters are manipulators.)

In this paper, we take the viewpoint of *cooperative game theory* to solve such games among the manipulators, and study computational aspects of the relevant solution concepts. As in most of the literature on voting, we assume that the agents do not have the ability to make or receive payments, so that we are in the *nontransferable utility (NTU)* case of cooperative game theory. (Recently, Bachrach, Elkind, and Faliszewski [2] studied a similar problem in the *transferable utility* setting, and obtained results linking coalitional manipulation [5] and bribery [6] problems with their cooperative game-theoretic model.) Moreover, in this setting, what one (sub)coalition of manipulators can achieve depends on the actions (votes) of the manipulators outside the coalition. We consider two different ways of addressing this—via the standard notions of the α- and the β-core [12].

2 Preliminaries

Let us now define the basic notions of (computational) social choice theory and coalitional game theory, as used in this paper.

An election E is a triple $(C, V, \mathcal{P})$, where $C = \{c_1, \dots, c_m\}$ is the set of candidates, $V = \{1, \dots, n\}$ is the set of voters, and $\mathcal{P} = (P_1, \dots, P_n)$ is a preference profile of voters in V. That is, each voter i, $1 \leq i \leq n$, is associated with preference order P_i from $\mathcal{P}$. A preference order is a linear order over the candidates in C. We will sometimes write $\succ_i$ instead of P_i. We write $L(C)$ to denote the set of all linear orders over C. Let U be some subset of V. By $\mathcal{P}_U$ we mean $(P_i)_{i \in U}$ and by $\mathcal{P}_{-U}$ we mean $(P_i)_{i \notin U}$. Using standard notational conventions, we have $\mathcal{P} = (\mathcal{P}_U, \mathcal{P}_{-U})$.

A *voting rule* $\mathcal{R}$ is a function that, given election E as input, returns one of the candidates, denoted $\mathcal{R}(E)$, as the election's winner. (We assume ties are resolved by some simple tie-breaking rule and that the manipulators, unwilling to rely on tie-breaking, always seek unique winners; we point the reader to the work of Obraztsova, Elkind, and Hazon [11], and Obraztsova and Elkind [10] for a detailed discussion of tie-breaking in voting manipulation.) We focus on the following (families of) voting rules:

Positional scoring rules. Let $\boldsymbol{s} = (s_1, \dots, s_m)$ be a vector of non-negative integers, such that $s_1 \geq s_2 \geq \dots \geq s_m$. For each voter, a candidate receives s_1 points if it is ranked first by the voter, s_2 points if it is ranked second, etc. The score of a candidate is the total number of points the candidate receives. The winner is the candidate with the maximum score. The scoring rules which we will consider here are k-approval, where $\boldsymbol{s} = (1, \dots, 1, 0, \dots, 0)$ $(s_1 = \dots = s_k = 1; s_{k+1} = \dots = s_m = 0)$, Plurality, where $\boldsymbol{s} = (1, 0, \dots, 0)$, and Borda, where $\boldsymbol{s} = (m-1, m-2, \dots, 0)$.

Plurality with Runoff. In this rule, a first round eliminates all candidates except the two with the highest plurality scores. The second round determines the winner between these two by their pairwise election.

Simplified Bucklin. A candidate c's Bucklin score is the smallest number k such that more than half of the votes rank c among the top k candidates. The winner is the candidate that has the smallest Bucklin score.[4]

We use the term *(coalitional) manipulation* to refer to a situation where a voter (a group of voters) casts votes not according to his (their) true preferences, but rather to obtain some goal. It is one of the best-studied forms of strategic behavior in elections (see the surveys [8, 7]). The definition below is taken from the paper of Bachrach, Elkind, and Faliszewski [2], which itself is inspired by the definition of Conitzer, Sandholm, and Lang [5].

Definition 1. *For any voting rule $\mathcal{R}$, an instance $I = (E, S, c)$ of the $\mathcal{R}$-*COALITIONAL MANIPULATION *problem is given by an election $E = (C, V, \mathcal{P})$, a set of* manipulators *S, $S \cap V = \emptyset$, and a distinguished candidate $c \in C$. It is a "yes"-instance if there is a vector $\mathcal{P}_S = (P_i)_{i \in S} \in (L(C))^{|S|}$ such that $\mathcal{R}(\mathcal{P}, \mathcal{P}_S) = c$, and a "no"-instance otherwise.*

Note that in the traditional definition of coalitional manipulation the manipulators, unlike honest voters, do not have preferences over the candidates; they simply want to get a particular candidate elected. Thus, this definition eliminates the problem of deciding which candidates the manipulators should support by making it external to the setting.

3 Manipulation Model

The goal of this paper is to study the unweighted constructive coalitional manipulation problem in the setting in which not all manipulators have identical preferences. Specifically, our focus is on the computational complexity of deciding which candidates the manipulators may support in a stable way, without breaking the coalition.

[4] The Nonsimplified Bucklin rule additionally breaks ties by the number of votes that rank a candidate among the top k candidates. In computational social choice it is common to focus on Simplified Bucklin instead of its full variant.

Formally, we consider the following setting. We are given an election $E = (C, V, \mathcal{P})$ where some of the voters are truthful and some are willing to manipulate. We denote the set of possible manipulators by M and we will refer to them as *colluders*. In our model the remaining voters, those in $H = V \setminus M$, are *honest voters* who vote truthfully. We assume that the manipulators have a way of communicating with one another (this is a standard assumption, though typically it is not mentioned explicitly). We ask which candidate the colluders should support. Naturally, the answer to this question depends strongly on the attitudes that the colluders have towards one another, and on the way they expect one another to behave. We consider two settings, depending on how the colluders react to the breaking out of the coalition by some subset M' of M: (a) the pessimistic model, where players in M' want to succeed irrespective of the reaction of players in $M \setminus M'$, and (b) the adaptive model, where players in M' can pick their votes depending on the votes of players in $M \setminus M'$.

Definition 2. *Let $\mathcal{R}$ be a voting rule and let $E = (C, V, \mathcal{P})$ be an election with colluder set M. Fix a candidate $c \in C$ and a coalition $S \subseteq M$.*

1. *We say that c is* feasible *for S if there is a profile $\mathcal{P}'_S$ such that $\mathcal{R}(\mathcal{P}'_S, \mathcal{P}_{-S}) = c$.*
2. *We say that c is* α-feasible *for a coalition S if there is a preference profile $\mathcal{P}'_S$ such that for all preference profiles $\mathcal{P}''_{M\setminus S}$ it holds that $\mathcal{R}(\mathcal{P}'_S, \mathcal{P}''_{M\setminus S}, \mathcal{P}_{-M}) = c$.*
3. *We say that $c \in C$ is* β-feasible *for a coalition $S \subseteq M$ if for all preference profiles $\mathcal{P}'_{M\setminus S}$ there exists a preference profile $\mathcal{P}''_S$ such that $\mathcal{R}(\mathcal{P}''_S, \mathcal{P}'_{M\setminus S}, \mathcal{P}_{-M}) = c$.*

We denote the set of all candidates that are feasible (respectively, α-feasible, β-feasible) for S by $F(S)$ (respectively, $F_\alpha(S)$, $F_\beta(S)$).

As a side remark, it is easy to see that $F(S)$ is never empty as it always contains $\mathcal{R}(E)$; but $F_\alpha(S)$ and $F_\beta(S)$ may be empty.

Using our feasibility notions, we can adapt the notions of α- and β-core to our setting [12].

Definition 3. *Let $E = (C, V, \mathcal{P})$ be an election, M be a subset of V, $\mathcal{R}$ be a voting rule, and c be a candidate. We say that c belongs to the α-core (respectively, β-core) if $c \in F(M)$ and there is no candidate $c' \in C$ and non-empty coalition $W' \subseteq M$ such that (a) $c' \in F_\alpha(W')$ (respectively, $c' \in F_\beta(W')$), and (b) each voter in W' prefers c' to c.*

One can see that these notions are related to the notion of strong Nash equilibrium (SNE) in the game played by the colluders. Recall that an SNE is a Nash equilibrium in which no coalition, taking the actions of its complements as given, can cooperatively deviate in a way that benefits all of its members [1]. In the context of SNE, we have the following scenario: the colluders have agreed upon some voting profile. The deviating coalition can privately communicate; when a coalition coordinates a deviation, the remaining players are unaware of it, so they stick to their agreed-upon strategies. In this model, it is relatively easy for a sub-coalition to break off since the deviating colluders know in advance how the rest of the colluders will vote. On the other hand, the α-core and the β-core model an election in which it is non-trivial for a subgroup of manipulators to break off from the coalition. (Consider, e.g., an election on choosing an acceptable debt level for a country and the manipulators being MPs from the ruling party. Even if they personally disagreed with some particular proposed debt level, they would need to obtain very strong support before breaking off from the coalition.) Here, if a subcoalition breaks off then it has to be ready to face every possible reaction of the colluders they abandon.

The difference between the α-core and the β-core is that in the case of the β-core the splitting subcoalition knows that it will know the remaining colluders' votes before having to cast their own. (In the parliamentary example from the previous paragraph, this means that MPs leaving the ruling party know they will be asked about their vote on the debt issue only after the ruling party will be.)

The above intuitions are supported by the observation that, by definition, if $\mathcal{P}'_M$ is SNE then $c = \mathcal{R}(\mathcal{P}_H, \mathcal{P}'_M) \in \beta$-core, and β-core $\subseteq \alpha$-core. (To see this, note that if for $c' \in C \setminus \{c\}$, and $W' \subseteq M$ such that everybody in W' prefers c' to c, $c' \in F_\beta(W')$ then W' can deviate from $\mathcal{P}'_M$ to make c' a winner, and so $\mathcal{P}'_M$ is not a SNE; and if for some candidate c' and set $W' \subseteq M$ it holds that $c' \in F_\alpha(W')$ then $c' \in F_\beta(W')$.) That is, settings where the α-core is an appropriate model are more stable than those where colluders do not counteract deviations by other colluders.

It is interesting that there are prominent voting rules and preference profiles for which the non-empty α-core/β-core does not contain the truthful winner.

Example 4. Consider a Maximin election with candidate set $C = \{a, b, c\}$, and the following preference profile (the numbers in parentheses are the serial numbers of the voters): (1) $b \succ c \succ a$, (2) $c \succ a \succ b$, (3) $a \succ c \succ b$, and (4) $b \succ a \succ c$.

Suppose that voter (1) is the honest voter, and the rest are the colluders. Suppose also that the order of the tie-breaking is c preferred to a preferred to b. When we apply the Maximin voting rule to the above election, c wins. Now we show that both the α- and the β-core of this election are equal to $\{a\}$. First we show that $c \notin \alpha$-core (and thus, $c \notin \beta$-core). For any candidates $x, y \in C$, denote by $W_{xy} \subseteq M$ the subset of colluders that prefer x to y. $W_{ac} = \{3, 4\}$. It is easy to verify that if the voters in W_{ac} both vote $a \succ b \succ c$, then no matter how the colluder (2) votes, a wins the election. This proves that $c \notin \alpha$-core. Similarly, $W_{ab} = \{2, 3\}$. It is easy to see that if the voters in W_{ab} both vote $a \succ b \succ c$, then no matter how voter (4) votes, a wins the election. So, $b \notin \alpha$-core. It remains to show that $a \in \beta$-core. $W_{ca} = \{2\}$. We already know that if the voters in $M \setminus W_{ca} = W_{ac}$ both vote $a \succ b \succ c$, then no matter how the colluder (2) votes, a wins the election. So, c is not β-feasible for W_{ca}. $W_{ba} = \{4\}$. As mentioned earlier, if the voters in $M \setminus W_{ba} = W_{ab}$ both vote $a \succ b \succ c$, then no matter how the voter (4) votes, a wins the election. Thus, b is not β-feasible for W_{ba}, and hence, $a \in \beta$-core.

Example 5. Consider a Borda election with candidate set $C = \{a, b, c, d, e, f\}$, and the following preference profiles: (1) $f \succ d \succ b \succ c \succ a \succ e$, (2) $a \succ f \succ b \succ c \succ d \succ e$, (3) $b \succ f \succ d \succ c \succ a \succ e$, and (4) $d \succ b \succ f \succ a \succ e \succ c$.

Here we have again one honest voter (1), and the rest are the colluders. In the above election, f wins. Now consider the set $W_{bf} = \{3, 4\}$. If (3) votes $b \succ e \succ c \succ a \succ d \succ f$, and (4) votes $b \succ a \succ e \succ c \succ d \succ f$, then no matter how voter (2) votes, b wins the election. Hence $b \in F_\alpha(W_{bf}) \Rightarrow f \notin \alpha$-core. We claim that $b \in \beta$-core (and hence $b \in \alpha$-core). Indeed, b is the Condorcet winner among the colluders. Hence, if the two colluders who prefer b to some candidate $x \in \{a, c, d, e, f\}$, vote $b \succ e \succ c \succ a \succ d \succ f$ and $b \succ a \succ e \succ c \succ d \succ f$, then no matter how the third colluder votes, b wins the election. Hence, no other candidate is β-feasible for the coalition that prefers him, which implies b is in the β-core.

To study our core notions computationally, we need to define an appropriate decision problem.

Definition 6. *For a voting rule $\mathcal{R}$, an instance $I = (E, M, c)$ of the $\mathcal{R}$-α-CORE (respectively, $\mathcal{R}$-β-Core) problem, is given by an election $E = (C, V, \mathcal{P})$, a set of colluders $M \subseteq V$ and a candidate $c \in C$. It is a "yes"-instance if c is in the α-core (respectively, β-core) of the corresponding game among the manipulators, and a "no"-instance otherwise.*

4 Main Results

In this section we present our main computational results. Namely, we provide a computational analysis of the notions of the α- and β-cores. We first provide a connection between our notions of the core and the standard problem of coalitional manipulation.

Theorem 7. *Given an election $E = (C, V, \mathcal{P})$, a set $M \subseteq V$, a voting rule $\mathcal{R}$ and a candidate $c \in C$, the problem of determining whether c is in the α- or β-core is at least as hard as the corresponding coalitional manipulation problem.*

Proof. We show a reduction from the coalitional manipulation problem. Given a voting rule $\mathcal{R}$ and an instance $I = (E, S, c)$, where $E = (C, V, \mathcal{P})$, of the $\mathcal{R}$-Coalitional Manipulation problem, we construct the core problem in the following way. We set $E' = (C, V', \mathcal{P}')$ where $V' = V \cup S$, $\mathcal{P}' = \mathcal{P} \cup \mathcal{P}'_S$, and $\mathcal{P}'_S$ is a profile of voters in S where c is ranked first in all the ballots, and the rest of the candidates are ranked in some arbitrary order. We set $M = S$. There exists a manipulation by voters in S making c win in the coalitional manipulation problem if and only if c is in the core (either α or β) of the core problem (which is defined by the instance (E', M, c)), since here c is in the core if and only if c is feasible for the coalition M (as there is no non-empty coalition W' where the voters in W' prefer c' to c). $\square$

However, could it be that the α- or the β-core problems are strictly harder than the coalitional manipulation problem? We will see that for many voting rules, the answer is "no". Due to Theorem 7, we focus on those voting rules for which the coalitional manipulation problem is polynomial-time solvable.

Theorem 8. *Let R be a positional scoring rule with scoring vector $s = (s_1, \ldots, s_m)$ such that the R-Coalitional Manipulation problem is polynomial-time solvable. Then there exists a polynomial-time algorithm solving the R-α-Core problem.*

Proof. Suppose we are given an instance $I = (E, M, c)$ of the R-α-Core problem, where $E = (C, V, \mathcal{P})$. Suppose w.l.o.g. that $s_m = 0$. The algorithm first checks whether $c \in F(M)$ by solving the R-Coalitional Manipulation problem. Then it goes over all the other candidates. For each candidate $x \neq c$ it checks what is the maximal set $W_{xc} \subseteq M$ of the colluders who prefer x to c. If $W_{xc} \neq \emptyset$, it builds a "profile"[5] $\mathcal{Q}_{M \setminus W_{xc}}$ for $M \setminus W_{xc}$ s.t. the score of x is $s(\mathcal{Q}_{M \setminus W_{xc}}, x) = 0$, and the score of any other candidate $y \neq x$ is $s(\mathcal{Q}_{M \setminus W_{xc}}, y) = s_1 \cdot |M \setminus W_{xc}|$. Finally, it solves the R-Coalitional Manipulation problem $(C, V \setminus W_{xc}, \mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}}), x, |W_{xc}|$ to see whether the votes in W_{xc} can be cast in order to make x win the election. Next we prove that we manage to make x win if and only if x is α-feasible for W_{xc}.

Claim. Let $\mathcal{Q}_{W_{xc}}$ be a profile. x is the winner under $\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}$ (where $\mathcal{Q}_{M \setminus W_{xc}}$ is as defined above) if and only if for any (real) profile $\mathcal{P}'_{M \setminus W_{xc}}$ x wins under $\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}$.

Proof. We first prove the "only if" part. Suppose that x wins under $\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}$. Let $\mathcal{P}'_{M \setminus W_{xc}}$ be any profile. For all y, $y \neq x$, it holds that $s(\mathcal{P}'_{M \setminus W_{xc}}, y) \leq s_1 \cdot |M \setminus W_{xc}|$. Also, $s(\mathcal{P}'_{M \setminus W_{xc}}, x) \geq 0$. Therefore, for each y, $y \neq x$, $s(\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, x) \geq s(\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, x) > s(\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, y) = s(\mathcal{P}_H \cup \mathcal{Q}_{W_{xc}}, y) + s_1 \cdot |M \setminus W_{xc}| \geq s(\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, y)$.[6] And so, x will win under $\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}$.

Now we prove the "if" part. Suppose that for any profile $\mathcal{P}'_{M \setminus W_{xc}}$, x is the winner under $\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}$. Let y, $y \neq x$ be a candidate. By our assumption, x is the winner under $\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}$, where $\mathcal{P}'_{M \setminus W_{xc}}$ is built as follows: y is ranked on top of each preference in $\mathcal{P}'_{M \setminus W_{xc}}$, x is ranked at the bottom of each preference in $\mathcal{P}'_{M \setminus W_{xc}}$, and the other candidates are ranked arbitrarily. From this construction we get: $s(\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, x) = s(\mathcal{P}_H, x) + s(\mathcal{Q}_{W_{xc}}, x) = s(\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, x) > s(\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, y) = s(\mathcal{P}_H, y) + s_1 \cdot |M \setminus W_{xc}| + s(\mathcal{Q}_{W_{xc}}, y) = s(\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, y)$. This is true for every $y \neq x$. Therefore, x is the winner under $\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}$. **(End of proof of claim.)** □

If we find a candidate $x \neq c$ and a non-empty coalition W_{xc} such that x is α-feasible for W_{xc}, then c is not in the α-core. Otherwise, c is in the α-core. □

Theorem 9. *Let $k \in \mathbb{N}$, $1 \leq k \leq |C| - 1 = m - 1$. There is a polynomial-time algorithm solving the k-approval-β-Core problem.*

Proof. Suppose we are given an instance $I = (E, M, c)$ of the k-approval-β-Core problem, where $E = (C, V, \mathcal{P})$. The algorithm first checks whether $c \in F(M)$ using, for example, the reduction to a particular scheduling problem that was observed by Xia et al. in [14]. Then it goes over all the other candidates. For each candidate $x \neq c$ it checks what the maximal set $W_{xc} \subseteq M$ of the colluders who prefer x to c is. If $W_{xc} \neq \emptyset$, it builds a profile $\mathcal{Q}_{M \setminus W_{xc}}$ for $M \setminus W_{xc}$ where x is ranked in the last place, and his k strongest opponents are ranked in the k first places, in all the votes. Finally, it runs the

---

[5] We used the double quotes for the word "profile" since it is not a real profile, but rather a score specification. Nevertheless, we use the above notation for simplicity, as if it were a real profile.

[6] We use here a strong inequality as we prove our results for the unique winner model. However, our results can be modified to work for the co-winner model as well.

above-mentioned algorithm on $(C, V \setminus W_{xc}, \mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}}), x, |W_{xc}|$ to see whether the votes in W_{xc} can be cast in order to make x win the election. If we fail to make x win, then by definition x is not β-feasible. In the next claim we prove that the opposite is also true, i.e., if the algorithm finds a manipulation, then x is β-feasible.

Claim. If there exists a profile $\mathcal{Q}_{W_{xc}}$ such that x wins the election under the profile $\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}$ (where $\mathcal{Q}_{M \setminus W_{xc}}$ is as defined above), then for each profile $\mathcal{P}'_{M \setminus W_{xc}}$ there exists a profile $\mathcal{P}'_{W_{xc}}$ such that x is the winner under $\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}$.

Proof. Suppose that x is the winner under $\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}$. Let $\mathcal{P}'_{M \setminus W_{xc}}$ be any profile of $M \setminus W_{xc}$. In the following algorithm we will convert $\mathcal{Q}_{W_{xc}}$ into some other profile $\mathcal{P}'_{W_{xc}}$ such that x will be the winner under $\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}$. We assume w.l.o.g. that $s(\mathcal{P}'_{M \setminus W_{xc}}, x) = 0$, since otherwise the situation of x can only improve.

1. $\mathcal{P}'_{W_{xc}} \leftarrow \mathcal{Q}_{W_{xc}}$ (initialization).
2. Let Z be the set of candidates who got 1 from each voter in $\mathcal{Q}_{M \setminus W_{xc}}$. Recall that $|Z| = k$, and also $x \notin Z$.
3. Foreach $P'_i \in \mathcal{P}'_{M \setminus W_{xc}}$:
   (a) Let Y_i be the set of candidates who got 1 in P'_i. $|Y_i| = k$. Also, by our assumption, $x \notin Y_i$.
   (b) Foreach $y \in Y_i \setminus Z$
      i. If there exists a vote $P_j \in \mathcal{P}'_{W_{xc}}$ and a candidate $z \in Z \setminus Y_i$ such that $s(P_j, y) = 1$ and $s(P_j, z) = 0$, then change the places of y and z in the vote P_j
4. Return $\mathcal{P}'_{W_{xc}}$.

Note that when processing the preference list $P'_i \in \mathcal{P}'_{M \setminus W_{xc}}$, if we changed the places of y and z in the vote $P_j \in \mathcal{P}'_{W_{xc}}$, then after the change we have $s(\{P'_i, P_j\}, y) = s(\{Q'_i, Q_j\}, y)$, and $s(\{P'_i, P_j\}, z) = s(\{Q'_i, Q_j\}, z)$, where Q'_i is the i-th vote in the profile $\mathcal{Q}_{M \setminus W_{xc}}$, and Q_j is the j-th vote in the profile $\mathcal{Q}_{W_{xc}}$.

If for all $P'_i \in \mathcal{P}'_{M \setminus W_{xc}}$ and for all $y \in Y_i \setminus Z$ we have changed between the places of y and some z in some vote $P_j \in \mathcal{P}'_{W_{xc}}$, then we get that for all $b \neq x$, $s(\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}, b) = s(\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, b)$. Since this equation also holds for x, we have found a profile, $\mathcal{P}'_{W_{xc}}$, under which x wins the election.

Now suppose that there exists a vote $P'_i \in \mathcal{P}'_{M \setminus W_{xc}}$ and a candidate $y \in Y_i \setminus Z$ such that there do not exist $z \in Z \setminus Y_i$ and $P_j \in \mathcal{P}'_{W_{xc}}$ as described above. It follows that the score of y increased relative to the score of any $z \in Z \setminus Y_i$ in profile $\mathcal{Q}_{W_{xc}}$ less than $|M \setminus W_{xc}|$ times. Therefore, for each $z \in Z \setminus Y_i$, $s(\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}, y) \leq s(\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, z) < s(\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, x) = s(\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}, x)$. Also, since for all $z \in Z \setminus Y_i$, $s(\{P'_i\}, z) = 0$ and for all $Q'_{i'} \in \mathcal{Q}_{M \setminus W_{xc}}$, $s(\{Q'_{i'}\}, z) = 1$, it follows that $s(\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}, z) < s(\mathcal{P}_H \cup \mathcal{Q}_{M \setminus W_{xc}} \cup \mathcal{Q}_{W_{xc}}, z)$. And so, we get in this case as well that x is the winner under $\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}$. **(End of proof of claim.)** □

We conclude the proof of the theorem by the remark that if we found a candidate $x \neq c$ and a non-empty coalition W_{xc} such that x is β-feasible for W_{xc}, then c is not in the β-core. Otherwise, c is in the β-core. □

Theorem 10. *There exists a polynomial-time algorithm solving the Simplified-Bucklin-α-Core problem.*

Proof. Suppose we are given an instance $I = (E, M, c)$ of the Simplified-Bucklin-α-Core problem, where $E = (C, V, \mathcal{P})$. The algorithm first checks whether $c \in F(M)$ by solving the Simplified-Bucklin-Coalitional Manipulation problem (see [15] for the algorithm for this). It then goes over all the other candidates; for each candidate $x \neq c$ it computes the maximal set $W_{xc} \subseteq M$ of colluders who prefer x to c. If $W_{xc} \neq \emptyset$, we need to check whether x is α-feasible for W_{xc}. We first introduce some new notation. Let $B(a, k, \mathcal{Q})$ denote the number of votes in the profile $\mathcal{Q}$ that rank the candidate a in the k first places. Let d be the minimal integer such that $B(x, d, \mathcal{P}_H) + |W_{xc}| > \frac{1}{2} \cdot |V|$. Let $\mathcal{P}''_{M \setminus W_{xc}}$ be a "profile" of $M \setminus W_{xc}$ such that $B(x, d, \mathcal{P}''_{M \setminus W_{xc}}) = 0$; and for all $y \neq x$, $B(y, 1, \mathcal{P}''_{M \setminus W_{xc}}) = |M \setminus W_{xc}|$.[7] In the following claim we prove that it is enough to check whether x can win vs. $\mathcal{P}''_{M \setminus W_{xc}}$ to determine whether x is α-feasible for W_{xc}.

Claim. x is the winner under $\mathcal{P}_H \cup \mathcal{P}''_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}$ where $\mathcal{P}''_{M \setminus W_{xc}}$ is as defined above, and $\mathcal{P}^x_{W_{xc}}$ is some fixed (real) profile for W_{xc} which ranks x in the first position in all the votes, if and only if for all profiles $\mathcal{P}'_{M \setminus W_{xc}}$ x is the winner under $\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}$.

Proof. Suppose x is the winner under $\mathcal{P}_H \cup \mathcal{P}''_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}$. Let $\mathcal{P}'_{M \setminus W_{xc}}$ be any profile. Let d be as defined above. Since for all $y \neq x$, $B(y, d, \mathcal{P}'_{M \setminus W_{xc}}) \leq |M \setminus W_{xc}|$, we have that $g'_y := \lfloor \frac{1}{2}|V| \rfloor - B(y, d, \mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}}) \geq \lfloor \frac{1}{2}|V| \rfloor - B(y, d, \mathcal{P}_H) - |M \setminus W_{xc}| = \lfloor \frac{1}{2}|V| \rfloor - B(y, d, \mathcal{P}_H \cup \mathcal{P}''_{M \setminus W_{xc}}) =: g''_y$. Therefore, $k'_y := \min\{g'_y, |W_{xc}|\} \geq \min\{g''_y, |W_{xc}|\} =: k''_y$. Since x is the winner under $\mathcal{P}_H \cup \mathcal{P}''_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}$, we have for all $y \neq x$, $B(y, d, \mathcal{P}^x_{W_{xc}}) \leq k''_y \leq k'_y$. So, by definition of k'_y, $B(y, d, \mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}) \leq \lfloor \frac{1}{2}|V| \rfloor$. On the other hand, $B(x, d, \mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}) \geq B(x, d, \mathcal{P}_H \cup \mathcal{P}^x_{W_{xc}}) = B(x, d, \mathcal{P}_H) + |W_{xc}| > \frac{1}{2}|V|$ (where the last inequality follows from the definition of d). And so, x wins under $\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}$.

For the opposite direction, suppose that for every profile $\mathcal{P}'_{M \setminus W_{xc}}$, x is the winner under $\mathcal{P}_H \cup \mathcal{P}'_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}$, where $\mathcal{P}^x_{W_{xc}}$ is some fixed profile of W_{xc} with x ranked at the top of each vote. Let y, $y \neq x$, be any candidate. $\mathcal{P}^x_{W_{xc}}$ makes x win also vs. the profile $\mathcal{P}^{y,x}_{M \setminus W_{xc}}$ where y is ranked first and x is ranked last in all the preferences. Since the score of x under $\mathcal{P}_H \cup \mathcal{P}^{y,x}_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}$ is d, it follows that the score of y under $\mathcal{P}_H \cup \mathcal{P}^{y,x}_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}$ is greater than d. Therefore, $B(y, d, \mathcal{P}^x_{W_{xc}}) \leq \min\{|W_{xc}|, \lfloor \frac{1}{2}|V| \rfloor - B(y, d, \mathcal{P}_H) - |M \setminus W_{xc}|\}$. Hence, the score of y under $\mathcal{P}_H \cup \mathcal{P}''_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}$ is $> d$. On the other hand, by definition of d and $\mathcal{P}^x_{W_{xc}}$, $B(x, d, \mathcal{P}_H \cup \mathcal{P}''_{M \setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}) = B(x, d, \mathcal{P}_H \cup \mathcal{P}^x_{W_{xc}}) = B(x, d, \mathcal{P}_H) + |W_{xc}| > \frac{1}{2} \cdot |V|$. So, the score of x under

[7] Of course, if $m > 2$ then $\mathcal{P}''_{M \setminus W_{xc}}$ cannot be a real profile; rather it is just a score specification. As before, we use the notation as if it were a real profile, for simplicity.

$\mathcal{P}_H \cup \mathcal{P}''_{M\setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}$ is d. Hence, x is the winner under $\mathcal{P}_H \cup \mathcal{P}''_{M\setminus W_{xc}} \cup \mathcal{P}^x_{W_{xc}}$. **(End of proof of claim.)** □

We now resume the proof of Theorem 10. With the above claim in hand, we can compute in polynomial time whether x is α-feasible for the coalition W_{xc}, in the following way. First compute d as the minimal integer such that $B(x, d, \mathcal{P}_H) + |W_{xc}| > \frac{1}{2}|V|$. Then define for each $y \neq x$, $g''_y = \lfloor\frac{1}{2}|V|\rfloor - B(y, d, \mathcal{P}_H) - |M \setminus W_{xc}|$, and $k''_y = \min\{g''_y, |W_{xc}|\}$. If $\sum_{y\neq x} k''_y < (d-1)|W_{xc}|$, then there does not exist a profile $\mathcal{P}'''_{W_{xc}}$ making x win under $\mathcal{P}_H \cup \mathcal{P}''_{M\setminus W_{xc}} \cup \mathcal{P}'''_{W_{xc}}$ (see [15] for details of the algorithm). Otherwise, we build the profile $\mathcal{P}^x_{W_{xc}}$ as follows. We first put x on top of all the preferences of $\mathcal{P}^x_{W_{xc}}$. Then for all $i = 1, \ldots, m-1$ we put the candidate c_i in the next k''_{c_i} available places in the votes of $\mathcal{P}^x_{W_{xc}}$, such that $B(c_i, d, \mathcal{P}^x_{W_{xc}}) \leq k''_{c_i}$, until we fill all the critical places. By Claim 4, if we have found a manipulation, then it works for all profiles $\mathcal{P}'_{M\setminus W_{xc}}$, and if we have not found a manipulation, then there does not exist a manipulation that works for all the profiles $\mathcal{P}'_{M\setminus W_{xc}}$. That is, we have found a manipulation if and only if $x \in F_\alpha(W_{xc})$. If we have found some $x \neq c$ such that $W_{xc} \neq \emptyset$ and $x \in F_\alpha(W_{xc})$, then c is not in the α-core. Otherwise, c is in the α-core. □

Theorem 11. *There exists a polynomial-time algorithm solving the Simplified-Bucklin-β-Core problem.*

Proof. We are given an instance $I = (E, M, c)$ of the Simplified-Bucklin-β-Core problem, where $E = (C, V, \mathcal{P})$. The algorithm first checks whether $c \in F(M)$ by solving the Simplified-Bucklin-Coalitional Manipulation problem. Then it goes over all the other candidates. For each candidate $x \neq c$ we compute the maximal set $W_{xc} \subseteq M$ of the colluders who prefer x to c. If $W_{xc} \neq \emptyset$, we need to check whether x is β-feasible for W_{xc}. Let $B(a, k, \mathcal{Q})$ again denote the number of votes in the profile $\mathcal{Q}$ that rank the candidate a in k first places. Let d be the minimal integer such that $B(x, d, \mathcal{P}_H) + |W_{xc}| > \frac{1}{2}|V|$. Let us build the profile $\mathcal{Q}_{M\setminus W_{xc}}$ as follows. Assume w.l.o.g. that the candidates $\{c_1, \ldots, c_{m-1}\} = C \setminus \{x\}$ are sorted according to $B(c_i, d, \mathcal{P}_H)$, in the descending order. In every vote in $\mathcal{Q}_{M\setminus W_{xc}}$, the candidates $c_1, \ldots, c_d$ are ranked at the first d places (for example, $c_1 \succ c_2 \succ \ldots \succ c_d$), and the rest of the candidates are ranked in some arbitrary order. In the next claim we prove that this is enough to check whether we can make x win vs. the profile $\mathcal{Q}_{M\setminus W_{xc}}$ to see whether x is β-feasible for W_{xc}.

Claim. There exists a profile $\mathcal{P}''_{W_{xc}}$ such that x is the winner under $\mathcal{P}_H \cup \mathcal{Q}_{M\setminus W_{xc}} \cup \mathcal{P}''_{W_{xc}}$ (where $\mathcal{Q}_{M\setminus W_{xc}}$ is as defined above) if and only if for each profile $\mathcal{P}'_{M\setminus W_{xc}}$ there exists a profile $\mathcal{P}'_{W_{xc}}$ such that x is the winner under $\mathcal{P}_H \cup \mathcal{P}'_{M\setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}$.

Proof. One direction is trivial: if for each profile $\mathcal{P}'_{M\setminus W_{xc}}$ there exists a profile $\mathcal{P}'_{W_{xc}}$ such that x is the winner under $\mathcal{P}_H \cup \mathcal{P}'_{M\setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}$, then, in particular, for $\mathcal{Q}_{M\setminus W_{xc}}$ there exists a profile $\mathcal{P}''_{W_{xc}}$ such that x is the winner under $\mathcal{P}_H \cup \mathcal{Q}_{M\setminus W_{xc}} \cup \mathcal{P}''_{W_{xc}}$.

For the other direction, suppose that there exists a profile $\mathcal{P}''_{W_{xc}}$ such that x is the winner under $\mathcal{P}_H \cup \mathcal{Q}_{M\setminus W_{xc}} \cup \mathcal{P}''_{W_{xc}}$. Let $\mathcal{P}'_{M\setminus W_{xc}}$ be any profile. For all $y \in C \setminus \{x\}$ define $g_y = \lfloor\frac{1}{2}|V|\rfloor - B(y, d, \mathcal{P}_H) - B(y, d, \mathcal{Q}_{M\setminus W_{xc}})$. And let $k_y = \min\{g_y, |W_{xc}|\}$. Since x wins under $\mathcal{P}_H \cup \mathcal{Q}_{M\setminus W_{xc}} \cup \mathcal{P}''_{W_{xc}}$, we have $\sum_{y\neq x} k_y \geq (d-1)|W_{xc}|$.

Similarly, define $g'_y = \lfloor \frac{1}{2}|V| \rfloor - B(y, d, \mathcal{P}_H) - B(y, d, \mathcal{P}'_{M \setminus W_{xc}})$. And define $k'_y = \min\{g'_y, |W_{xc}|\}$. It is enough for us to show that $\sum_{y \neq x} k'_y \geq \sum_{y \neq x} k_y$, since then we have $\sum_{y \neq x} k'_y \geq (d-1)|W_{xc}|$, and so we can fill the critical places in the votes of $\mathcal{P}'_{W_{xc}}$ by enumerating the candidates and placing each candidate in the next position in k'_y of the votes. We would like first to transform the above problem (whether $\sum_{y \neq x} k'_y \geq \sum_{y \neq x} k_y$) into an equivalent one. We define $h_y = \frac{1}{2}|V| - g_y, h'_y = \frac{1}{2}|V| - g'_y$. Also define $A = \lfloor \frac{1}{2}|V| \rfloor - |W_{xc}|$. Furthermore, define $t_y = \frac{1}{2}|V| - k_y = \max\{A, h_y\} = \max\{A, B(y, d, \mathcal{P}_H) + B(y, d, \mathcal{Q}_{M \setminus W_{xc}})\}$, and $t'_y = \frac{1}{2}|V| - k'_y = \max\{A, h'_y\} = \max\{A, B(y, d, \mathcal{P}_H) + B(y, d, \mathcal{P}'_{M \setminus W_{xc}})\}$. It is easy to see that $\sum_{y \neq x} k'_y \geq \sum_{y \neq x} k_y$ if and only if $\sum_{y \neq x} t'_y \leq \sum_{y \neq x} t_y$.

So, our goal is to prove that $\sum_{y \neq x} t'_y \leq \sum_{y \neq x} t_y$. Denote, for shortness, $s = |M \setminus W_{xc}|$. Then, by definition, for all $i \leq d$, $B(c_i, d, \mathcal{Q}_{M \setminus W_{xc}}) = s$, and, for all $i > d$, $B(c_i, d, \mathcal{Q}_{M \setminus W_{xc}}) = 0$. Also, for $y \in C \setminus \{x\}$, denote $l_y = B(y, d, \mathcal{P}'_{M \setminus W_{xc}})$. By definition, for all y, $0 \leq l_y \leq s$, and $\sum_{y \neq x} l_y \leq d \cdot s$. We divide the proof into 3 cases.

Case 1. $A \geq B(c_1, d, \mathcal{P}_H) + s$. Then $\sum_{y \neq x} t_y = A(m-1) = \sum_{y \neq x} t'_y$.

Case 2. $A \leq B(c_d, d, \mathcal{P}_H)$. Let us denote $D = \{c_1, \ldots, c_d\}$, $S = \{y \in C \setminus (D \cup \{x\}) \mid B(y, d, \mathcal{P}_H) \geq A\}$, and $T = C \setminus (D \cup S \cup \{x\})$. We have

$$\sum_{y \neq x} t_y = \sum_{i=1}^{d} B(c_i, d, \mathcal{P}_H) + d \cdot s + \sum_{y \in S} B(y, d, \mathcal{P}_H) + A \cdot |T|.$$

On the other hand,

$$\begin{aligned}
\sum_{y \neq x} t'_y &= \sum_{i=1}^{d} [B(c_i, d, \mathcal{P}_H) + l_{c_i}] + \sum_{y \in S} [B(y, d, \mathcal{P}_H) + l_y] \\
&\quad + \sum_{y \in T} \max\{A, B(y, d, \mathcal{P}_H) + l_y\} \\
&\leq \sum_{y \in D \cup S} [B(y, d, \mathcal{P}_H) + l_y] + \sum_{y \in T} [A + l_y] \\
&= \sum_{y \in D \cup S} B(y, d, \mathcal{P}_H) + A|T| + \sum_{y \neq x} l_y \\
&\leq \sum_{y \in D \cup S} B(y, d, \mathcal{P}_H) + A|T| + ds = \sum_{y \neq x} t_y.
\end{aligned}$$

Case 3. $B(c_d, d, \mathcal{P}_H) < A < B(c_1, d, \mathcal{P}_H) + s$. Recall that $D = \{c_1, \ldots, c_d\}$. Denote $Z^* = \{z \in D \mid B(z, d, \mathcal{P}_H) + s > A\}$, and $Y^* = D \setminus Z^* = \{y \in D \mid B(y, d, \mathcal{P}_H) + s \leq A\}$.

Case 3.a. $Y^* \neq \emptyset$. Then

$$\sum_{y \neq x} t_y = \sum_{z \in Z^*} B(z, d, \mathcal{P}_H) + s \cdot |Z^*| + A \cdot (m - 1 - |Z^*|).$$

Whereas,

$$\sum_{y \neq x} t'_y = \sum_{z \in Z^*} \max\{A, B(z, d, \mathcal{P}_H) + l_z\} + A(m - 1 - |Z^*|)$$
$$\leq \sum_{z \in Z^*} [B(z, d, \mathcal{P}_H) + s] + A(m - 1 - |Z^*|) = \sum_{y \neq x} t_y.$$

Case 3.b. $Y^* = \emptyset$. Here we further define $Z_1 = \{z \in D \mid B(z, d, \mathcal{P}_H) + l_z > A\}$, $Y_1 = \{y \in D \mid B(y, d, \mathcal{P}_H) + l_y \leq A\}$, $Z_2 = \{z \in C \setminus (D \cup \{x\}) \mid B(z, d, \mathcal{P}_H) + l_z > A\}$, and $Y_2 = \{y \in C \setminus (D \cup \{x\}) \mid B(y, d, \mathcal{P}_H) + l_y \leq A\}$.

We have the following equalities:

$$\sum_{i=1}^{m-1} t_{c_i} = \sum_{z \in Z_1} [B(z, d, \mathcal{P}_H) + s] + \sum_{y \in Y_1} [B(y, d, \mathcal{P}_H) + s] + A|Z_2| + A|Y_2| \tag{1}$$

$$\sum_{i=1}^{m-1} t'_{c_i} = \sum_{z \in Z_1} [B(z, d, \mathcal{P}_H) + l_z] + A|Y_1| + \sum_{z \in Z_2} [B(z, d, \mathcal{P}_H) + l_z] + A|Y_2|. \tag{2}$$

Recall that we need to prove that $\sum_{i=1}^{m-1} t_{c_i} \geq \sum_{i=1}^{m-1} t'_{c_i}$. Removing the common parts from both equalities (1) and (2), we need to prove that

$$s \cdot d + \sum_{y \in Y_1} B(y, d, \mathcal{P}_H) + A|Z_2| \geq \sum_{z \in Z_1} l_z + A|Y_1| + \sum_{z \in Z_2} [B(z, d, \mathcal{P}_H) + l_z] \tag{3}$$

Here we also have two cases.

Case 3.b.i. $|Z_2| \geq |Y_1|$. Denote $Y_1 = \{y_1, \ldots, y_{|Y_1|}\}$ and $Z_2 = \{z_1, \ldots, z_{|Y_1|}, z_{|Y_1|+1}, \ldots, z_{|Z_2|}\}$. For each i, $1 \leq i \leq |Y_1|$ we have $B(y_i, d, \mathcal{P}_H) \geq B(z_i, d, \mathcal{P}_H)$ due to the sorting of the candidates we assume. Therefore,

$$\sum_{i=1}^{|Y_1|} B(y_i, d, \mathcal{P}_H) \geq \sum_{i=1}^{|Y_1|} B(z_i, d, \mathcal{P}_H). \tag{4}$$

Also,

$$s \cdot d \geq \sum_{z \in Z_1} l_z + \sum_{z \in Z_2} l_z \tag{5}$$

Plugging (4) and (5) into (3), we get that it is enough to prove that $A|Z_2| \geq A|Y_1| + \sum_{i=|Y_1|+1}^{|Z_2|} B(z_i, d, \mathcal{P}_H)$, or, equivalently, that $A(|Z_2| - |Y_1|) \geq \sum_{i=|Y_1|+1}^{|Z_2|} B(z_i, d, \mathcal{P}_H)$. And indeed, this inequality is true since for each i, $B(z_i, d, \mathcal{P}_H) \leq B(c_d, d, \mathcal{P}_H) < A$.

Case 3.b.ii. $|Z_2| < |Y_1|$. Since $D = Z_1 \cup Y_1$ and $Z_1 \cap Y_1 = \emptyset$, we have $d = |Z_1| + |Y_1|$. So, rephrasing inequality (3), we have to prove that

$$\begin{aligned} &s|Z_1| + s|Y_1| + \sum_{y \in Y_1} B(y, d, \mathcal{P}_H) + A|Z_2| \\ &\geq \sum_{z \in Z_1} l_z + A|Y_1| + \sum_{z \in Z_2} [B(z, d, \mathcal{P}_H) + l_z]. \end{aligned} \tag{6}$$

Since for each i, $1 \leq i \leq m-1$ it holds that $s \geq l_{c_i}$, we have $s(|Z_1| + |Z_2|) \geq \sum_{z \in Z_1} l_z + \sum_{z \in Z_2} l_z$. By subtracting it from (6) and moving $A|Z_2|$ to the right-hand side, we get that it is enough to prove that

$$\begin{aligned} &s(|Y_1| - |Z_2|) + \sum_{y \in Y_1} B(y, d, \mathcal{P}_H) \\ &\geq A(|Y_1| - |Z_2|) + \sum_{z \in Z_2} B(z, d, \mathcal{P}_H). \end{aligned}$$

Let us denote $Z_2 = \{z_1, \ldots, z_{|Z_2|}\}$ and $Y_1 = \{y_1, \ldots, y_{|Z_2|}, y_{|Z_2|+1}, \ldots, y_{|Y_1|}\}$. For each $i, 1 \leq i \leq Z_2$, $B(y_i, d, \mathcal{P}_H) \geq B(z_i, d, \mathcal{P}_H)$ because of our sorting. Hence, $\sum_{i=1}^{|Z_2|} B(y_i, d, \mathcal{P}_H) \geq \sum_{i=1}^{|Z_2|} B(z_i, d, \mathcal{P}_H)$. It remains to prove that $s(|Y_1| - |Z_2|) + \sum_{i=|Z_2|+1}^{|Y_1|} B(y_i, d, \mathcal{P}_H) \geq A(|Y_1| - |Z_2|)$. And indeed, we assume that $Y^* = \emptyset$, i.e., for each j, $1 \leq j \leq d$, $B(c_j, d, \mathcal{P}_H) + s > A$. In particular, for each i, $|Z_2| + 1 \leq i \leq |Y_1|$, $B(y_i, d, \mathcal{P}_H) + s > A$. Therefore, $\sum_{i=|Z_2|+1}^{|Y_1|} [B(y_i, d, \mathcal{P}_H) + s] > A(|Y_1| - |Z_2|)$, which completes the proof of the claim. **(End of proof of claim.)** □

If we have found a candidate $x \neq c$ and a non-empty coalition W_{xc} such that x is β-feasible for W_{xc}, then by definition, c is not in the β-core. Otherwise, c is in the β-core. □

Theorem 12. *Let R be the Plurality with Runoff voting rule. There exists a polynomial-time algorithm solving the R-β-Core problem.*

Proof. Suppose we are given an instance $I = (E, M, c)$ of the R-β-Core problem, where R is the Plurality with Runoff voting rule, and $E = (C, V, \mathcal{P})$. The algorithm first checks whether $c \in F(M)$ by solving the R-Coalitional Manipulation problem (see [17] for the algorithm for this). Then it goes over all the other candidates. For each candidate $x \neq c$ it computes the maximal set $W_{xc} \subseteq M$ of the colluders who prefer x to c. If $W_{xc} \neq \emptyset$, we need to check whether x is β-feasible for W_{xc}. To do so, we iterate over all candidates $a \neq x$, and for each a, we check whether W_{xc} can make x win vs. the profile $\mathcal{P}^{a,x}_{M \setminus W_{xc}}$ where everybody ranks a first, x last, and the other candidates in some arbitrary order. Then we iterate over all pairs of candidates (a, b), and for each pair (a, b) and for all $i = 1, \ldots, |M \setminus W_{xc}| - 1$, we check whether W_{xc} can make x win vs. the profile $\mathcal{P}^{i,a,b,x}_{M \setminus W_{xc}}$ where i voters in $M \setminus W_{xc}$ rank a first, the $|M \setminus W_{xc}| - i$ remaining voters rank b first, all these voters rank x last, and the other candidates are ranked in some arbitrary order. In the next claim we prove that it is enough to check these profiles to determine whether x is β-feasible for W_{xc}.

Claim. If:

1. for each $a \neq x$, there exists a profile $\mathcal{P}''_{W_{xc}}$ such that x is the winner under $\mathcal{P}_H \cup \mathcal{P}^{a,x}_{M\setminus W_{xc}} \cup \mathcal{P}''_{W_{xc}}$, and
2. for each pair (a,b) (where $x \notin \{a,b\}$) and for each i, $1 \leq i \leq |M \setminus W_{xc}| - 1$, there exists a profile $\mathcal{P}'''_{W_{xc}}$ such that x is the winner under $\mathcal{P}_H \cup \mathcal{P}^{i,a,b,x}_{M\setminus W_{xc}} \cup \mathcal{P}'''_{W_{xc}}$,

then x is β-feasible. Otherwise x is not β-feasible.

Proof. It is clear that if (1) there is a candidate $a \neq x$, such that for each profile $\mathcal{P}'_{W_{xc}}$ x is not a winner under $\mathcal{P}_H \cup \mathcal{P}^{a,x}_{M\setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}$, or (2) there is a pair of candidates (a,b) and i, $1 \leq i \leq |M \setminus W_{xc}| - 1$, such that for each profile $\mathcal{P}'_{W_{xc}}$ x is not a winner under $\mathcal{P}_H \cup \mathcal{P}^{i,a,b,x}_{M\setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}$, then by definition, x is not β-feasible for W_{xc}. For the other direction, suppose that for each $a \neq x$, there exists a profile $\mathcal{P}''_{W_{xc}}$ such that x is the winner under $\mathcal{P}_H \cup \mathcal{P}^{a,x}_{M\setminus W_{xc}} \cup \mathcal{P}''_{W_{xc}}$, and for each (a,b) and i there exists a profile $\mathcal{P}'''_{W_{xc}}$ such that x is the winner under $\mathcal{P}_H \cup \mathcal{P}^{i,a,b,x}_{M\setminus W_{xc}} \cup \mathcal{P}'''_{W_{xc}}$. Let $\mathcal{P}'_{M\setminus W_{xc}}$ be any profile for $M \setminus W_{xc}$. We will show that there exists a profile $\mathcal{P}'_{W_{xc}}$ such that x is the winner of the election under $\mathcal{P}_H \cup \mathcal{P}'_{M\setminus W_{xc}} \cup \mathcal{P}'_{W_{xc}}$. Let $y, z \in C$ be the winners of the first round under the partial profile $\mathcal{P}_H \cup \mathcal{P}'_{M\setminus W_{xc}}$. Denote by $\gamma(\mathcal{P}, b)$ the plurality score of candidate b under the profile $\mathcal{P}$. Suppose w.l.o.g. that

$$\gamma(\mathcal{P}_H \cup \mathcal{P}'_{M\setminus W_{xc}}, y) \geq \gamma(\mathcal{P}_H \cup \mathcal{P}'_{M\setminus W_{xc}}, z). \tag{7}$$

Denote by $N(\mathcal{P}, a, b)$ the number of votes in the profile $\mathcal{P}$ who prefer a over b. We divide the proof into 3 cases:

Case 1. $x \notin \{y, z\}$. Let us consider the profile $\mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}}$, where $i = |M \setminus W_{xc}| - \gamma(\mathcal{P}'_{M\setminus W_{xc}}, z)$. That is, we choose i such that $\gamma(\mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}}, z) = \gamma(\mathcal{P}'_{M\setminus W_{xc}}, z)$. Let $\mathcal{P}'''_{W_{xc}}$ be the profile such that x is the winner under $\mathcal{P}_H \cup \mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}} \cup \mathcal{P}'''_{W_{xc}}$. Let x and b be the two candidates who proceed to the second round under $\mathcal{P}_H \cup \mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}} \cup \mathcal{P}'''_{W_{xc}}$. Here we have 3 cases:

Case 1.a. $b = y$. We may assume that $\mathcal{P}'''_{W_{xc}}$ has in the first places only x's and y's (if it does not, then we can change it appropriately, and the winners in the first and the second round will not change). Then under $\mathcal{P}_H \cup \mathcal{P}'_{M\setminus W_{xc}} \cup \mathcal{P}'''_{W_{xc}}$ the only possible winners of the first round are y, z and x. Denote for brevity, $\mathcal{Q}^1 = \mathcal{P}_H \cup \mathcal{P}'_{M\setminus W_{xc}} \cup \mathcal{P}'''_{W_{xc}}$, and $\mathcal{Q}^2 = \mathcal{P}_H \cup \mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}} \cup \mathcal{P}'''_{W_{xc}}$. Since x is ranked last by all the voters in $\mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}}$, we have $\gamma(\mathcal{Q}^1, x) \geq \gamma(\mathcal{Q}^2, x)$. Also, by the definition of $\mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}}$, we have $\gamma(\mathcal{Q}^1, y) \leq \gamma(\mathcal{Q}^2, y)$ and $\gamma(\mathcal{Q}^1, z) = \gamma(\mathcal{Q}^2, z)$. As mentioned before, x is one of the winners of the first round under $\mathcal{Q}^2$. It follows that also under the profile $\mathcal{Q}^1$ x will be one of the winners of the first round. By assumption (7), and by the inequality $\gamma(\mathcal{P}'''_{W_{xc}}, y) \geq \gamma(\mathcal{P}'''_{W_{xc}}, z)$, we have $\gamma(\mathcal{Q}^1, y) \geq \gamma(\mathcal{Q}^1, z)$, and so the second winner of the first round is y. Now, $N(\mathcal{P}'_{M\setminus W_{xc}}, x, y) \geq 0 = N(\mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}}, x, y)$. So, since x beats y in the second round under the profile $\mathcal{Q}^2$, x also beats y in the second round under $\mathcal{Q}^1$. So, we found a profile ($\mathcal{P}'''_{W_{xc}}$), such that x is the winner under $\mathcal{Q}^1 = \mathcal{P}_H \cup \mathcal{P}'_{M\setminus W_{xc}} \cup \mathcal{P}'''_{W_{xc}}$.

Case 1.b. $b = z$. Here we may assume that $\mathcal{P}'''_{W_{xc}}$ contains in the first places only x's and z's. Here we also have that the only possible winners of the first round under $\mathcal{Q}^1$ are x, y and z. Again, we have $\gamma(\mathcal{Q}^1, x) \geq \gamma(\mathcal{Q}^2, x)$, $\gamma(\mathcal{Q}^1, y) \leq \gamma(\mathcal{Q}^2, y)$ and $\gamma(\mathcal{Q}^1, z) = \gamma(\mathcal{Q}^2, z)$. So, since under $\mathcal{Q}^2$ the winners of the first round are x and z, we have that under $\mathcal{Q}^1$ the winners are also x and z. x beats z in the second round under $\mathcal{Q}^2$, and $N(\mathcal{P}'_{M\setminus W_{xc}}, x, z) \geq 0 = N(\mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}}, x, z)$. Hence, x beats z in the second round under $\mathcal{Q}^1$ as well.

Case 1.c. $b \notin \{y, z\}$. We may assume that $\mathcal{P}'''_{W_{xc}}$ contains in the first places only x's and b's. The only possible winners of the first round under $\mathcal{Q}^1$ are y, z, x and b. We have the following inequalities: $\gamma(\mathcal{P}'_{M\setminus W_{xc}}, x) \geq 0 = \gamma(\mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}}, x)$, $\gamma(\mathcal{P}'_{M\setminus W_{xc}}, b) \geq 0 = \gamma(\mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}}, b)$, $\gamma(\mathcal{P}'_{M\setminus W_{xc}}, y) \leq \gamma(\mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}}, y)$, and $\gamma(\mathcal{P}'_{M\setminus W_{xc}}, z) = \gamma(\mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}}, z)$. And so, since x and b are the winners of the first round under $\mathcal{Q}^2$, they are also the winners of the first round under $\mathcal{Q}^1$. In the second round, x beats b under the profile $\mathcal{Q}^2$, and $N(\mathcal{P}'_{M\setminus W_{xc}}, x, b) \geq 0 = N(\mathcal{P}^{i,y,z,x}_{M\setminus W_{xc}}, x, b)$. It follows that x beats b in the second round under the profile $\mathcal{Q}^1$ as well.

Case 2. $x = y$. So x and z are the two winners of the first round under the partial profile $\mathcal{P}_H \cup \mathcal{P}'_{M\setminus W_{xc}}$. Recall that $\mathcal{P}^{z,x}_{M\setminus W_{xc}}$ is a profile where everybody in $M \setminus W_{xc}$ ranks z first, x last, and the other candidates in some arbitrary order. Let $\mathcal{P}''_{W_{xc}}$ be the profile such that x is the winner under $\mathcal{P}_H \cup \mathcal{P}^{z,x}_{M\setminus W_{xc}} \cup \mathcal{P}''_{W_{xc}}$. Denote, for shortness, $\mathcal{Q}^3 = \mathcal{P}_H \cup \mathcal{P}'_{M\setminus W_{xc}} \cup \mathcal{P}''_{W_{xc}}$ and $\mathcal{Q}^4 = \mathcal{P}_H \cup \mathcal{P}^{z,x}_{M\setminus W_{xc}} \cup \mathcal{P}''_{W_{xc}}$. Let x and d be the two candidates who are the winners of the first round under $\mathcal{Q}^4$.

Case 2.a. $d = z$. Here we can assume that $\mathcal{P}''_{W_{xc}}$ contains only x's and z's in the first places. Hence, under $\mathcal{Q}^3$, x and $z = d$ are the winners of the first round.

Case 2.b. $d \neq z$. We can assume that $\mathcal{P}''_{W_{xc}}$ contains only x's and d's in the first places. Therefore, under $\mathcal{Q}^3$ the only possible winners of the first round are x, z and d. Since under $\mathcal{Q}^4$ x and d proceed to the second round, and due to the fact that in $\mathcal{P}^{z,x}_{M\setminus W_{xc}}$ z is ranked on top of all the preferences, we have $\gamma(\mathcal{Q}^3, d) \geq \gamma(\mathcal{Q}^4, d) > \gamma(\mathcal{Q}^4, z) \geq \gamma(\mathcal{Q}^3, z)$, and $\gamma(\mathcal{Q}^3, x) \geq \gamma(\mathcal{Q}^4, x) > \gamma(\mathcal{Q}^4, z) \geq \gamma(\mathcal{Q}^3, z)$. Therefore, under $\mathcal{Q}^3$, x and d are the winners of the first round.

Case 2 (continued). Now, as x is the winner of the second round under $\mathcal{Q}^4$, and since $N(\mathcal{P}'_{M\setminus W_{xc}}, x, d) \geq 0 = N(\mathcal{P}^{z,x}_{M\setminus W_{xc}}, x, d)$, we have $N(\mathcal{Q}^3, x, d) \geq N(\mathcal{Q}^4, x, d) > N(\mathcal{Q}^4, d, x) \geq N(\mathcal{Q}^3, d, x)$. Hence, x beats d in the second round under $\mathcal{Q}^3$, and so x wins the election.

Case 3. $x = z$. This case is handled similarly to Case 2. **(End of proof of claim.)** □

If we found $x \neq c$ such that $W_{xc} \neq \emptyset$ and $x \in F_\beta(W_{xc})$, then by definition $c \notin \beta$-core. Otherwise, $c \in \beta$-core. □

5 Conclusions and Future Work

In this paper we have provided a computational analysis of the following question: given a coalition of manipulative voters, which candidate should they manipulate in favor of, given that they might not have identical preferences. To perform our analysis we have used the notions of α- and β-core, which—under various assumptions—describe the sets of candidates that the manipulators can support in a stable manner (i.e., without running the risk of breaking the coalition).

Our main results are the following. The complexity of determining membership in the α- and β-cores is at least as high as the complexity of the constructive coalitional manipulation problem for the same rule. On the other hand, for the several prominent voting rules for which coalitional manipulation is easy, we have also provided polynomial time algorithms for determining membership in the α- and β-cores. One direction for future work is to extend the above research to other voting rules. Another interesting direction is to try to find voting rules that produce an outcome in the core, if the core is non-empty, when applied to the truthful preferences of the voters. Yet another direction is to investigate the computational complexity of finding a voting profile of the colluders which is a strong Nash equilibrium, if one exists.

6 Acknowledgments

Michael Zuckerman and Jeffrey S. Rosenschein were partially funded by Israel Science Foundation grant #898/05, Israel Ministry of Science and Technology grant #3-6797, and, along with Vincent Conitzer, United States-Israel Binational Science Foundation grant #2006-216. Piotr Faliszewski was Supported by AGH University of Technology Grant no. 11.11.120.865, by Polish Ministry of Science and Higher Education grant N-N206-378637, and by Foundation for Polish Science's program Homing/Powroty. Vincent Conitzer acknowledges NSF IIS-0812113, IIS-0953756, and CCF-1101659, as well as an Alfred P. Sloan fellowship, for support.

References

1. R. Aumann. Acceptable points in general cooperative n-person games. *Contributions to the Theory of Games*, 4:287–324, 1959.
2. Y. Bachrach, E. Elkind, and P. Faliszewski. Coalitional voting manipulation: A game-theoretic perspective. In *The Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI 2011)*, pages 49–54, 2011.
3. J. Bartholdi and J. Orlin. Single Transferable Vote resists strategic voting. *Social Choice and Welfare*, 8:341–354, 1991.
4. J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989.
5. V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *JACM*, 54(3), June 2007.
6. P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. How hard is bribery in elections? *Journal of AI Research*, 35:485–532, 2009.
7. P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Using complexity to protect elections. *Communications of the ACM*, 53(11):74–82, 2010.
8. P. Faliszewski and A. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.
9. A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.

10. S. Obraztsova and E. Elkind. On the complexity of voting manipulation under randomized tie-breaking. In *The Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI 2011)*, pages 319–324, 2011.
11. S. Obraztsova, E. Elkind, and N. Hazon. Ties matter: Complexity of voting manipulation revisited. In *The Tenth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, pages 71–78, 2011.
12. B. Peleg and P. Sudholter. *Introduction to the Theory of Cooperative Games*. Kluwer Academic Publishers, Boston, 2003.
13. M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.
14. L. Xia, V. Conitzer, and A. D. Procaccia. A scheduling approach to coalitional manipulation. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC'10)*, pages 275–284, 2010.
15. L. Xia, M. Zuckerman, A. D. Procaccia, V. Conitzer, and J. S. Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *The Twenty-First International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pages 348–353, Pasadena, California, July 2009.
16. M. Zuckerman, O. Lev, and J. S. Rosenschein. An algorithm for the coalitional manipulation problem under maximin. In *The Tenth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, pages 845–852, Taipei, Taiwan, 2011.
17. M. Zuckerman, A. D. Procaccia, and J. S. Rosenschein. Algorithms for the coalitional manipulation problem. *AI Journal*, 173(2):392–412, February 2009.