

Complexity of (Non-)Convergence in Iterative Voting

Paul W. Goldberg
University of Oxford
Oxford, United Kingdom
paul.goldberg@cs.ox.ac.uk

Marios Mavronicolas
University of Cyprus
Nicosia, Cyprus
mavronicolas.marios@ucy.ac.cy

Tomasz Was
University of Oxford
Oxford, United Kingdom
tomasz.was@cs.ox.ac.uk

ABSTRACT

Iterative voting is a well-studied model of repeated decision-making, introduced in 2010 by Meir *et al.* [19], where *strategic* voters may repeatedly revise their votes, given information about other voters' interim votes, before convergence to a stable state where no voter has an incentive to revise. Despite considerable previous convergence and non-convergence results for iterative voting under various voting rules and information and behavioral assumptions in the last 15 years, the computational complexity of detecting convergence and non-convergence has not been explored before.

In this work, we present the *first* such complexity results. Specifically, we establish, as our main results, the NP-completeness of the following two decision problems about plurality elections when an arbitrary group of voters can revise their votes as long as the updates are *direct* and *beneficial* for each member of the group:

- Does a given election converge to a *strong Nash equilibrium* within at most ℓ revision steps? We exhibit instances where a strong Nash equilibrium does exist but is unreachable by any number of such steps.
- Does a given election create voting *cycles* of ℓ revision steps?

We also prove general results for *Pareto efficient* voting rules. Specifically, for two voters and three candidates, if in each step a single voter updates their vote using the natural TB heuristic [16], then for every Pareto-efficient voting rule there is no cycle of length more than two. In contrast, when both voters update their votes using the same heuristic simultaneously, and we have $m \geq 3$ candidates, every rule satisfying a slight refinement of Pareto efficiency can end up in a cycle of length m .

KEYWORDS

Iterative Voting, Computational Complexity, Strong Nash Equilibrium

ACM Reference Format:

Paul W. Goldberg, Marios Mavronicolas, and Tomasz Was. 2026. Complexity of (Non-)Convergence in Iterative Voting. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/MDLZ8000>



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/MDLZ8000>

1 INTRODUCTION

Voting has long been the classical method for aggregating diverse individual preferences into a collective social decision. Voting systems have more recently gained interest in the context of AI and multi-agent systems. Voting usually consists of a single *round* of ballot submissions by *voters* and results in a winning *candidate* as the outcome. It has long been known that voting rules are typically susceptible to *manipulation* by voters, whereby voters misreport their preferences. *Iterative voting*, initiated in an influential work by Meir *et al.* [19] fifteen years ago, represents a voting scenario in which voters may revise their ballots in response to ballots submitted previously by their peers. After each iteration, voters have the opportunity to manipulate in pursuit of their own interest; potentially, a new winner is derived as a result of the manipulations, which initiates the next iteration, and so on. The voting procedure terminates at a given round once it reaches a *stable* state, in which no voter has an incentive to change their round's ballot. Since [19], iterative voting has attracted considerable interest and attention; see, for example, [2, 3, 9, 10, 13, 14, 16, 18, 22, 24–28, 32].

A widespread assumption in this literature is that voters begin with their truthful preferences. One motivation for this assumption is that without it, stable states are usually trivial to compute: typically a stable state results from taking an arbitrary candidate c and having every voter rank c in the top position. Consequently, any candidate can win, even an unpopular one. If we start at the truthful state, it raises interesting questions about whether, for example, some given candidate can win (or lose) via a sequence of deviations, or whether a sequence of deviations can cycle, or converge to a stable state. The answers to these questions depend on the voting rule in use, and what sorts of deviations are allowed. In this work, we focus on computational issues related to the convergence or non-convergence of the iterative voting process. We present the first hardness results for the following types of question:

- *What is the computational complexity of finding convergence to a stable state in iterative voting?*
- *What is the computational complexity of detecting cycles in iterative voting?*

1.1 Contribution

We begin by analysing a scenario in which not only a single voter, but an arbitrary group of voters can change their ballots simultaneously, as long as every member of the group benefits from the new winner. Additionally, we require that the action of each voter in the group is *direct*, by which we mean that the new winner is placed on the top of each modified ballot. Such dynamics are said to follow *group beneficial and direct actions* [19]. We show that in this setting, under the plurality voting rule (in which the winner is simply the candidate that is most often ranked first), deciding if a

stable state can be reached within at most ℓ steps, when starting from the truthful ballots, is NP-complete (Theorem 3.3). Furthermore, we observe that there are instances where a stable state is unreachable from the truthful one, even though a stable state exists in general (Theorem 3.1).

Next, in the same setting of group beneficial and direct actions, we study the decision problem of the existence of cycles. We extend the basic voting model by assuming that voters are *weighted*: each voter i has a *weight* $w_i > 0$. An additional extension to the voting model consists of assuming that each candidate j is supported by some *fixed voters* with total weight s_j : they exhibit a “stubborn” voting behavior, having already decided to cast ballots voting for candidate j in all iterations. Then, the plurality winner is the candidate with the highest sum of weights of (active) voters that placed it at the top of their ballots increased by the score from the fixed voters. Thus, we arrive at a formulation of the decision problem about group beneficial and direct cycles, whose instances include the voting instance equipped with weights and fixed scores, and the length of the sought cycle. We show that the cycle problem is NP-complete (Theorem 4.1).

We next consider a model in which the voters update their ballots using a simple heuristic called *top-bottom* (TB) [16]. Specifically, in each step, the most preferred candidate of a given voter is placed on top of the ranking, and, if the current winner is a different candidate, it is placed at the bottom of the ranking. We analyze two classes of voting rules in this context, namely *Pareto efficient* voting rules and a refinement of this class that we introduce, namely *second-choice Pareto efficient* rules. We show that for two agents and three candidates, if in each step only a single voter updates their vote, then for every Pareto efficient voting rule there is no cycle of length more than two (Theorem 5.3). In contrast, if both voters can update their votes simultaneously, then for $m \geq 3$ candidates, every second-choice Pareto efficient rule may end up in a cycle of length m (Theorem 5.4). These results offer insights that may lead to further complexity results in the style of Theorems 3.3 and 4.1.

1.2 Related Work

The study of complexity issues in voting was initiated in the influential work of Bartholdi *et al.* [11] back in 1989. There has followed an impressive volume of complexity results about voting problems in general; typical examples concern the determination of the winner in a one-shot election and the corresponding power of manipulative actions, control, and bribery — see the complexity surveys in [29, Chapters 4 and 5]. To our knowledge, the presented NP-hardness results are the *first* hardness results about convergence and non-convergence in iterative voting. Neither has earlier work on iterative voting processes (cf. [17, Chapter 10]) for deciding allocations of public goods considered issues of complexity.

The closest relatives to our complexity results in the area of iterative voting are a couple of NP-hardness results for decision problems other than detecting convergence and non-convergence. In the context of opinion polls announced by polling agencies for iterative elections, Baumeister *et al.* [2, Theorems 3.1, 3.2, 3.4 and 4.1] consider a few decision problems related to the influence by polls constrained by limits on manipulability, such as not differing too much from the correct poll. Extending the notion of *possible*

and *necessary* winners in a one-shot election [15] to the setting of iterative voting, a recent manuscript by Mousseau *et al.* [21] considers the decision problem of the existence of a possible (resp., necessary) iterative winner, defined as a candidate that is an iterative winner in *some (all)* sequence of improvement steps; they prove contrasting complexity results: deciding the existence of a necessary iterative winner is polynomial-time solvable [21, Theorem 2] whereas deciding the existence of a possible iterative winner is NP-hard [21, Theorem 3]. Notably, both [2] and [21] use reductions from Exact Cover by 3 Sets, or X3C [8]. X3C has been used in NP-hardness reductions in various other settings of voting; for example, Obraztsova *et al.* [23, Theorems 1, 2 and 3] look at plurality elections with *truth-biased* voters under *k-approval* scoring rules for $k > 2$, and establish the NP-hardness of deciding the existence of a Nash equilibrium with or without a particular winner.

Koolyk *et al.* [16, Section 5] showed that several rules (e.g., Maximin, Copeland and Ranked Pairs) may not converge under best-response dynamics and several heuristics, including TB. To our knowledge, no other (non-)convergence results are known for TB. A few convergence and non-convergence results are known for other action classes; for example, the action class $M1$ [10, Section 2.2] converges under every voting rule [10, Theorem 1]; the action class *k-pragmatist* [28] converges under any *positional scoring rule* paired with the lexicographic tie-breaking rule [28, Theorem 5]; the action class $M2$ [10, Section 2.2] converges under any positional scoring rule, Copeland and Maximin [10, Theorem 2].

Obraztsova *et al.* [25] identify two interesting conditions on dynamics, namely *Function Monotonicity* and *Set Monotonicity*, that are sufficient for their convergence — see [25, Theorems 1 and 2], respectively. They proceed to identify several classes of voting rules, including Positional Scoring Rules, Maximin, Copeland and Bucklin, for which at least one of the conditions holds.

The model with weighted users and fixed scores, which we consider in Section 4, was introduced to iterative voting in the original paper of Meir *et al.* [19], along with a counter-example to convergence. Voter dynamics involving weighted voters under the plurality rule was further studied in [6, 7].

Under various voting rules, strong Nash equilibria were shown to be closely connected to the concept of *Condorcet winner*, i.e., a candidate that is preferred over each other candidate by at least half of the voters [30].

2 PRELIMINARIES

For an integer $k \in \mathbb{N}$, let $[k]$ denote the set $\{1, \dots, k\}$.

2.1 Elections and Voting Rules

For a set of candidates C , we denote by $\mathcal{L}(C)$ the set of all linear orders over C . For a set of agents (or *voters*) V , a *profile* $\mathbb{P} \in \mathcal{L}(C)^V$, is a collection of linear orders over C , one for each agent, also known as *ballots*. $\mathcal{P}_{C,V}$ denotes the set of all profiles for given C and V .

An *election instance* is a 3-tuple $\mathcal{E} = (C, V, \mathbb{P}^*)$ where the profile \mathbb{P}^* , called the *preference profile* or the *truthful profile*, includes the *preference orders* of the agents, which represent their *true* preferences. Agents might misreport their preferences; then, the reported profile will be different from the truthful one. For a ballot $\succ \in \mathcal{L}(C)$ and candidate $c \in C$, we denote the *position* of c in the ordering \succ by

$\text{pos}_{\succ}(c) = |\{c' \in C : c' \succ c\}| + 1$. We let $\text{top}(\succ)$ denote candidate c such that $\text{pos}_{\succ}(c) = 1$.

A *voting rule* (also called a social choice function) is a function, r , that takes a profile \mathbb{P} as input and returns a single *winner* $r(\mathbb{P}) \in C$ (we consider only resolute rules in this paper). *Scoring rules* constitute a class of voting rules defined using a *scoring function* $w: [C] \rightarrow \mathbb{R}$ that assigns a score to each candidate based on their position on the ballot (the scoring function is usually assumed to be nonincreasing). The winner is a candidate with the highest summed scores across all ballots in the profile; that is, $r_w((\succ_i)_{i \in V}) = \arg \max_{c \in C} \sum_{i \in V} w(\text{pos}_{\succ_i}(c))$. If there is more than one such candidate, a *tie-breaking* mechanism has to be used. Throughout the paper, we use *lexicographic tie-breaking* to obtain a unique winner: we assume a fixed ordering \triangleright of candidates in C and take the candidate that is highest in \triangleright among those with the highest score. *Plurality* is the scoring rule with scoring function $w(i) = \mathbb{1}_{\{i=1\}}$ (where $\mathbb{1}$ is a characteristic function).

2.2 State Graphs and Action Classes

Following [5], a *state graph* is a graph $G = (\mathcal{P}_{C,V}, \mathcal{A})$, whose set of nodes is the set of all profiles, or *states*, and an arc $(\mathbb{P}, \mathbb{P}') \in \mathcal{A}$ means that some agents in profile \mathbb{P} can strategically change their ballots, resulting in profile \mathbb{P}' . We assume that \mathcal{A} is irreflexive, i.e., $(\mathbb{P}, \mathbb{P}) \notin \mathcal{A}$, for every $\mathbb{P} \in \mathcal{P}_{C,V}$.

An *action class* is a function that takes as input an election instance $\mathcal{E} = (C, V, \mathbb{P}^*)$ and a voting rule r , and outputs a state graph.

We will consider the following action classes. Here, and throughout the paper, we will use the notational convention in which $\mathbb{P} = (\succ_i)_{i \in V}$, $\mathbb{P}' = (\succ'_i)_{i \in V}$, and $\mathbb{P}^* = (\succ_i^*)_{i \in V}$.

- *Beneficial* actions are those in which a single agent can change their ballot if the winner in the new profile is more preferred over the previous winner. Formally, $(\mathbb{P}, \mathbb{P}') \in \mathcal{A}$ if and only if there is $j \in V$ such that $r(\mathbb{P}') \succ_j^* r(\mathbb{P})$ and $\succ_i = \succ'_i$, for every $i \in V \setminus \{j\}$.
- For *beneficial and direct* actions [19] we additionally require that the new winner is also on the top of the new ballot of the agent that makes the change. Formally, $(\mathbb{P}, \mathbb{P}') \in \mathcal{A}$ if and only if there is $j \in V$ such that $r(\mathbb{P}') \succ_j^* r(\mathbb{P})$, $\text{pos}_{\succ'_j}(r(\mathbb{P}')) = 1$, and $\succ_i = \succ'_i$, for every $i \in V \setminus \{j\}$.

Both beneficial as well as beneficial and direct actions assume that voters, while being myopic, have full knowledge of the current profile and can perfectly reason which changes of the profile will be beneficial for them. In contrast, in the following action class, we assume that voters follow a simple heuristic to update their ballots in each step.

- In *top-bottom* actions [16], or TB, for short, an agent puts the currently winning candidate at the bottom of their new ballot and their most preferred candidate on the top. Formally, $(\mathbb{P}, \mathbb{P}') \in \mathcal{A}$ if and only if there is $j \in V$ such that $\text{top}(\succ'_j) = \text{top}(\succ_j^*)$, $\text{pos}_{\succ'_j}(r(\mathbb{P})) = |C|$, and $\succ_i = \succ'_i$, for every $i \in V \setminus \{j\}$.

The action classes above, allowing only one agent to change their ballot at each step, can be extended to allow *groups* of agents to change their ballots.

- For *group beneficial* actions, $(\mathbb{P}, \mathbb{P}') \in \mathcal{A}$ if and only if for every $j \in V$ either $r(\mathbb{P}') \succ_j^* r(\mathbb{P})$ or $\succ_j = \succ'_j$.

- For *group beneficial and direct* actions, $(\mathbb{P}, \mathbb{P}') \in \mathcal{A}$ if and only if for every $j \in V$ either $r(\mathbb{P}') \succ_j^* r(\mathbb{P})$ and $\text{pos}_{\succ'_j}(r(\mathbb{P}')) = 1$, or $\succ_j = \succ'_j$.

We note that without the assumption that actions should be direct, group beneficial actions are very permissive and allow, for example, a voter to cast a ballot that is reversing its true preference order as long as, due to the modifications in ballots of other voters, the winner changes to a more preferred candidate from the perspective of this voter.

Finally, for action classes that represent simple heuristics with which voters update their ballot, such as TB, it makes sense to talk about *simultaneous actions*, where all voters change their ballots accordingly at the same time, in contrast to the group actions, in which a group of voters coordinate to change the outcome for a more preferred one:

- In *simultaneous* TB actions all agents for which the currently winning candidate is not their top choice put it at the bottom of their new ballot and then put their most preferred candidate on the top. Formally, $(\mathbb{P}, \mathbb{P}') \in \mathcal{A}$ if and only if for every $i \in V$, if $r(\mathbb{P}) \neq \text{top}(\succ_i^*)$, then $\text{top}(\succ'_i) = \text{top}(\succ_i^*)$, $\text{pos}_{\succ'_i}(r(\mathbb{P})) = |C|$, and $\succ_i = \succ'_i$, otherwise.

A profile $\mathbb{P} \in \mathcal{P}_{C,V}$ is a *stable state* if there is no profile $\mathbb{P}' \in \mathcal{P}_{C,V}$ such that $(\mathbb{P}, \mathbb{P}') \in \mathcal{A}$. In other words, stable states are equivalent to sinks in the state graph. For beneficial actions they correspond to *Nash equilibria* [12] while for group beneficial actions to *strong Nash equilibria* [1].¹

A *cycle of length ℓ* , is a sequence of ℓ pairwise distinct states in a state graph $(\mathbb{P}_1, \dots, \mathbb{P}_\ell)$ such that $(\mathbb{P}_\ell, \mathbb{P}_1) \in \mathcal{A}$ and $(\mathbb{P}_i, \mathbb{P}_{i+1}) \in \mathcal{A}$, for every $i \in [\ell - 1]$.

3 FASTEST CONVERGENCE

In this section, we analyse how fast we can converge to a stable state when starting from the truthful state. In other words, what is the shortest sequence of actions in the state graph from the truthful state to a sink? We focus on the plurality rule and group beneficial and direct actions.

Since both Sections 3 and 4 consider plurality, where the winner is solely determined by the top choices in the ballots of the agents, without loss of generality we will simplify the notion of the state graph in these two sections. Instead of having profile $\mathbb{P} \in \mathcal{P}_{C,V}$ as a state, we will model states as vectors of the top choices in the ballots of the voters, i.e., $\mathbf{x} \in C^V$.

3.1 No Convergence

We begin by showing that there are situations in which there exists a strong Nash equilibrium, but there is no possibility of reaching it through beneficial and direct actions.

THEOREM 3.1. *There exists an election (C, V, \mathbb{P}^*) that under Plurality, has a strong Nash equilibrium, but a strong Nash equilibrium is unreachable via beneficial and direct actions from the truthful state.*

¹Note that a strong Nash equilibrium usually requires that in the new strategy profile, agents changing strategy are not worse-off and at least one of them is better-off. However, in our case, the better-off agent implies that the winner has to change; since we consider strict orderings, the other agents must also be better-off.

PROOF. Consider the election $\mathcal{E} = (C, V, \mathbb{P}^*)$ with $C = \{a, b, c\}$ and $V = \{v_1, v_2, v_3, v_4, v_5\}$, with the following truthful profile \mathbb{P}^* :

v_1	v_2	v_3	v_4	v_5
$a \succ b \succ c$	$a \succ b \succ c$	$b \succ a \succ c$	$c \succ a \succ b$	$c \succ b \succ a$

Assume that ties are broken in reverse alphabetical order, i.e., c wins on ties with a and b , and b wins with a . First, observe that a is the Condorcet winner. Thus, the state where everyone votes for a is stable as there is no group of voters that can change the winner to the benefit of members of the group. Claim 1 characterizes all stable states in this election.

CLAIM 1. *The state is stable if and only if all voters v_1, v_2, v_3 , and v_4 vote for candidate a .*

PROOF. First, observe that such a state would indeed be stable. Only voters v_3 and v_5 prefer candidate b over candidate a , but if they both vote for b it will still receive less support than a for which v_1, v_2 , and v_4 would vote. Similarly, only v_4 and v_5 prefer c over a , but if they both vote for c , it is still a that wins with the support of v_1, v_2 , and v_3 .

Thus, it suffices to show that there are no other stable states. To this end, observe that a is the Condorcet winner, hence only states in which a is winning can be stable [30]. Moreover, if a is winning in a state \mathbf{x} , but at least one of the voters v_1, v_2 , or v_3 are not voting for a in \mathbf{x} , then voters v_4 and v_5 can both vote for c and c will start winning (v_4 and v_5 cannot both vote for c in \mathbf{x} as that would mean that a is not winning in \mathbf{x}). Thus, each voter v_1, v_2 , and v_3 has to vote for a in a stable state. Finally, if a is winning in a state \mathbf{x} , but v_4 is not voting for a , then voters v_3 and v_5 can both vote for b and b will start winning. Thus, indeed in every stable state, voters v_1, v_2, v_3 , and v_4 vote for candidate a . \square

Now, assume for a contradiction that there is a sequence of states $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k$ such that each consecutive step can be obtained from the previous one by a direct and beneficial action of a coalition of voters, \mathbf{x}_0 is the truthful state, and \mathbf{x}_k is stable. In the truthful state \mathbf{x}_0 , voters v_3 and v_4 vote for b and c , respectively. On the other hand, by Claim 1, they both have to vote for a in \mathbf{x}_k . Thus, there exist $t, t' \in [k]$ such that \mathbf{x}_t is the last state in which v_3 is not voting for a and $\mathbf{x}_{t'}$ the last state in which this is the case for v_4 . Because their actions to start voting for a must be direct and beneficial, it must be that c is winning in \mathbf{x}_t and b in $\mathbf{x}_{t'}$ (as otherwise, changing the winner to a would not be beneficial for v_3 and v_4 , respectively). In particular this means that $t \neq t'$. We observe the following.

CLAIM 2. *It holds that $t < t'$.*

PROOF. Observe that c can only win in some state if both v_4 and v_5 vote for c (other voters cannot vote for c as it is their least preferred candidate, and each candidate needs at least two votes to be winning). Hence, if $t > t'$, then v_4 would have to vote for c in \mathbf{x}_t (so that c can win in \mathbf{x}_t), but also v_4 has to vote for a (by the definition of t')—a contradiction. \square

CLAIM 3. *It holds that a is not winning in state $\mathbf{x}_{t'-1}$.*

PROOF. Towards a contradiction, assume a is winning in state $\mathbf{x}_{t'-1}$. By Claim 2, voter v_3 must already vote for a in $\mathbf{x}_{t'-1}$. In order for a to win, at least two voters must vote for a :

- v_4 cannot vote for a in $\mathbf{x}_{t'-1}$. Since b is winning in $\mathbf{x}_{t'}$, it would mean that v_4 is not changing its vote between $\mathbf{x}_{t'-1}$ and $\mathbf{x}_{t'}$ (as the change is not beneficial for v_4), so it is voting for a as well in $\mathbf{x}_{t'}$. But that contradicts the definition of t' .
- v_5 would never vote for a as it is its least preferred candidate

Thus, v_1 or v_2 need to vote for a in $\mathbf{x}_{t'-1}$. Since they are identical, without loss of generality, assume that v_1 is voting for a .

If v_2 is voting for a as well, the only way for b to win in $\mathbf{x}_{t'}$ would be if v_3 and v_5 start voting for b in $\mathbf{x}_{t'}$. However, we assumed that $t < t'$ was the last state in which v_3 did not vote for a . Hence, v_2 cannot vote for a , so it must vote for b (it would never vote for c as it is its least preferred candidate). Thus, $\mathbf{x}_{t'-1}$ is of the form (a, b, a, \cdot, \cdot) .

As we discussed, v_4 is not voting for a and it cannot vote for b (it is its least preferred candidate), so $\mathbf{x}_{t'-1} = (a, b, a, c, \cdot)$.

However, this means that no matter who v_5 is voting for (and it cannot be a —its least preferred candidate), this candidate is winning as a loses on ties with other candidates. A contradiction. \square

By Claim 3, it has to be c that is winning in $\mathbf{x}_{t'-1}$ (it cannot be b as b wins in $\mathbf{x}_{t'}$ and the winner has to change with each step). However, in order for c to win, both v_4 and v_5 must vote for it (other voters cannot vote for c as it is their least preferred candidate, and each candidate needs at least two votes to be winning). But then, in order for b to win in $\mathbf{x}_{t'}$, voters v_1, v_2 , and v_3 must start voting for it, but that would contradict Claim 2. Therefore, we arrive at a contradiction, which concludes the proof. \square

3.2 Finding Fastest Convergence is NP-hard

Next, we show that deciding whether there exists a path of length at most ℓ from the truthful state to a strong Nash equilibrium under group beneficial and direct actions and plurality is NP-complete. First let us formally define the problem.

\exists PATH TO STABLE STATE OF AT MOST GIVEN LENGTH UNDER GROUP DIRECT AND BENEFICIAL ACTIONS AND PLURALITY

Instance: An election $\mathcal{E} = (C, V, \mathbb{P}^*)$ and a number $\ell \in \mathbb{N}$ (in unary).

Question: Does there exist a sequence of $k + 1$ states $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k$ for $k \leq \ell$, such that $\mathbf{x}_0 = (\text{top}(\succ_i^*))_{i \in V}$ and under plurality and group direct and beneficial actions, \mathbf{x}_i is obtainable from \mathbf{x}_{i-1} , for each $i \in [k]$, and \mathbf{x}_k is a stable state?

Then, let us consider the verification problem. For general normal-form games, verification of strong Nash equilibria has been shown recently to be co-NP-complete [31]. Fortunately, it turns out that in this setting, verification of stable states can be done in polynomial time.

PROPOSITION 3.2. *Verifying if a state is stable under group beneficial and direct actions and plurality can be done in polynomial time.*

PROOF. Let us assume that in some state \mathbf{x} in which candidate a is the plurality winner, there is a coalition of voters, all preferring candidate b over a , that can change their votes to obtain state \mathbf{y} in which b is the new winner (basically, a witness that \mathbf{x} is not a stable state). Then, consider state \mathbf{z} obtained from \mathbf{x} by switching the vote

to b for all voters that prefer b over a . In state z the plurality score of b is at least as high as in y (since we switched to b at least as many voters) and the plurality score of every other candidate in z is at most as high as in y (for the same reason). Hence, if b is winning in y then it has to be winning in z as well. Conversely, if it is not winning in z then it cannot be winning in y . Therefore, if for every candidate b other than the current winner a , we check whether b can become the new winner after we switch all voters preferring b over a to vote for b , and the answer is negative, we know that the current state is stable. And if for some candidate b such action results in b being the new winner, then clearly the current state is not stable. \square

We next move on to the main result of this section.

THEOREM 3.3. \exists PATH TO STABLE STATE OF AT MOST GIVEN LENGTH UNDER GROUP DIRECT AND BENEFICIAL ACTIONS AND PLURALITY is NP-complete.

PROOF. Given a witness of ℓ steps from the truthful state and the final state x , we can verify in polynomial time whether x is stable by Proposition 3.2. Additionally, by computing the plurality winner in each of the ℓ steps we can check in polynomial time whether all actions were beneficial and direct for all voters changing their ballots. Thus, the problem is in NP.

We proceed to show NP-hardness by reduction from the NP-complete problem Restricted Exact Cover by 3-Sets (RX3C) [8]. Here, we are given a universe of n elements $\mathcal{U} = \{u_1, \dots, u_n\}$ for $n = 3k$ for some $k \in \mathbb{N}$, and a family of 3-subsets $\mathcal{S} = \{S_1, \dots, S_n\}$ such that every element $u_i \in \mathcal{U}$ belongs to exactly 3 subsets in \mathcal{S} . Let us assume that $n \geq 6$. For every $i, j \in [n]$, let us denote $S_j = \{u_{j,1}, u_{j,2}, u_{j,3}\}$ and $u_i \in S_{i,1}, S_{i,2}, S_{i,3}$. The question is whether there exists a subfamily of subsets $K \subset \mathcal{S}$ such that $|K| = k$ and K is a cover of \mathcal{U} , i.e., $\bigcup_{S_j \in K} S_j = \mathcal{U}$.

For every instance of RX3C we construct an instance of our problem as follows. Let $\ell = 2k$. There will be a total of $n + 6$ candidates: n set candidates s_1, \dots, s_n corresponding to sets in \mathcal{S} , 5 beginning candidates b_1, \dots, b_5 , one of which will be a top choice of most of the voters, but they will not play a role later on, and also one winning candidate w , which will be the Condorcet winner and the winner in the final state. Let us assume that w loses on ties with all other candidates.²

Furthermore, there will be a total of $6n + 1$ voters. First, we will have n element voters corresponding to elements of the universe, u_1, \dots, u_n . The top choice of each voter $u_i \in \mathcal{U}$ will be the beginning candidate b_1 , followed by all set candidates except for the sets to which u_i belongs, then remaining beginning candidates, the winning candidate w , and finally the set candidates for the sets to which u_i belongs, i.e.,

$$u_i : b_1 \succ s_1 \succ \dots \succ s_n \succ b_2 \succ \dots \succ b_5 \succ w \succ s_{i,1} \succ s_{i,2} \succ s_{i,3}.$$

Next, we will have $n + 4$ proposing voters p_1, \dots, p_{n+4} with identical preference orders starting with the winning candidate w . Specifically, for every $i \in [n + 4]$, we have:

$$p_i : w \succ s_1 \succ \dots \succ s_n \succ b_1 \succ \dots \succ b_5.$$

²This assumptions can be dropped, in a slightly more complicated construction in which we double the number of voters and remove one of the *proposing* voters (defined later on).

Then, we will have $n + 3$ opposing voters o_1, \dots, o_{n+3} with identical preference orders that start with b_2 and put w as well as other beginning candidates at the end of their ranking. Formally, for every $i \in [n + 3]$, we have:

$$o_i : b_2 \succ s_1 \succ \dots \succ s_n \succ w \succ b_1 \succ b_3 \succ b_4 \succ b_5.$$

Further, we have n counter voters c_1, \dots, c_n with preference orders identical to those of opposing voters, except the positions of b_2 and b_3 is switched:

$$c_i : b_3 \succ s_1 \succ \dots \succ s_n \succ w \succ b_1 \succ b_2 \succ b_4 \succ b_5.$$

Finally, we have $2n - 6$ finishing voters f_1, \dots, f_{2n-6} , half of which prefer b_4 the most and half b_5 , but all have w as their second choice. Formally, for $i \in [n - 3]$, we have:

$$f_i : b_4 \succ w \succ s_1 \succ \dots \succ s_n \succ b_1 \succ b_2 \succ b_3 \succ b_5,$$

and for $i \in [2n - 6] \setminus [n - 3]$, we have:

$$f_i : b_5 \succ w \succ s_1 \succ \dots \succ s_n \succ b_1 \succ b_2 \succ b_3 \succ b_4.$$

In total, we have:

- n element voters u_1, \dots, u_n ,
- $n + 4$ proposing voters p_1, \dots, p_{n+4} ,
- $n + 3$ opposing voters o_1, \dots, o_{n+3} ,
- n counting voters c_1, \dots, c_n , and
- $2n - 6$ finishing voters f_1, \dots, f_{2n-6} .

Consequently, at the truthful state, candidate w is the winner, where the full scores of candidates are as follows:

w	b_1	b_2	b_3	b_4	b_5
$n + 4$	n	$n + 3$	n	$n - 3$	$n - 3$

First, assume that there is an exact cover $K = \{S_1^*, \dots, S_k^*\}$ in the RX3C instance. For each $i \in [k]$, by s_i^* let us denote the set candidate with index corresponding to subset S_i^* from the cover. Let us show that the existence of a cover implies that there is a sequence of $2k$ steps to a stable state. In each odd step $2i - 1$, for $i \in [k]$, we will switch to voting for candidate s_i^* all opposing voters o_1, \dots, o_{n+3} , and counter voters c_1, \dots, c_{3i} . This will give us the total support for s_i^* of $(n + 3) + 3i$. In each even step $2i$, for $i \in [k - 1]$, we will switch voters that correspond to elements of set S_i^* to voting for candidate w . As we will not switch them later on, and all proposing candidates always vote for w , this will give us the total support for w of $(n + 4) + 3i$ after each such step. Thus, after each odd step $2i - 1$ the winner is s_i^* and after each even step $2i$ the winner is w . In the final $\ell = 2k$ -th step, we switch to voting for w also all of the finishing voters f_1, \dots, f_{2n-6} as well as the element voters in set S_k^* . Observe that in each step we indeed switch the vote of only those voters that benefit from the new winner.

It remains to prove that we indeed arrive at a stable state. To this end, we consider each candidate in $C \setminus \{w\}$ and show that there is no coalition that can change the winner to this candidate. Observe that in final state all element, proposing, and finishing voters are voting for w .

- For b_1 , only element voters prefer it over w . If we switch all of them to vote for b_1 , b_1 will have the score of n , but w will have score $3n - 2$ (all proposing and finishing voters).

- For b_2 , only the opposing and element voters prefer it over w . If we switch all of them to vote for b_2 , b_2 will have the score of $2n + 3$, but w will have score $3n - 2$ (as we assume $n \geq 6$, w is the winner).
- For b_3 , only the counting and element voters prefer it over w . Switching them to vote for b_3 will give it the score of $2n$, which is less than $3n - 2$ of w .
- For b_4 (or analogously for b_5), only $n - 3$ finishing voters and element voters prefer it over w . If we switch them to vote for b_4 , it will receive the score of $2n - 3$, while w will have $2n + 1$ (proposing voters and remaining finishing voters).
- Finally, each set candidate s_i is preferred over w by $n - 3$ element voters, all opposing voters, and all counter voters. Switching them all to vote for s_i will give it the score of $(n - 3) + (n + 3) + n = 3n$. On the other hand, w will have the score of $3 + (n + 4) + (2n - 6) = 3n + 1$ for 3 element voters that did not switch their vote and all proposing and finishing voters.

We now assume that there is a sequence of at most $2k$ steps that leads to a stable state and prove that this implies the existence of an exact cover in the RX3C instance. To this end, we first prove that in a stable state, all element, proposing, and finishing voters vote for w .

CLAIM 4. *A state is stable if and only if all element, proposing, and finishing voters vote for w .*

PROOF. First, observe that w is a Condorcet winner, i.e., for every other candidate $c \in C \setminus \{w\}$, there is at least $3n + 1$ voters that prefer w over c . Thus, if any other candidate c is the winner in a given state, then it is not a stable state, as we can take those $3n + 1$ voters and switch their vote to w (or make them remain voting for w). Hence, it is enough to consider states in which w is the winner.

Now, assume that at least one proposing or finishing voter is not voting for w . Then, let us take candidate s_1 and switch all opposing and counter voters to vote for s_1 as well as all element voters $u_i \in \mathcal{U} \setminus S_1$ (all of them prefer s_1 over w). As a result, we will have $(n + 3) + n + (n - 3) = 3n$ voters voting for s_1 . On the other hand, for w votes at most $3 + (n + 4) + (2n - 6) - 1 = 3n$ voters. Since we assumed that w loses on ties with all other candidates, s_1 is the new winner, hence the state was not stable.

Finally, assume that at least one element voter, u_i , is not voting for w . Let S_j be one of the sets such that $u_i \in S_j$. Then, let us switch to candidate s_j all element voters in $\mathcal{U} \setminus S_j$, all opposing voters, and all counter voters. Then, the total support for s_j will be $(n - 3) + (n + 3) + n = 3n$. On the other hand, w will be supported by at most $2 + (n + 4) + (2n - 6) = 3n$ voters, as at most 2 element voters vote for w (those in $S_j \setminus \{u_i\}$, to be exact). Since we assumed that w loses on ties with all other candidates, the state was not stable. \square

Since all element voters vote for b_1 initially, Claim 4 implies that each of them must start voting for w at some point in the $2k$ steps that lead to a strong Nash equilibrium. Take arbitrary element voter u_i and let t be the first step in which u_i started voting for w . This means that after step t the winner is w (as otherwise their action to switch to w would not be direct) and after step $t - 1$ the winner is one of the set candidates s_j such that $u_i \in S_j$ (as otherwise the action would not be beneficial for u_i).

Let $T \subset [\ell]$ be a set of indices of steps that result in w being the winner and let $K \subset \mathcal{S}$ be a set of all subsets S_j such that s_j is a winner after step $t - 1$ for some $t \in T$. Then, observe that $\bigcup_{S_j \in K} S_j = \mathcal{U}$ as all element voters u_i must start voting for w at some step. Also, $|K| \leq \ell/2 = k$. Thus, K is the exact cover in the corresponding RX3C instance. \square

4 DETECTING CYCLES IS NP-HARD

In this section, we continue our analysis of plurality elections and group beneficial and direct actions. However, instead of finding a shortest path to a stable state, we want to check whether there exists a cycle of a given length starting from a given state.

\exists CYCLE OF GIVEN LENGTH UNDER GROUP BENEFICIAL AND DIRECT ACTIONS AND PLURALITY

Instance: An election (C, V, \mathbb{P}^*) endowed with weights $\{w_i\}_{i \in V}$ and fixed scores $\{s_j\}_{j \in C}$, an initial state $x_0 \in C^V$ and a number $\ell \in \mathbb{N}$ (in unary).

Question: Is there a cycle with length ℓ and initial profile x_0 under group beneficial and direct actions and plurality for (C, V, \mathbb{P}^*) ?

We note that the problem of existence of cycles is connected to the problem of the slowest convergence. If there exists at least one cycle, then the time to convergence is infinite in the worst case. Therefore, the problem of detecting cycles is, in some sense, complementary to the problem of the fastest convergence studied in Section 3.

In what follows, we present the main result of this section—detecting cycles of a given length under group beneficial and direct actions and plurality is NP-complete.

THEOREM 4.1. \exists CYCLE OF GIVEN LENGTH UNDER GROUP BENEFICIAL AND DIRECT ACTIONS AND PLURALITY is NP-complete.

PROOF. A certificate of a positive instance is a cycle of states with length ℓ and initial state x_0 . Given the certificate, we verify as follows: For each state in the cycle, compute the winner under plurality and verify that the associated action is a direct and beneficial one. The verification can be done in time polynomial in $|V|$ and $|C|$ and linear in ℓ . Hence, \exists CYCLE WITH GIVEN LENGTH UNDER DIRECT AND BENEFICIAL ACTIONS AND PLURALITY \in NP.

To show hardness, we reduce from PARTITION. In this problem, we are given a set of n integers $\mathcal{U} = \{u_1, \dots, u_n\}$ and an integer $T \in \mathbb{N}$ such that $\sum_{u_i \in \mathcal{U}} u_i = 2T$. The question is whether there exists a partition of \mathcal{U} into two disjoint subsets that cover the whole \mathcal{U} with equal sum of integers in each subset. Formally, whether there exists $S \subseteq \mathcal{U}$ such that $\sum_{u_i \in S} u_i = T$. This is known to be NP-complete [4].

Based on an instance of PARTITION we construct the corresponding instance of \exists CYCLE OF GIVEN LENGTH UNDER GROUP BENEFICIAL AND DIRECT ACTIONS AND PLURALITY as follows.

Let there be three candidates $C = \{a, b, c\}$. We will assume that ties are broken lexicographically between them (i.e., a wins with both b and c on ties, while b wins with c on ties). Furthermore, with a slight abuse of notation, let us associate each integer from \mathcal{U} with a voter and add one additional voter u_0 . The preference order of u_0

is

$$u_0 : c \succ a \succ b,$$

while the preference order of each voter $u_i \in \mathcal{U}$ is

$$u_i : b \succ c \succ a.$$

For voter weights, we set the weight of voter u_0 to $w_0 = 1$, and the weights of the remaining voters to their integers, i.e., $w_i = u_i$, for each $i \in [n]$. Also, we set the fixed scores of candidates to $s_a = T$, $s_b = 1$, and $s_c = 0$. In the initial profile \mathbf{x}_0 all voters vote for candidate c , i.e., $\mathbf{x}_0 = (c, c, \dots, c)$. Finally, we set the length of a cycle in question to $\ell = 3$. Observe that the total plurality scores of candidates in \mathbf{x}_0 are thus $(s_a, s_b, s_c + \sum_{u_i \in \mathcal{U} \setminus \{u_0\}} w_i) = (T, 1, 2T+1)$. Hence, c is the initial winner.

First, let us show that if there is a solution to the PARTITION instance, then there is a cycle of length 3 starting from \mathbf{x}_0 . Let $S \subseteq \mathcal{U}$ be the subset such that $\sum_{u_i \in S} w_i = T$. Then, let \mathbf{x}_1 be a state obtained from \mathbf{x}_0 by changing the votes of voters in S to b . Then, the scores of candidates in \mathbf{x}_1 are $(s_a, s_b + \sum_{u_i \in S} w_i, s_c + \sum_{u_i \in \mathcal{U} \setminus \{u_0\} \setminus S} w_i) = (T, T+1, T+1)$. Since b wins with c on ties, b is the winner in \mathbf{x}_1 . Thus, indeed the action that changes \mathbf{x}_0 to \mathbf{x}_1 is beneficial for each voter u_i for $i \in [n]$ and it is direct as well.

Then, let \mathbf{x}_2 be a state obtained from \mathbf{x}_1 by changing the vote of voter u_0 from c to a . Then, the scores of candidates in \mathbf{x}_2 are $(s_a + w_0, s_b + \sum_{u_i \in S} w_i, s_c + \sum_{u_i \in \mathcal{U} \setminus \{u_0\}} w_i) = (T+1, T+1, T)$. Since a wins with b on ties, a is the winner in \mathbf{x}_2 . Thus, the action that changes \mathbf{x}_1 to \mathbf{x}_2 is beneficial (as u_0 prefers a over b) and direct.

Finally, observe that the action that changes \mathbf{x}_2 to \mathbf{x}_0 is also beneficial (as all voters prefer c over a) and direct. Therefore, we have a cycle of length 3.

We will conclude the proof by showing that if there is no solution to the PARTITION problem, then there is no cycle at all (of any length) in the corresponding instance of \exists CYCLE OF GIVEN LENGTH UNDER GROUP BENEFICIAL AND DIRECT ACTIONS AND PLURALITY.

Let us analyze what states can be reached from state \mathbf{x}_0 within a single step. Observe that the top candidate of voter u_0 , namely c , is winning in \mathbf{x}_0 . Hence, u_0 cannot change its vote as such action would not be beneficial for u_0 . For voters u_1, \dots, u_n , the current winner in \mathbf{x}_0 is their second top candidate. Thus, the only beneficial and direct action can consist of some number of voters from u_1, \dots, u_n changing their vote to their top candidate, i.e., b . Observe that the sum of weights of voters that change their vote to b has to be at least T . Otherwise, c would still be the winner and the action would not be beneficial.

For every subset $S \in \{u_1, \dots, u_n\}$ let \mathbf{x}_S be a state obtained from \mathbf{x}_0 by changing the vote of voters in S to b . As we discussed in the previous paragraph, every state reachable from \mathbf{x}_0 within one step is a state \mathbf{x}_S for some S such that $\sum_{u_i \in S} w_i \geq T$. In what follows, we will show that each such state is stable, which will imply that there is no cycle starting from \mathbf{x}_0 .

To this end, fix arbitrary S such that $\sum_{u_i \in S} w_i = X \geq T$. Observe that actually $X > T$, as otherwise we would have a solution to the PARTITION instance. Then, the scores of candidates in \mathbf{x}_S are $(s_a, s_b + \sum_{u_i \in S} w_i, s_c + \sum_{u_i \in \mathcal{U} \setminus \{u_0\} \setminus S} w_i) = (T, X+1, 2T-X+1)$. Hence, b is winning. Observe that for every voter from $\{u_1, \dots, u_n\}$ their top candidate is winning in \mathbf{x}_S , which means that they cannot change their vote or the action would not be beneficial. Thus, the only voter that can change their vote is u_0 . However, since $X > T$,

it follows that $X+1 > T+1$, thus no matter how u_0 changes its vote, the winner would not change. Hence, there is no beneficial and direct action, which means that the state \mathbf{x}_S is stable. This concludes the proof. \square

5 TB ACTIONS AND PE VOTING RULES

In this section, we focus on TB actions instead of beneficial and direct ones. We provide two general results about existence of cycles in the state graph for two classes of rules: *Pareto efficient* rules and a slight refinement of this class *second-choice Pareto efficient* rules.

Let us first provide a formal definition of Pareto efficiency [20], which requires that we never select a candidate that is preferred by every voter over another candidate.

Definition 5.1. For a set of candidates C , set of voters V , and a profile $\mathbb{P} \in \mathcal{P}_{C,V}$, a candidate $c \in C$ is said to be *Pareto dominated* by a candidate $c' \in C$, if for every voter $i \in V$ it holds that $c' \succ_i c$. A voting rule r is *Pareto efficient*, or PE for short, if it never selects a candidate dominated by another candidate.

Next, we introduce *second-choice Pareto efficiency* as a refinement of Pareto efficiency, which additionally requires that if there exists a unique candidate that for every agent is its top-choice or second-top-choice, then this candidate must be chosen.

Definition 5.2. A voting rule r is *second-choice Pareto efficient* if it is Pareto efficient and satisfies: For every set of candidates C , set of voters V , and profile $\mathbb{P} \in \mathcal{P}_{C,V}$, if there is a unique candidate $c \in C$ such that $\text{pos}_i(c) \in \{1, 2\}$ for every $i \in V$, then $r(\mathbb{P}) = c$.

We first show that under TB dynamics in the case of 2 agents and 3 candidates there is no cycle of length larger than 2 under any Pareto efficient rule, election instance, and initial state.

THEOREM 5.3. *For every election instance with 2 voters and 3 candidates and Pareto efficient voting rule, under TB actions, there is no cycle of length larger than 2.*

PROOF. Fix an arbitrary election instance with 2 voters $V = \{1, 2\}$, 3 candidates $C = \{a, b, c\}$, and a Pareto efficient rule r . Assume, for a contradiction, that there exists a cycle of length $\ell > 2$.

Let us first show that both voters need to change their votes in the cycle. Otherwise, since after the first update the single voter changing their vote puts their top-choice at the top of their ballot and it stays there, there are only two rankings of the three candidates, between which it can alternate. This contradicts the assumption that $\ell > 2$.

Then, since both voters change their ballots within the cycle, the first position in each ballot must stay the same in all profiles in the cycle and coincide with the most preferred candidate of each voter. Without loss of generality, let us assume that a is the top choice of voter 1 and b of voter 2 (we know that their top-choices are distinct, as otherwise we have a quick convergence to a stable state, since a Pareto efficient rule has to select a candidate that is unanimously at the first position in all rankings).

Thus, the cycle can alternate between four profiles, namely

$$\begin{aligned}\mathbb{P}_1 &= (a \succ b \succ c, b \succ a \succ c), \\ \mathbb{P}_2 &= (a \succ b \succ c, b \succ c \succ a), \\ \mathbb{P}_3 &= (a \succ c \succ b, b \succ c \succ a), \quad \text{and} \\ \mathbb{P}_4 &= (a \succ c \succ b, b \succ a \succ c).\end{aligned}$$

Since each time a single voter updates their ballot, and $\ell > 2$, there are only two possible cycles: $(\mathbb{P}_1, \mathbb{P}_2, \mathbb{P}_3, \mathbb{P}_4)$ and $(\mathbb{P}_4, \mathbb{P}_3, \mathbb{P}_2, \mathbb{P}_1)$.

For cycle $(\mathbb{P}_1, \mathbb{P}_2, \mathbb{P}_3, \mathbb{P}_4)$ to exist under TB actions, it has to hold that candidate c is winning in \mathbb{P}_4 (as voter 1 is putting c at the bottom of their ballot when moving to profile \mathbb{P}_1). However, this is not possible as c is Pareto dominated by a in that profile.

Similarly, for cycle $(\mathbb{P}_4, \mathbb{P}_3, \mathbb{P}_2, \mathbb{P}_1)$ to exist under TB actions, candidate c must be winning in \mathbb{P}_2 (as voter 2 is putting c at the bottom of their ballot when moving to profile \mathbb{P}_1). However, this is not possible since c is Pareto dominated by b . Thus, we arrive at a contradiction, which concludes the proof. \square

In contrast, under simultaneous TB dynamics for 2 agents and $m \geq 3$ candidates, there is an election instance and an initial state, such that for every second-choice Pareto efficient rule there is a cycle of length m .

THEOREM 5.4. *For every $m \geq 3$, there is an election instance with 2 voters and m candidates such that for every second-choice Pareto efficient voting rule, under simultaneous TB actions, there is a cycle of length m starting from the truthful state.*

PROOF. Let us consider an election instance with two voters $V = \{1, 2\}$ and m candidates $C = \{c_1, \dots, c_m\}$ and the following truthful preferences of the voters \mathbb{P}_1 :

$$1 : c_1 \succ c_2 \succ c_3 \succ \dots \succ c_m, \quad 2 : c_2 \succ c_1 \succ c_3 \succ \dots \succ c_m.$$

Let us fix an arbitrary second-choice Pareto efficient rule r .

Observe that the only candidates that are not Pareto dominated are c_1 and c_2 . Thus, one of this candidates has to be selected by r . Without loss of generality, let us assume that $r(\mathbb{P}_1) = c_1$. Then, under simultaneous TB actions, voter 2 puts c_1 at the bottom of its ballot, while 1 does not change its ballot. Hence, we obtain the following profile \mathbb{P}_2 in the next step:

$$1 : c_1 \succ c_2 \succ c_3 \succ \dots \succ c_m, \quad 2 : c_2 \succ c_3 \succ \dots \succ c_m \succ c_1.$$

Since r is second-choice Pareto efficient, it must be that $r(\mathbb{P}_2) = c_2$. Thus, under simultaneous TB actions, voter 1 puts c_2 at the bottom of its ballot, while 2 does not change its ballot. Hence, the following profile \mathbb{P}_3 in next:

$$1 : c_1 \succ c_3 \succ \dots \succ c_m \succ c_2, \quad 2 : c_2 \succ c_3 \succ \dots \succ c_m \succ c_1.$$

In general, for $i \in \{3, 4, \dots, m\}$, let us define profile \mathbb{P}_i as follows:

$$\begin{aligned}1 : c_1 \succ c_i \succ \dots \succ c_m \succ c_2 \succ c_3 \succ \dots \succ c_{i-1}, \\ 2 : c_2 \succ c_i \succ \dots \succ c_m \succ c_1 \succ c_3 \succ \dots \succ c_{i-1}.\end{aligned}$$

Since r is second-choice Pareto efficient, for every $i \in \{3, 4, \dots, m\}$ it holds that $r(\mathbb{P}_i) = c_i$. Then, under simultaneous TB actions, both voters 1 and 2 put candidate c_i at the bottom of their ballots. Consequently, for each $i \in \{3, 4, \dots, m-1\}$, after profile \mathbb{P}_i we get profile \mathbb{P}_{i+1} . Additionally, we get \mathbb{P}_1 after \mathbb{P}_m . Thus, $(\mathbb{P}_1, \mathbb{P}_2, \dots, \mathbb{P}_m)$ is a cycle under simultaneous TB actions, which concludes the proof. \square

6 DISCUSSION AND OPEN PROBLEMS

This work barely scratches the surface of the exciting yet totally unexplored area of computational complexity issues about convergence and non-convergence in iterative voting. Besides the considered settings of plurality elections with group beneficial and direct actions, there remains a plethora of (combinations) of voting rules and action classes to investigate in this exciting area and go beyond the two particular complexity results (Theorems 3.3 and 4.1) we obtained. Yet interesting extensions of these results are in sight; for example, does restricting the number of agents or the number of candidates to a constant make the two NP-complete decision problems any easier? It would also be exciting to consider special cases for the preference profiles, such as *single-peaked* preferences, and investigate whether they break the NP-completeness results. The decision problem about the existence of stable states that remain unreachable under a certain action class and voting rule, as in Theorem 3.1, offers further possibilities for complexity investigations.

An immediate open problem is whether Theorem 5.3 extends to more than 2 agents or to more than 3 candidates (or to both) under the TB dynamics. Also, a characterization of dynamics that guarantee convergence under every Pareto efficient voting rule remains a tantalizing open problem.

ACKNOWLEDGMENTS

Marios Mavronicolas' work was performed during a sabbatical leave at University of Oxford; he is partially supported by funds for the promotion of research at University of Cyprus. Tomasz Wąs and Paul Goldberg are supported by UK Engineering and Physical Sciences Research Council (EPSRC) under grant EP/X038548/1.

REFERENCES

- [1] R. J. Aumann. 1959. Acceptable Points in General Cooperative n -Person Games. In *Contributions to the Theory of Games IV*, R. D. Luce and A. W. Tucker (Eds.). Annals of Mathematics Studies, Vol. 30. Princeton University Press, 287–324.
- [2] D. Baumeister, A. K. Selker, and A. Wilczynski. 2020. Manipulation of Opinion Polls to Influence Iterative Elections. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*. 132–140.
- [3] U. Endriss, S. Obraztsova, M. Polukarov, and J. S. Rosenschein. 2016. Strategic Voting with Incomplete Information. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*. 236–242.
- [4] M. R. Garey and D. S. Johnson. 2002. *Computers and Intractability*. Vol. 29. WH Freeman.
- [5] M. Goemans, V. Mirrokni, and A. Vetta. 2005. Sink equilibria and convergence. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*. 142–151.
- [6] N. Gohar. 2012. *Manipulative Voting Dynamics*. Ph.D. Dissertation. University of Liverpool, UK.
- [7] N. Gohar, S. Noor, F. F. Babar, A. Malik, and S. Shaheen. 2019. Dynamics of manipulation in voting, veto and plurality. *Cluster Computing* 22 (2019), 7333–7345.
- [8] T. Gonzales. 1985. Clustering to Minimize the Maximum Intercluster Distance. *Theoretical Computer Science* 38 (1985), 293–306.
- [9] U. Grandi, J. Lang, A. Ozkes, and S. Airiau. 2024. Voting Behavior in One-Shot and Iterative Multiple Referenda. *Social Choice and Welfare* 63, 3 (2024), 641–675.
- [10] U. Grandi, A. Loreggia, F. Rossi, K. B. Venable, and T. Walsh. 2013. Restricted Manipulation in Iterative Voting: Condorcet Efficiency and Borda Score. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory*. 181–192.
- [11] J. J. Bartholdi III, C. A. Tovey, and M. A. Trick. 1989. The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare* 6, 3 (1989), 227–241.
- [12] J. F. Nash Jr. 1950. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences* 36, 1 (1950), 48–49.
- [13] J. Kavner, R. Meir, F. Rossi, and L. Xia. 2023. Convergence of Multi-Issue Iterative Voting Under Uncertainty. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence*. 2783–2791.

- [14] J. Kavner and L. Xia. 2021. Strategic Behavior is Bliss: Iterative Voting Improves Social Welfare. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems*. 19021–19032.
- [15] K. Konczak and J. Lang. 2005. Voting Procedures with Incomplete Preferences. In *Proceedings of the 1st Multidisciplinary Workshop on Advances in Preference Handling*.
- [16] A. Koolyk, T. Strangway, O. Lev, and J. S. Rosenschein. 2017. Convergence and Quality of Iterative Voting Under Non-Scoring Rules. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. 273–279.
- [17] J. J. Laffont. 1987. Incentives and the Allocation of Public Goods. In *Handbook of Public Economics*. Elsevier, Chapter 10, 537–569.
- [18] O. Lev and J. S. Rosenschein. 2016. Convergence of Iterative Scoring Rules. *Journal of Artificial Intelligence Research* 57 (2016), 573–591.
- [19] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. 2017. Iterative Voting and Acyclic Games. *Artificial Intelligence* 252 (2017), 100–122.
- [20] H. Moulin. 1991. *Axioms of cooperative decision making*. Number 15. Cambridge University Press.
- [21] V. Mousseau, H. Surugue, M. Tydrichová, and A. Wilczynski. 2025. On Iterative Voting Outcomes in Plurality Election. (2025). Unpublished manuscript.
- [22] L. Naamani-Dery, S. Obraztsova, Z. Rabinovich, and M. Kalech. 2015. Lie on the Fly: Iterative Voting Center with Manipulative Voters. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*. 2033–2039.
- [23] S. Obraztsova, O. Lev, E. Markakis, Z. Rabinovich, and J. S. Rosenschein. 2015. Beyond Plurality: Truth-Bias in Binary Scoring Rules. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*. 451–468.
- [24] S. Obraztsova, O. Lev, M. Polukarov, Z. Rabinovich, and J. S. Rosenschein. 2016. Non-Myopic Voting Dynamics: An Optimistic Approach. In *Proceedings of the 10th Multidisciplinary Workshop on Advances in Preference Handling*.
- [25] S. Obraztsova, E. Markakis, M. Polukarov, Z. Rabinovich, and N. R. Jennings. 2015. On the Convergence of Iterative Voting: How Restrictive Should Restricted Dynamics Be?. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*. 993–999.
- [26] M. Polukarov, S. Obraztsova, Z. Rabinovich, A. Kruglyi, and N. R. Jennings. 2015. Convergence to Equilibria in Strategic Candidacy. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*. 624–630.
- [27] Z. Rabinovich, S. Obraztsova, O. Lev, E. Markakis, and J. S. Rosenschein. 2015. Analysis of Equilibria in Iterative Voting Schemes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*. 1007–1013.
- [28] A. Reijngoud and U. Endriss. 2012. Voter Response to Iterated Poll Information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*. 635–642.
- [29] J. Rothe. 2024. *Economics and Computation – An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division* (second ed.). Springer.
- [30] M. R. Sertel and M. R. Sanver. 2004. Strong equilibrium outcomes of voting games are the generalized Condorcet winners. *Social Choice and Welfare* 22, 2 (2004), 331–347.
- [31] J. Soundy, M. T. Irfan, and H. Chan. 2025. Pure and Strong Nash Equilibrium Computation in Succinctly Representable Aggregate Games. In *Proceedings of the 41st Conference on Uncertainty in Artificial Intelligence*. 4013–4033.
- [32] A. Wilczynski. 2019. Poll-Confident Voters in Iterative Voting. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*. 2205–2212.