

Analysis of Equilibria in Iterative Voting Schemes

Zinovi Rabinovich¹
zr@zinovi.net

Svetlana Obraztsova^{2a,2b}
svetlana.obraztsova@gmail.com

Omer Lev³
omerl@cs.huji.ac.il

Evangelos Markakis⁴
markakis@gmail.com

Jeffrey S. Rosenschein³
jeff@cs.huji.ac.il

¹Independent Researcher, Jerusalem, Israel

^{2a}National Technical University of Athens, Greece and ^{2b}Tel-Aviv University, Israel

³Hebrew University of Jerusalem, Israel

⁴Athens University of Economics and Business, Greece

Abstract

Following recent studies of iterative voting and its effects on plurality vote outcomes, we provide characterisations and complexity results for three models of iterative voting under the plurality rule. Our focus is on providing a better understanding regarding the set of equilibria attainable by iterative voting processes. We start with the basic model of plurality voting. We first establish some useful properties of equilibria, reachable by iterative voting, which enable us to show that deciding whether a given profile is an iteratively reachable equilibrium is NP-complete. We then proceed to combine iterative voting with the concept of *truth bias*, a model where voters prefer to be truthful when they cannot affect the outcome. We fully characterise the set of attainable truth-biased equilibria, and show that it is possible to determine all such equilibria in polynomial time. Finally, we also examine the model of *lazy voters*, in which a voter may choose to abstain from the election. We establish convergence of the iterative process, albeit not necessarily to a Nash equilibrium. As in the case with *truth bias*, we also provide a polynomial time algorithm to find all the attainable equilibria.

1 Introduction

There are many aspects to coordination in multiagent systems that have engaged researchers in recent years, including questions related to the aggregation of multiple agents' preferences into a single system-wide choice. Researchers looking at group decision-making have explored the properties of voting schemes, which provide well-founded preference aggregation techniques; at times, voting can also provide intuitive and efficient prescriptive algorithms for resolving disagreements among participants.

Alas, as the Gibbard-Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975) famously states, voting rules are susceptible to manipulation. Given this negative result, there has been a solid body of research over the last two decades that has focused on the complexity of manipulation, as a potential barrier to the Gibbard-Satterthwaite theorem. In a different direction, and given the impossibility of having strategyproof voting rules, another body of research has emerged

on game-theoretic analysis of voting, initiated by Farquharson (1969). Viewing voters as strategic agents, it is then natural to examine the Nash equilibria (NE) of the underlying voting games, as a potential solution concept for preference aggregation scenarios.

However, Nash equilibria in this context, and without any further refinement, end up being a poor tool to predict voting behaviour. For example, even if all voters rank the same candidate last, it is still a Nash equilibrium (in most common voting rules) that all voters vote for this disliked candidate. More generally, there can be a very large number of equilibria in most voting games, and many of them are irrelevant to the effective analysis of the voting rule. Several methods have been proposed to handle this multitude of equilibria, and in this work, we focus on the following three:

- **Iterative voting:** Instead of examining all Nash equilibria, one can consider only equilibria that are “reachable” from certain initial voting profiles; that is, we can study sequences of improvement moves, where during each move one of the players who is dissatisfied with the current result can change its vote to achieve a better outcome. This process is reminiscent of a group of friends trying to find a movie that they would all like to see, or a restaurant they would all want to go to; dissatisfied with the current result, they change their declared preference seeking to modify the outcome (see, e.g., Meir et al. 2010). Several online services, such as Doodle, enable this kind of iterative preference aggregation.
- **Truth bias:** Another approach for filtering out equilibria is to add a small utility gain for voters who vote truthfully. This extra “bonus” should be small enough so that agents would still prefer to manipulate the election in cases where they can affect the outcome for their own benefit. The rationale of the model is that agents who are not pivotal tend to simply declare their true preferences. Introducing this bias in the model dramatically reduces the number of equilibria by several orders of magnitude (Thompson et al. 2013).
- **Voting with abstentions:** A different way to simulate real-world incentives is to add a small utility gain to voters who choose to abstain. The rationale is that coming to

the election may incur a cost in time, effort, etc. Hence the abstention avoids this cost, but just as in truth bias, the agents would still have an incentive to manipulate if they can affect the outcome. Such voters are also referred to as *lazy voters* (Desmedt and Elkind 2010).

Contribution: Our main goal is to study in more depth the above three techniques and their relationship, separately and combined, to Nash equilibria. We first examine iterative voting in a simple form, namely plurality voting with agents playing a best-response strategy. After first establishing some useful properties, we find that even in this straightforward model (which is known to eventually converge (Meir et al. 2010)), checking whether a given profile is a reachable Nash equilibrium is NP-hard.

We then turn to the truth-bias approach, and combine it with iterative voting. In this model, iterative voting converges to equilibria that are arguably more natural; truth bias reflects intuitive voter leanings, and eliminates many undesirable equilibria. We give a characterisation of the reachable equilibrium profiles, and detail a polynomial-time algorithm for finding all such equilibria of a game.

Finally, we examine the lazy-voter approach, combined with iterative voting. We study two versions of the iterative process, depending on whether a voter that decides to abstain at some step is allowed to come back to the election at a later step. We mostly focus on the more intuitive version where abstention is interpreted as not being able to come back at later steps. Convergence is guaranteed under this model, albeit not necessarily to a Nash equilibrium. However, we are still able to fully characterise the Nash equilibria that can be reached by this process and obtain a polynomial algorithm.

This work enhances our understanding of the effects of truth-bias and lazy-bias, which have only recently been introduced in the literature. It also serves as a means to compare and select the relevant voting model, on the basis of properties satisfied by the equilibria of each model.

Related Work

There have been many attempts to escape the multitude of Nash equilibria in voting games, and we focus on the three that have been at the forefront of recent research in the AI community and outside it. The iterative voting model we utilise is based on the one introduced by Meir et al. (2010) and expanded by additional researchers (Lev and Rosenschein 2012; Branzei et al. 2013; Kukushkin 2011; Reyhani and Wilson 2012). That model followed previous research into iterative and dynamic mechanisms, much of it summarised by (Laffont 1987).

Another line of work studies restricted dynamics in iterative processes, i.e., limitations on the allowed moves by the voters (Reijngoud and Endriss 2012; Grandi et al. 2013; Obraztsova et al. 2014; Meir, Lev, and Rosenschein 2014). However, these papers concentrate on convergence, while we study the set of equilibria obtainable in this manner.

The notion of adding a truth bias to games was introduced (for a specific case) by Laslier and Weibull (2013), and was proposed for a specific voting rule (with limited results) by Dutta and Laslier (2010). A more robust model was sug-

gested by Thompson et al. (2013), which introduced a more general framework, and contained various empirical results when using plurality in truth-biased games. The theoretical side of that work was recently enhanced by Obraztsova et al. (2013). The notion of lazy voting was studied by Desmedt and Elkind (2010), as another way of eliminating some of the undesirable Nash equilibria.

The rest of the paper is organised as follows. Definitions and notation are in Section 2. Section 3 presents our first results, on equilibrium characterisation of iterative plurality voting, including our NP-completeness result for a profile reachability decision. This is contrasted by the results of Section 4, which describe the properties of *truth-biased* iterative plurality voting, leading to a polynomial-time algorithm for equilibrium computation. Turning our attention to *lazy voters*, we study in Section 5 analogous questions for two iterative versions of this principle. We conclude in Section 6 with some final remarks. While omitting many of the proofs due to space limitations, we provide key proofs and proof outlines.

2 Definitions and Notation

We consider a set of m candidates $C = \{c_1, \dots, c_m\}$ and a set of n voters $V = \{1, \dots, n\}$. Each voter i has a strict linear *preference order* (i.e., a ranking) over C , which we denote by a_i . For notational convenience in comparing candidates, we will sometimes use \succ_i instead of a_i . When $c_k \succ_i c_j$ for some $c_k, c_j \in C$, we say that voter i prefers c_k to c_j .

At an election, each voter submits a preference order b_i , which does not necessarily coincide with a_i . We refer to b_i as the vote or ballot of voter i . The vector of submitted ballots $\mathbf{b} = (b_1, \dots, b_n)$ is called a *preference profile*. At a profile \mathbf{b} , voter i has voted truthfully if $b_i = a_i$. Any other vote from i will be referred to as a non-truthful vote. Similarly the vector $\mathbf{a} = (a_1, \dots, a_n)$ is the *truthful preference profile*, whereas any other profile is a non-truthful one.

Given a voter $i \in V$, and its vote b_i under a profile \mathbf{b} , we denote by $\text{top}(b_i)$ the top choice of the vote. A *voting rule* \mathcal{F} is a mapping that, given a preference profile \mathbf{b} , outputs a candidate $c \in C$; we write $c = \mathcal{F}(\mathbf{b})$. In this paper we will consider only Plurality¹ under lexicographic tie-breaking. This is one of the most well-studied and widely-used rules. Under Plurality, each candidate is assigned a *score* equal to the number of ballots where it has appeared as the top choice. The winner of the election is then the candidate with the maximum score. That is, the winner of the election is the candidate who appears as the top choice in the maximum number of votes. In case of ties, we assume that tie-breaking is resolved by the linear order $c_1 \succ c_2 \succ \dots \succ c_m$.

Given \mathbf{b} , let s be the maximum score achieved by a candidate. We denote by $W(\mathbf{b})$ the set of tied candidates with score equal to s , i.e., all the potential winners before tie-breaking is applied. Also, let $H(\mathbf{b})$ be the set of candidates

¹For the Plurality rule, it is enough for a voter to submit its top choice, and not a whole linear order. However, here we provide general definitions so that our models are applicable to other voting rules.

that receive $s - 1$ votes in \mathbf{b} , but would win a tie-break against any candidate in $W(\mathbf{b})$ (these are candidates who would need one extra vote to become a winner). These two sets play an important role in our analysis, as they, together, define the “runner-ups” — the candidates that can win with an additional point. Finally, we denote by $sc(c, \mathbf{b})$ the score of candidate c in \mathbf{b} .

Game theoretic considerations

In this work, we view elections as non-cooperative games. The standard way to do this is to associate a utility function u_i with every voter i , which is consistent with its true preference order. That is, we require that $u_i(c_k) \neq u_i(c_j)$ for every $i \in V$, $c_j, c_k \in C$, and also that $u_i(c_k) > u_i(c_j)$, if and only if $c_k \succ_i c_j$.

We study and compare three game-theoretic models. We refer to the first one as the *basic model* (following Meir et al. 2010) since it is the most standard approach. Under the basic model, the strategy space of each voter is the set of all linear orders, and the payoff function of voter i when its real preference is a_i is:

$$p_i(a_i, \mathbf{b}, \mathcal{F}) = u_i(c_j), \text{ if } c_j = \mathcal{F}(\mathbf{b}),$$

where \mathbf{b} is the submitted profile.

The second model is a variation of the first one, and we refer to it as the *truth-biased model*, following Thompson et al. 2013. In this model, we suppose that voters have a slight preference for voting truthfully when they cannot unilaterally affect the outcome of the election. This bias is captured by inserting a small extra payoff, when the voter votes truthfully. This extra gain is small enough so that voters may still prefer to be non-truthful in cases where they *can* affect the outcome. If \mathbf{a} is the real profile and \mathbf{b} is the submitted one, then the payoff function of voter i is given by:

$$p_i(a_i, \mathbf{b}, \mathcal{F}) = \begin{cases} u_i(c_j), & \text{if } c_j = \mathcal{F}(\mathbf{b}) \wedge a_i \neq b_i, \\ u_i(c_j) + \epsilon, & \text{if } c_j = \mathcal{F}(\mathbf{b}) \wedge a_i = b_i. \end{cases} \quad (1)$$

The value of ϵ should not effect the preference between candidates. Hence, $0 < \epsilon < \min_{i \in V} \min_{c, c' \in C} |u_i(c) - u_i(c')|$. We note

also that the exact numbers for the utilities $u_i(\cdot)$ do not matter in the analysis that follows, as long as ϵ lies within the specified range.

The third model in this paper is also a variation of the first one. We refer to it as the *model with lazy voters*, following Desmedt and Elkind 2010. In this model, it is assumed that voters have a slight preference for not coming to the election if they are not pivotal. The ability to *abstain* from the election is captured by introducing a special ballot, which \mathcal{F} simply omits from its calculations. However, the bias itself is captured similarly to the truth bias in the previously described model variation. Let \perp denote the abstention ballot. If \mathbf{a} is the real profile and \mathbf{b} is the submitted one, then the payoff function of voter i is given by:

$$p_i(a_i, \mathbf{b}, \mathcal{F}) = \begin{cases} u_i(c_j), & \text{if } c_j = \mathcal{F}(\mathbf{b}) \wedge b_i \neq \perp, \\ u_i(c_j) + \epsilon, & \text{if } c_j = \mathcal{F}(\mathbf{b}) \wedge b_i = \perp. \end{cases} \quad (2)$$

In the pathological case that the submitted profile is the vector (\perp, \dots, \perp) , we assume that no candidate is elected and each voter has a payoff of ϵ . This clearly cannot be realised as a stable state.

A Nash equilibrium in these games is a profile \mathbf{b} , where no voter can unilaterally improve its payoff, i.e.,

for every i and every b'_i , we have $p_i(a_i, \mathbf{b}, \mathcal{F}) \geq p_i(a_i, (b'_i, \mathbf{b}_{-i}), \mathcal{F})$ (\mathbf{b}_{-i} being the vector \mathbf{b} without player i 's vote). In addition, for a general non-truthful profile \mathbf{b} we will define the *best response set* of player i as: $BR(i, \mathbf{b}) = \{b'_i | \forall b p_i(a_i, (b, \mathbf{b}_{-i}), \mathcal{F}) \leq p_i(a_i, (b'_i, \mathbf{b}_{-i}), \mathcal{F})\}$

Improvement Step Dynamics

Consider an election game, either in the basic or in the truth-biased model. We focus on an iterative process, where, starting from the truthful preference profile, voters can change their strategy by making improvement steps. An *improvement step* for a voter i at a profile \mathbf{b} is a switch to another strategy (i.e., vote) b'_i , leading to the profile (b'_i, \mathbf{b}_{-i}) , in which the payoff for i is strictly higher than before. A *best response improvement step* is one in which the voter changing its strategy achieves its currently best possible payoff.

We focus on *best response dynamics*, allowing only best response improvement steps. In particular, we will refer to a *best response improvement path* (for simplicity, *improvement path*), as any sequence $(\mathbf{b}^0 \rightarrow \mathbf{b}^1 \rightarrow \dots)$ of voting profiles satisfying that for every $k \geq 1$, there exists a unique agent, say voter i , such that $\mathbf{b}^k = (b'_i, \mathbf{b}_{-i}^{k-1})$, where b'_i is a best response for voter i , and $b'_i \neq b_i^{k-1}$. We do not make any restrictions on the order in which the agents apply their improvement moves. We only assume that they start from the truthful profile, i.e., $\mathbf{b}^0 = \mathbf{a}$ (a natural starting point for such a process, as also argued by Branzei et al 2013), and then make their best response updates in an arbitrary order.

The process in general can lead to different outcomes, or may not even converge. We are interested in studying the Nash equilibria of the election games defined before, that can be reached by best response improvement paths. Notice, we do not study the equilibria of the overall dynamic setup or the effects of changing the dynamics, as, for instance, the work of Brafman-Tennenholtz 2004 or Obraztsova et al. 2014 would suggest.

3 Basic model analysis

We begin by analysing the properties of Nash equilibria that are reachable under iterative voting in the basic model. First, we recall the following property of best response improvement paths.

Lemma 1. [Quoted from Lemma 4 in (Branzei et al. 2013)] *An improvement step can only take a vote for a non-winning candidate and transfer it to the winner of the newly formed voting profile.*

We can now consider the implications of voting iterations on Nash equilibrium profiles.

Definition 1. *Given a profile \mathbf{b} , we denote by $CS(\mathbf{b})$ (and refer to it as the chasing set of \mathbf{b}), the set: $CS(\mathbf{b}) = (W(\mathbf{b}) \cup H(\mathbf{b})) \setminus \{\mathcal{F}(\mathbf{b})\}$, i.e., the candidates who could become a winner if they were to receive one additional vote.*

The following fact which we use repeatedly in the sequel follows from the analysis in (Branzei et al. 2013).

Fact 1. *Let \mathbf{b} be an equilibrium profile obtained by a sequence of improvement steps from the truthful profile \mathbf{a} . Then $\mathcal{F}(\mathbf{b}) \in W(\mathbf{a}) \cup H(\mathbf{a})$.*

Lemma 2. *Let \mathbf{b} be an equilibrium profile, obtained by a sequence of improvement steps from \mathbf{a} . Then $CS(\mathbf{b}) \subseteq W(\mathbf{a}) \cup H(\mathbf{a})$, and if $\mathbf{b} \neq \mathbf{a}$, then $CS(\mathbf{b}) \neq \emptyset$.*

The properties identified in the previous lemmas help us formulate the following necessary conditions for reachable equilibria.

Theorem 1. *Let \mathbf{b} be an equilibrium profile obtained by a sequence of improvement steps from the truthful profile \mathbf{a} . Then the following holds.*

1. *For every voter i , $\mathcal{F}(\mathbf{b}) \succ_i c$, $\forall c \in CS(\mathbf{b}) \setminus \{top(b_i)\}$.*
2. *For every voter i such that $top(b_i) \in CS(\mathbf{b})$, $top(b_i) \succ_i c$, $\forall c \in CS(\mathbf{b}) \setminus \{top(b_i)\}$.*

We can now proceed to the algorithmic question of whether a given profile can be reached by a sequence of improvement steps. We obtain the following negative result for the basic model.

Theorem 2. *Given a truthful profile \mathbf{a} and a profile \mathbf{b} , distinct from \mathbf{a} , it is NP-complete to decide if \mathbf{b} is reachable by iterative best-response updates, starting from \mathbf{a} .*

Proof Sketch. To show that the problem is in NP, it is enough to provide as a certificate the sequence of best-response updates that leads from profile \mathbf{a} to profile \mathbf{b} . It is important to note here that this is indeed a polynomial length certificate.

To prove NP-hardness, we provide a reduction from the Hitting Set (HS) problem, which is the following: we are given a set of ground elements $G = \{g_1, \dots, g_n\}$, and a family of subsets of G , $W = \{w_1, \dots, w_m\}$, $w_i \subseteq G$, $|w_i| = l_i$. We are also given a number $k \leq n$. The decision problem is to ascertain that there is a hitting set $U \subset G$, so that $|U| \leq k$, and $\forall i \in [m]$, $U \cap w_i \neq \emptyset$.

We assume that we are given an instance that satisfies: 1) $|w_1| \geq |w_i|$, $\forall i \in [m]$, 2) $|w_1| \geq 3$, and 3) $m \geq n$ (we can always pad an instance by replicating a set to satisfy this). These three assumptions do not impact the complexity of the HS problem.

Given such an instance of the HS problem, we construct an instance of our problem, i.e., a truthful profile \mathbf{a} , and a matching (non-truthful) profile \mathbf{b} , so that a sequence of iterative best-response updates going from \mathbf{a} to \mathbf{b} exists if and only if the HS instance has a solution. \square

4 Truth-Biased Iterative Voting

Having fully characterised the Nash equilibria in the iterative voting scheme under plurality, and in light of the negative result of Theorem 2, we now introduce the assumption of truth bias as described in Section 2. For the remainder of this section we will consider only truth-biased agents and investigate how this property affects the outcome of iterative voting. Once again, we begin our analysis by recalling some basic and already established properties of equilibrium profiles under truth bias, namely, the following lemma, which is proved in (Obraztsova et al 2013).

Lemma 3. *Suppose that $\mathbf{b} \neq \mathbf{a}$ is a non-truthful, Nash equilibrium profile. Let $c = \mathcal{F}(\mathbf{b})$. Then all non-truthful votes in \mathbf{b} have c as the top candidate.*

Let $s = sc(\mathcal{F}(\mathbf{a}), \mathbf{a})$. The following lemma is an easy corollary of Lemma 2 by Obraztsova et al. 2013.

Lemma 4. *Suppose that $\mathbf{b} \neq \mathbf{a}$ is a non-truthful, Nash equilibrium profile. Let $c = \mathcal{F}(\mathbf{b})$. Then $sc(c, \mathbf{b}) \leq s + 1$.*

We let W_{\succ_c} denote the set of all candidates $c_j \in W(\mathbf{a})$ such that $c_j \succ c$. H_{\succ_c} is defined similarly. It is easy to see that Fact 1 continues to hold in this model, too. Hence, the winner at an equilibrium $\mathbf{b} \neq \mathbf{a}$, belongs to the set $W(\mathbf{a}) \cup H(\mathbf{a})$. The next lemmas shed more light on each of the two possible cases (that $\mathcal{F}(\mathbf{b}) \in W(\mathbf{a})$ or $\mathcal{F}(\mathbf{b}) \in H(\mathbf{a})$) and they also highlight some important differences between the basic model and the truth-biased one.

Lemma 5. *Suppose that $\mathbf{b} \neq \mathbf{a}$ is an equilibrium profile, and that $c \in W(\mathbf{a}) \cup H(\mathbf{a})$ is the winner in \mathbf{b} . Then*

- *there is only one voter (say i) who submits a non-truthful vote in \mathbf{b} ;*
- *for every $c_j \in W(\mathbf{a}) \cup H(\mathbf{a}) \setminus \{top(a_i)\}$, $c \succ_i c_j$;*
- *If $c \in W(\mathbf{a})$, for every voter k and every $c_j \in W_{\succ_c} \setminus \{top(a_k)\}$, $c \succ_k c_j$. If $c \in H(\mathbf{a})$, for every voter k and every $c_j \in W(\mathbf{a}) \cup H_{\succ_c} \setminus \{top(a_k)\}$, $c \succ_k c_j$.*

These properties yield as direct corollaries the following:

Corollary 1. *If there exists a sequence of improvement steps that leads to the Nash equilibrium $\mathbf{b} \neq \mathbf{a}$ with $c = \mathcal{F}(\mathbf{b})$, then there exists such a chain consisting of only 1 step.*

As a consequence of Corollary 1 we can obtain an algorithm that finds all reachable Nash equilibria in time $\mathcal{O}(m^2 n^2)$. Using the following characterisation, which gives a more accurate description of attainable equilibria, we can construct an algorithm with even better running time.

Corollary 2. *Let $\mathbf{b} \neq \mathbf{a}$ be a Nash equilibrium with winner $c \in W(\mathbf{a})$. There exists a chain of improvement steps that leads to \mathbf{b} from \mathbf{a} if and only if the following conditions hold:*

1. *There is at most one candidate $c_j \neq \mathcal{F}(\mathbf{a}) \in W_{\succ_c}$ for which there exists at least one voter i , with $c_j \succ_i c$ and $c_j \neq top(a_i)$.*
2. *If no such candidate c_j , as described above exists, then there exists a voter i such that $c \succ_i c_k$ for every $c_k \in W(\mathbf{a}) \cup H(\mathbf{a}) \setminus \{top(a_i)\}$ and $top(a_i) = \mathcal{F}(\mathbf{a})$. Otherwise, there exists a voter i with $top(a_i) = c_j$ and $c \succ_i c_k$ for every $c_k \in W(\mathbf{a}) \cup H(\mathbf{a}) \setminus \{c_j\}$.*

Corollary 3. *Let $\mathbf{b} \neq \mathbf{a}$ be a Nash equilibrium with winner $c \in H(\mathbf{a})$. There exists a chain of improvement steps that leads to \mathbf{b} from \mathbf{a} if and only if the following conditions hold:*

1. *There exists at most one candidate $c_j \neq \mathcal{F}(\mathbf{a}) \in W(\mathbf{a}) \cup H_{\succ_c}$ such that there exists at least one voter i with $c_j \succ_i c$ and $c_j \neq top(a_i)$.*
2. *if no such candidate c_j , as described above exists, then there exists a voter i such that $c \succ_i c_k$ for every $c_k \in W(\mathbf{a}) \cup H(\mathbf{a}) \setminus \{top(a_i)\}$ and $top(a_i) = \mathcal{F}(\mathbf{a})$. Otherwise, there exists a voter i with $top(a_i) = c_j$ and $c \succ_i c_k$ for every $c_k \in W(\mathbf{a}) \cup H(\mathbf{a}) \setminus \{c_j\}$.*

The above characterisations entail the following.

Theorem 3. *There exists an $\mathcal{O}(mn)$ algorithm that finds all reachable Nash equilibria under truth-bias.*

5 Lazy-biased iterative voting

We now turn to consider the lazy voters model under iterative voting. We first consider a direct composition of the two concepts, where at each game stage all voting options (including abstaining) are available to a voter. Then, noting that the spirit of lazy voting would imply leaving the voting process entirely, we also investigate a different iterative process with limited voting dynamics, where abstention, if chosen, is permanent, i.e., voters cannot return to the election if they choose to abstain.

Note that NE do not always exist in the one-shot game with lazy voters (Desmedt and Elkind 2010). Nonetheless, it is possible to characterise those that do exist. In particular, for lexicographic tie-breaking, (Elkind et al. 2014) provides the following characterisation:

Theorem 4 (after (Elkind et al. 2014)). *Let \mathbf{b} be a NE under the lazy model. Then, i) \mathbf{b} consists of the truthful vote of exactly 1 voter, and abstention by all other voters; ii) if c is the top choice of the active voter in \mathbf{b} , it is ranked higher in the tie-breaking rule than all other candidates who are more preferable than c by any other voter.*

Interestingly, using Theorem 4, we can show that all NE can be reached by a sequence of improvement steps. Hence, our main finding for this model is summarised as follows:

Theorem 5. *If a game with lazy voters has a NE, then there is a sequence of improvement steps that converges to it.*

The proof of Theorem 5 is based on similar arguments as Theorem 6 below and we omit it. Since we know that not all such games possess a NE, there is no *general* guarantee that all sequences of improvement steps would converge.

Restricting the dynamics

We now consider an alternative version of iterative voting with lazy voters, which, as we will show, exhibits a different behaviour regarding convergence to stable states. Specifically, we follow the intuition that abstention expresses the fact that a voter has lost interest in the election. Hence, we feel it is more natural to study the following dynamics in the remainder of this section: if at some point in the sequence of best responses a voter decides to abstain, then he never comes back to the election.

More formally, the restriction we impose is that in any move $\mathbf{b} \rightarrow \mathbf{b}'$ along the best response improvement path, it holds that for every i , if $b_i = \perp$, then $b'_i = \perp$.

Under this restriction, we can in fact guarantee convergence, albeit not necessarily to a NE of the one-shot game, but to a state as defined below:

Definition 2. *Let A^t be the set of active voters after t steps of the iterative process, i.e., the set of voters who have not chosen to abstain. Obviously $A^t \subseteq A^{t-1}$, for every t . We say that a profile \mathbf{b} is a stable state at time t , if no voter from A^t has an incentive to change its current vote.*

Theorem 6. *Under the restricted dynamics for lazy voting, every sequence of improvement steps converges, and the final state is a stable state.*

Proof. Clearly if we have convergence, then the final state has to be a stable state. To prove convergence, we will utilise arguments from the work of Meir et al. 2010.

Consider an improvement path. We will first show that no cycles can occur within a segment of the path, located between 2 abstentions. To see this, note that Meir et al. 2010 showed that in the basic model of plurality, every sequence of best-response steps, beginning from *any* profile, always converges. This is shown for a restricted type of best response, where any improvement step *has to* give a point to the new winner. Coming back to our model, observe that under lazy-bias, any improvement step that is not an abstention gives a point to the new winner. Otherwise, the voter may gain additional utility by abstaining. Hence, after every abstention, our process evolves just like the restricted best response process in the basic model, with the remaining set of active voters. Since Meir et al. 2010 show convergence from any arbitrary profile, we cannot have cycles between any two abstentions.

Finally, since there are at most $n-1$ abstentions, there will be a final segment in the path where by the same argument as above, the active voters will converge to a stable state. \square

Observation 1. *Since we have guaranteed convergence to stable states, obviously not all stable states are Nash equilibria, as Nash equilibria do not always exist in the lazy model. Furthermore, as in the truth-biased model, there are Nash equilibria which are not reachable using a sequence of improvement steps.*

The following example demonstrates Observation 1.

Example 1. Figure 1 shows a game where no improvement path can lead to a Nash equilibrium. The tie-breaking rule is $c_1 \succ c_2 \succ c_3$, and we can see that the profile where voter 1 votes his true preference, and the other two voters abstain is the only NE. Yet, from the truthful profile of Figure 1, the only available improvement move is for voter 1 to deviate to \perp and leave the election. Afterwards, voter 2 or 3 will also leave the election, and since we are in the model where voters that leave do not return to the election, the process converges to electing c_2 , which is not a NE.

1	2	3
c_1	c_2	c_2
c_3	c_3	c_3
c_2	c_1	c_1

Figure 1: A game without convergence to a NE.

Similarly to Theorem 4, the following holds regarding the reachable stable states.

Lemma 6. *If a profile \mathbf{b} is a stable state reachable by a sequence of improvement steps, then it consists of the truthful vote of exactly 1 voter, and an abstention by all other voters.*

Comparing the two versions of dynamics we have studied, we see that the unrestricted best response dynamics with lazy voters do not generally converge, but all NE can be reached. On the other hand, the restricted dynamics manage to achieve stability, at the cost of making only a subset of

all NE reachable. Nevertheless, one more advantage in the model of restricted dynamics is that this subset of NE can be fully computed in polynomial time. This is achieved by Algorithm 1, presented below. A closer look at Algorithm 1 also shows that only under strong restrictions on the truthful profile \mathbf{a} , can a NE become unreachable. In fact, we would conjecture that the set of such profiles is small, and leave this as a question for future research.

Theorem 7. *There exists a polynomial algorithm that finds all reachable Nash equilibria from the truthful state with lazy voters with $\mathcal{O}(m^2n^2)$ complexity.*

To prove Theorem 7, we need to run Algorithm 1 over all voters. This algorithm is given a truthful profile \mathbf{a} and a truthful voter v whose top choice is z and returns “yes” if there exists a path of best-response moves to a Nash equilibrium terminating in this truthful voter voting for z (according to Lemma 6, all others abstain). The algorithm basically goes over all options of reaching a NE, as long as the requirements of Theorem 4 are satisfied, first checking whether \mathbf{a} , the initial state, is a Nash equilibrium in the usual sense (without abstentions) or not (lines 1-5), and finding the players that would enable the voter v to have a shot at being the only participant left in the game.

6 Conclusions and Future Work

Voting schemes have been challenged by the manipulability issue, which in turn paves the way for game-theoretic approaches. Thus arises a natural research direction: the study of Nash equilibria as a solution concept for preference aggregation over strategic agents. This leads to the need for characterising and computing the resulting set of equilibria.

In this paper, we investigated these issues for iterative voting under the plurality rule. Though it has been previously shown that iterative plurality voting converges to an equilibrium, the set of equilibria has not been further analysed or characterised. Our work addresses this by considering three different models of iterative voting—the basic model, along with two recently proposed tweaks, namely truth-bias and lazy-bias. We studied the properties of reachable Nash equilibria under these models and obtained characterisation results. Furthermore, we considered the algorithmic question of identifying reachable profiles. In the basic model, deciding on the reachability of a specific NE profile is NP-hard, but we obtained polynomial time algorithms for computing all reachable equilibria in the other models.

This work enhances our understanding of the effects of truth-bias and lazy-bias on plurality voting. It also serves as a means to compare and select the relevant voting model, on the basis of properties satisfied by the equilibria of each model.

7 Acknowledgments

This research has been co-financed by the European Union (European Social Fund — ESF) and Greek national funds through the Operational Program “Education and Lifelong Learning” of the National Strategic Reference Framework (NSRF) — Research Funding Program: THALES; and by ESRC grant RES-000-22-2731; and by Microsoft Research

Algorithm 1 Checking reachability of NE under lazy voting

Input: The initial profile \mathbf{a} , and the voter v to be checked, with preferences in the form $z \succ \dots$

- 1: **if** z cannot be a winning candidate in a NE according to the conditions of Theorem 4 **then**
 - 2: **return** No
 - 3: **end if**
 - 4: **if** \mathbf{a} is a NE in the (no-abstention) basic model **then**
 - 5: **if** z is a winner of \mathbf{a} **then**
 - 6: **return** Yes
 - 7: **end if**
 - 8: **if** $\exists \tilde{c} \neq z, v' \in V$ s.t. $\tilde{c} \succ_{v'} z \wedge z \succ_{v'} \mathcal{F}(\mathbf{a})$ **then**
 - 9: **return** Yes
 - 10: **end if**
 - 11: **if** \triangleright At this point, every voter whose top choice is not z , prefers $\mathcal{F}(\mathbf{a})$ to z .
 - 12: $C' \leftarrow \{c \in C \setminus \{z, \mathcal{F}(\mathbf{a})\} \text{ s. t. } \text{sc}(c, \mathbf{a}) \geq 2 \text{ or } \text{sc}(c, \mathbf{a}) = 1 \text{ and } c \succ \mathcal{F}(\mathbf{a}) \text{ in tie-breaking.}\}$ \triangleright
 - 13: Potential NE winners.
 - 14: **if** there is a voter $z \succ \dots \succ c \succ \dots \succ \mathcal{F}(\mathbf{a}) \succ \dots$ for some $c \in C'$ **then**
 - 15: **return** Yes
 - 16: **end if**
 - 17: **if** there is $\tilde{c} \neq z$, and $c \in C'$ such that there exists a vote in the form $\tilde{c} \succ \dots \succ c \succ \dots \succ \mathcal{F}(\mathbf{a}) \succ \dots$ **then**
 - 18: **return** Yes
 - 19: **end if**
 - 20: **return** No
 - 21: **end if**
 - 22: **end if**
 - 23: \triangleright From now we can assume \mathbf{a} is not a NE
 - 24: **if** z is not the winner nor a runner-up in \mathbf{a} **then**
 - 25: **return** Yes
 - 26: **end if**
 - 27: **if** z is a runner-up in \mathbf{a} **then**
 - 28: **if** For a runner up $b \neq z$ there is a voter $v \neq v$ with preference $\dots b \succ \dots \succ \mathcal{F}(\mathbf{a})$ **then**
 - 29: **return** Yes
 - 30: **end if**
 - 31: **if** $|V| \geq 4$ **then** Goto Line 6
 - 32: **else return** No
 - 33: **end if**
 - 34: **end if**
 - 35: \triangleright We can now assume $z = \mathcal{F}(\mathbf{a})$
 - 36: $\mathbf{b} \leftarrow$ Voter profile after running iterative plurality (without abstentions) while preventing v from deviating \triangleright
 - 37: from Meir et al. 2010 this is polynomial
 - 38: **if** $z = \mathcal{F}(\mathbf{b})$ **then**
 - 39: **return** Yes
 - 39: **end if**
 - 40: **if** z is not a runner-up in \mathbf{b} **then**
 - 41: **return** Yes
 - 41: **else** Goto Line 19 using \mathbf{b} instead of \mathbf{a} . $\triangleright z$ a runner-up in \mathbf{b}
 - 42: **end if**
-

through its PhD Scholarship Programme, Israel Science Foundation grant #1227/12, the Israel Ministry of Science

and Technology — Knowledge Center in Machine Learning and Artificial Intelligence grant #3-9243, the Google Inter-University Center for Electronic Markets and Auctions, and the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI); and RFFI 14-01-00156-a.

References

- Brafman, R. I., and Tennenholtz, M. 2004. Efficient learning equilibrium. *Artificial Intelligence* 159(1):27–47.
- Branzei, S.; Caragiannis, I.; Morgenstern, J.; and Procaccia, A. D. 2013. How bad is selfish voting? In *The Twenty-Seventh AAAI Conference on Artificial Intelligence*, 138–144.
- Desmedt, Y., and Elkind, E. 2010. Equilibria of plurality voting with abstentions. In *ACM EC*, 347–356.
- Dutta, B., and Laslier, J.-F. 2010. Costless honesty in voting. in 10th International Meeting of the Society for Social Choice and Welfare, Moscow.
- Elkind, E.; Markakis, E.; Obraztsova, S.; and Skowron, P. 2014. Equilibria of plurality voting: Lazy and truth-biased voters. Available at www.arxiv.org, Tech. Report number: 1409.4132, September 2014.
- Farquharson, R. 1969. *Theory of Voting*. Yale University Press.
- Gibbard, A. 1973. Manipulation of voting schemes. *Econometrica* 41(4):587–602.
- Grandi, U.; Loreggia, A.; Rossi, F.; Venable, K. B.; and Walsh, T. 2013. Restricted manipulation in iterative voting: Condorcet efficiency and borda score. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory*, 181–192.
- Kukushkin, N. S. 2011. Acyclicity of improvements in finite game forms. *International Journal of Game Theory* 40(1):147–177.
- Laffont, J.-J. 1987. Incentives and the allocation of public goods. In *Handbook of Public Economics*, volume 2. Elsevier. chapter 10, 537–569.
- Laslier, J.-F., and Weibull, J. W. 2013. An incentive-compatible Condorcet jury theorem. *Scandinavian Journal of Economics* 115(1):84–108.
- Lev, O., and Rosenschein, J. S. 2012. Convergence of iterative voting. In *The Eleventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, 611–618.
- Meir, R.; Polukarov, M.; Rosenschein, J. S.; and Jennings, N. 2010. Convergence to equilibria of plurality voting. In *The Twenty-Fourth National Conference on Artificial Intelligence*, 823–828.
- Meir, R.; Lev, O.; and Rosenschein, J. S. 2014. A local-dominance theory of voting equilibria. In *Proceedings of the 15th ACM Conference on Economics and Computation*, 313–330.
- Obraztsova, S.; Markakis, E.; Polukarov, M.; Rabinovich, Z.; and Jennings, N. R. 2014. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *The Fifth International Workshop on Computational Social Choice (COMSOC 2014)*.
- Obraztsova, S.; Markakis, E.; and Thompson, D. R. M. 2013. Plurality voting with truth-biased agents. In *The 6th International Symposium on Algorithmic Game Theory (SAGT)*, 26–37.
- Obraztsova, S.; Rabinovich, Z.; and Madunts, A. 2014. Faustian dynamics in Sarkar’s social cycle. In *The 21st European Conference on Artificial Intelligence (ECAI)*, 1071–1072.
- Reijngoud, A., and Endriss, U. 2012. Voter response to iterated poll information. In *The Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 635–644.
- Reyhani, R., and Wilson, M. 2012. Best reply dynamics for scoring rules. In *The 20th European Conference on Artificial Intelligence (ECAI)*, 672–677.
- Satterthwaite, M. A. 1975. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10(2):187–217.
- Thompson, D. R. M.; Lev, O.; Leyton-Brown, K.; and Rosenschein, J. S. 2013. Empirical analysis of plurality election equilibria. In *The Twelfth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, 391–398.