Farsighted Voting Dynamics

Svetlana Obraztsova[1], Omer Lev[2], Maria Polukarov[3], Zinovi Rabinovich[4], and Jeffrey S. Rosenschein[2]

[1] Tel-Aviv University, Israel
svetlana.obraztsova@gmail.com
[2] Hebrew University of Jerusalem, Israel
{omerl,jeff}@cs.huji.ac.il
[3] University of Southampton, UK
mp3@ecs.soton.ac.uk
[4] Mobileye Vision Technologies Ltd., Jerusalem, Israel
zr@zinovi.net

Abstract. Iterative voting has presented, in the past few years, a voting model in which a player is presented with an election poll, and changes their vote to influence the result immediately. Several extensions have been presented for this model, including some attempts to handle the uncertainty facing the players, but all of them retained the myopic assumption – players change their vote only when they believe they might be changing the outcome by their move.

This paper tackles this assumption by bounding the farsightedness of the players. Players will change their vote if they believe that if a certain number of other voters will change as well, the outcome might change. We show that players with the same farsightedness will converge to a Nash equilibrium with plurality, and with veto, even players with varying farsightedness degree will always converge. However, we show farsightedness is not necessarily a positive feature — in several cases it is better for the player to be myopic.

1 Introduction

The process by which multiple agents reach a decision has long been an important research area, dealing with elections, voting mechanisms and various techniques that allow a group of agents, each with its own preferences over a set of options, to reach a decision for all of them. In doing so, one of the main stumbling blocks has been that agents might not report their preference truthfully to the voting mechanism, and will seek to influence it by strategically reporting preferences that will change the mechanism's outcome in a way they find beneficial. Sadly, the Gibbard-Satterthwaite theorem [7,17] tells us that there is no voting system that can guarantee that voters will be truthful (i.e., no strategyproof voting mechanism).

Several methods have been proposed to deal with this issue. One of the common approaches has been to examine the complexity of the voting mechanism and its manipulations [2,19]. However, many common voting mechanisms, like plurality and veto, are very easily manipulable, and yet are still very widely used. In an attempt to improve our understanding of the ultimate outcome in the real world, Meir et al. [11] suggested examining the model of iterative voting, in which voters manipulated one-by-one (though not in a fixed order), and only if their move could make an immediate change, i.e, the voters were myopic, looking only at the next step.

The basic model suggested in Meir et al. [11] was further developed and extended in several papers by various authors, but the myopic aspect of the model remained in its various extensions, and voters were manipulating only if they believed their manipulation could benefit them directly, in the next step. In a sense, this was the last major assumption made in that paper and not further examined, due, in a large part, to the difficulty in effectively modeling non-myopic behavior.

In this paper we begin to tackle the myopic assumption and assume our voters can look ahead, using, in part, the model of uncertainty presented in [10], changed for the non-myopic setting. As in the real world, non-myopicness does not mean that voters are omnipotent: they can only see several steps ahead, and they do not know the inner preferences of other players. In a sense, this is a type of "bounded rationality" in which our rational players are still constrained by a certain bound (that might be different for each voter), limiting their ability to look ahead.

In this setting we discuss convergence for both plurality and veto (previous results in [9] show no other scoring rule converges even in myopic settings) when players use a best-response strategy (previous results [12] showed the scope for other strategies is quite limited). Furthermore, we explore whether being non-myopic is a clear advantage, and find, surprisingly, that being far-sighted might hinder a voter — a voter might be better off myopic, as more far-sighted manipulations might not pan out.

1.1 Related Work

There has been plenty of research done on attempting to analyze voting, particularly with plurality, an overview of which can be seen in [10]. Therefore, we shall focus here on the various extensions to the iterative voting model presented in [11] (which, in itself, was building on somewhat different iterative models, such as [1]). That paper assumed players which were myopic, used best-response strategy, plurality, and linear tie-breaking rules, and showed convergence in that case.

The necessity of linearity of tie-breaking rules was shown in [9], which also extended the convergence results to the veto rule (a result repeated by [16]) , while showing that no convergence can be achieved in any other scoring rule. More recently, [12] analyzed in a deeper manner the strategies that might ensure convergence in an iterative voting process (and not just best-response), a direction that was also explored in [8]. This left just myopicness as the only requirement of the original paper that has not been explored and examined.

Other research added other elements to the model: in [14], truth-bias (as modeled in [6,18,13]) and lazy-bias (as modeled in [5]) were added to the iterative process, and convergence results were quite different. Research has also been done [3] on the price of anarchy in these games, showing winner have a high score in the truthful profile .

The model most closely related to ours is the extension offered in [10]. While we adhere closer to the model of [11], the techniques used there have some applicability here. s While earlier work on locally dominant response do exist (e.g., [15,4]), it is Meir et al. [10] that offers the first analysis of the dynamics of such a response. In more detail, their paper extended the model to include uncertainty, and, in a way similar to ours, they look at uncertainty from the point of view of what the possible states could be, should some number of votes change. However, as Meir et al. [10] examines uncertainty, their paper pursues a dominant strategy for its players, who continue to play only when there is a state where a single move can change the election outcome. On the other hand, in our work, while voters continue to pursue a best-response strategy, they do not manipulate only when the outcome can be immediately changed, but rather if it could be changed in several steps. Practically, while the "locality" of the strategy is similar in both papers (i.e., considering states that are close to the current one), the meaning of those states is different, and therefore players pursue much different strategies, leading to different results.

Example 1. The difference between our model and that of [10], can be seen in the following case: Consider an iterative election with 3 candidates (tie breaking is lexicographic). Following several steps, we are now at a stage where there are 10 votes for candidate a, 9 votes for candidates b and c. All votes with a radius of 2. One voter has a preference $a \succ b \succ c$, and is currently voting for b. Using local dominance, this voter will stay at b, as all potential situations are possible and the player does not know which state it is in. As it has no dominating strategy, it stays put.

On the other hand, in the non-myopic model the voter knows it gains nothing from staying with b, as it knows the current situation, and b is not a winner in it. It would rather strengthen the winner it supports, and would change its vote to a.

2 Notation and Model

In general we will adopt the notation of [10]. We will denote a discrete set of x elements by $[x] = \{1, ..., x\}$. In particular, the set candidates and the set of voters are $M = [m]$ and $N = [n]$, respectively. For an abstract set X we will use $|X|$ to denote the number of elements in the set, e.g. $|M| = m$. Let $\pi(M)$ be the set of all linear orders over M. A *preference profile* $\mathbf{Q} \in (\pi(M))^n$ is the vector of all voter preferences, so that $Q_i \in \pi(M)$ is the (true) preference order of voter i. In particular, $Q_i(a) \in M$ is the rank of candidate $a \in M$, and $q_i = Q_i^{-1}(1)$ is the (true) top choice of voter $i \in N$.[5] We will also say that candidate $a \in M$ is preferred to candidate $b \in M$ by voter $i \in N$ and denote it by $a \succ_i b$ if $Q_i(a) < Q_i(b)$. A special preference order $\widehat{Q} \in \pi(M)$ will be used to denote a lexicographic tie-breaking when necessary.

A *partial strategic (or voting) profile* is then a set of ballots submitted by voters, $\mathbf{b} \in \pi(M)^V, V \subset N$. *A complete* strategic profile is one that corresponds to a full set of voters $V = N$. Notice that given a voting profile, the set of voters, V, who have submitted the ballots is implicit. If $b_i = Q_i$ we say that the voter $i \in V$ is truthful in $\mathbf{b} \in M^V$. In general, a voting rule determines the winner of an elections based on the submitted (complete) voting profile. However, in this paper we will be dealing solely with *scoring voting rules* characterised by a vector of parameters $\alpha = (\alpha_1, \ldots, \alpha_{m-1}, 0)$, $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_{m-1} \geq 0$. These parameters essentially give a numerical points value to each position within a ranking order, be that the preference order of a voter, $Q_i \in \pi(M)$, or a submitted ballot, $b_i \in \pi(M)$. So the highest ranked candidate would

[5] Notice that the smaller the rank of the candidate the more it is preferred.

receive α_1 points, the second highest α_2 points and so on. By aggregating these points, a winner of an election is then determined.

Formally, we will define a *scoring profile* or *state* as a partial statistic, $\mathbf{s_b}$, of a strategic profile $\mathbf{b} \in \pi(M)^V$, which assigns a score to each candidate $c \in M$, so that

$$\mathbf{s_b}(c) = \sum_{i \in V} \alpha_{b_i(c)}.$$

A *true* score of a candidate is $\mathbf{s_Q}(c) = \sum_{i \in N} \alpha_{Q_i(c)}$. Note we can view $\mathbf{s_b}$ as a vector in $\mathbb{N}^m$. Given such a score $\mathbf{s}$, the winner, $f(\mathbf{s})$, is the candidate with the maximal number of points in $\mathbf{s}$ (with the tie-breaking order $\widehat{Q}$ is used in case there are several candidates with the maximal score). In other words, for any $c \neq f(\mathbf{s}) \in M$ holds $\mathbf{s}(c) \leq \mathbf{s}(f(\mathbf{s}))$ and, furthermore, if $\mathbf{s}(c) = score(f(\mathbf{s}))$ then $\widehat{Q}(f(\mathbf{s})) < \widehat{Q}(c)$.

Now, we will concentrate on two of the most popular scoring rules: Plurality, defined by parameters $\alpha = (1, 0, \ldots, 0)$; and Veto, defined by parameters $\alpha = (1, \ldots, 1, 0)$. For these rules, a complete ranking order is superfluous as a ballot. In fact, a ballot's effect on the score profile and, hence, on the winner, is fully determined by the candidate getting a point (in Plurality) or not getting a point (in Veto). This allows us to "summarise" ballots, and for the remainder of this paper we write, with a slight abuse of notation, $b_i \in M$ for a single voter's ballot and $\mathbf{b} \subseteq M^V$ for a (partial) voting profile.

In this paper we adopt the iterative voting point of view. In other words, a voter $i \in N$ has access to a (partial) scoring profile $\mathbf{s}$ before she chooses her vote a_i. Then the vote can be strategically selected based on the outcome that would follow by concatenating a_i with $\mathbf{s}$ into new scoring profile $\mathbf{s}' = (\mathbf{s}, a_i)$. Similarly, $\mathbf{s} \setminus a_i$ is such that $\mathbf{s} = (\mathbf{s} \setminus a_i, a_i)$. In particular, standard notions of *better-response* and *dominance* are defined as follows. A candidate a_i is a *better-response* w.r.t. $\mathbf{s}$ if $f(s, a_i) \succ_i f(s, a')$ for all $a' \neq a_i$. If $f(\mathbf{s}, a) \succ_i f(\mathbf{s}, b)$ for all scoring profiles $\mathbf{s}$, then a *dominates* b from the point of view of voter i. These allow us to define *voter response function* and a *stable state*.

A *voter response function* $\rho : N \times \mathbb{N}^m \to M$ maps a scoring profile into a ballot of voter i. We will naturally shorthand $\rho_i(\mathbf{s}) = \rho(i, \mathbf{s})$. For instance, the voter best response function can be defined as follows:

Definition 1. *Formally, let* $\mathbf{s_b}$ *be a given scoring profile. For a given voter* i*, let* $\mathbf{s}_{-i} = \mathbf{s}_{\mathbf{b} \setminus i}$ *and* $\mathcal{B} \subset M$ *the set of all better responses of voter* i *to* $\mathbf{s}_{-i}$*. Then the best response function* ρ^{BR} *is such that* $\rho^{BR}(i, \mathbf{s_b}) \in \mathcal{B}$ *and for all* $b'_i \in \mathcal{B}$ *it holds* $Q_i(\rho(i, \mathbf{s_b})) < Q_i(b'_i)$*.*

In turn, a *stable state* is defined in terms of a voter response function.

Definition 2. *Let* $\mathbf{b}$ *be a voting profile of voters in some* $V \subset N$*, and* $\mathbf{s_b}$ *a scoring profile that corresponds to it. Then* $\mathbf{b}$ *is a stable voting profile and* $\mathbf{s_b}$ *is a stable state (with respect to voter response function* ρ*) if for all* $i \in V$ *it holds that* $\rho(i, \mathbf{s_b}) = b_i \in \mathbf{b}$*.*

Notice that if the response function is ρ^{BR}, then the stable state is a Nash Equilibrium in pure strategies. In general, a voter response function is an extremely flexible tool. It can be as simple as calculating the best response to the scoring profile formed by votes of others, or it can take into account a degree of uncertainty as is done in the locally dominant response in Meir et al. [10]. It can even be non-myopic in nature, a particular variation of which we present in this paper.

2.1 Farsighted voter model

A *farsighted* voter considers the possibility that their vote might change the situation, prompting other voters to change their vote (e.g., it would cause others to vote similarly, as there are other similarly minded voters), resulting in a better result for them in a few steps. However, our voters are not omnipotent – they are still limited by the common iterative voting constraints. Therefore, the underlying votes are still *opaque* — only the final score outcomes at each stage are known, and not which voter changed their votes, limiting their information (see Section 6 which expands on this point). Furthermore, we use a simple metric to define how optimistic is each voter of the chance others will follow them, by having each voter consider that up to $r_i \in \mathbb{N}$ voters might change their votes according to similar preferences to theirs.

We wish to note that while this definition has some similarities to [10] locality model (with ℓ_1 metric), in that model, the voter is unsure of what the *current* state is, and desires to be pivotal, while in our case they know what the current state is, and are trying to improve the outcome in several steps. This makes it a dominant strategy (as in [10]), in many cases, simply to stay put (Section 6 expands on further potential models).

We deal only with plurality and veto; only they converge with best response strategies, and we do not require a complex discussion of what it means that other voters move according to a voter's desires. There is only an option of whom to vote for (or veto), and not a more intricate division of points. Therefore, we can use the following definition:

Definition 3. *In general, a farsighted better response to a scoring profile* $\mathbf{s}$ *is a ballot* a_i *of radius* r_i*, if after the adoption of* a_i *it is possible that additional* r_i *ballot changes by other voters would make* a_i *the election's winner. In particular, we will instantiate this intuition for the Plurality and Veto voting rules as follows.*

Farsighted Plurality Response. Let $\mathbf{b}$ *be a strategic profile, and* $\mathbf{s} = (\mathbf{s}_{-i}, b_i)$ *be the score profiles that corresponds to it. In particular,* $\mathbf{s}_{-i}$ *denotes here the partial score profile formed from* $\mathbf{s}$ *by removing the* i*'th current ballot. A ballot* a_i *is a better* **farsighted plurality (FP)** *vote of radius* r_i *the following two conditions hold. First,* $a_i \succ_i w = f(\mathbf{s})$*. Second, letting* $\mathbf{s}' = (\mathbf{s}_{-i}, a_i)$ *and* $w = f(\mathbf{s}')$*, if* $\widehat{Q}(w) < \widehat{Q}(a_i)$*, then* $\mathbf{s}'(w) - \mathbf{s}'(a_i) \leq r_i$*, and if* $\widehat{Q}(w) > \widehat{Q}(a_i)$*, then* $\mathbf{s}'(w) - \mathbf{s}'(a_i) + 1 \leq r_i$*.*

A ballot a_i *is the* best FP response *for voter* $i \in N$ *of radius* r_i *if* a_i *is a better FP response for this voter and there is no ballot* b_i *that is also a better FP response and* $b_i \succ_i a_i$*.*

Farsighted Veto Response. Let $\mathbf{s}$ *be the current score profile,* $w = f(\mathbf{s})$ *be the current winner and denote by* $C_i^w \subset M$ *the subset of all candidates that* i *prefers over* w*,* *and* $c_i^* = \arg\max_{c \in C_i^w} \mathbf{s}(c)$, *where tie-breaking is taken into account, that is* $\arg\max$ *is a singleton under the score profile* $\mathbf{s}$.

A ballot a_i *is a better* **farsighted veto (FV)** *response of radius* r_i *if there is* $c \in C_i^w$ *and a set of voters* $V \subset N \setminus \{i\}$, $|V| \leq r_i$ *who can veto candidates in* $M \setminus \{c\}$ *without changing c's score and make it the new winner.*

A ballot a_i *is a* **best FV** *response of radius* r_i *if there is a set of voters* $V \subset N \setminus \{i\}$, $|V| \leq r_i$ *who can veto candidates in* $M \setminus \{c_i^*\}$ *without changing* c_i^* *score and make it the new winner. We also require that the size of the minimal such* V *reduces when* i *adopts* a_i. [6]

We denote the policy of farsighted best response under Plurality by ρ^{FP}, and under Veto by ρ^{FV}. For the latter, we require that if the current ballot is a best FV, then ρ^{FV} prefers it. Notice that this response is coherent with the standard definition of a better and the best responses—in the sense that when the farsighted radius is zero, the farsighted and the standard responses coincide. In addition, it is necessary to notice that a stable state with respect to ρ^F is not necessarily a Nash Equilibrium, nor the other way around. A simple intuition behind this is that, while ρ^{BR} is myopic, ρ^{FP} and ρ^{FV} are not.

Example 2. Assume that the Plurality voting rule is used, and that we have 4 voters and 4 candidates (named a through d), with the tiebreaking order $a \succ b \succ c \succ d$. Let the truthful preference profile be as follows:

$$\begin{array}{rl} \text{voter 1}: & a \succ_1 c \succ_1 b \succ_1 d \\ \text{voter 2}: & c \succ_2 a \succ_2 b \succ_2 d \\ \text{voter 3}: & d \succ_3 a \succ_3 b \succ_3 c \\ \text{voter 4}: & b \succ_3 d \succ_3 c \succ_3 a \end{array}$$

In the myopic version, voter 4 is the voter incentivized to manipulate, and does so by voting for d. Depending on which voter move first between voters 1 and 2, the winner will be a or c.

Now, suppose all voters have a farsighted radii of 2. Now, any candidate has a chance of winning, and therefore no voter will move from it strategy, and therefore the winner will remain a (in this this is a good result, as it ensures the winner will be a Condorcet winner).

3 Farsighted Plurality

In this section we will investigate the convergence of iterative voting processes guided by farsighted response functions under the Plurality voting rule. In particular, we will show that farsighted response leads to convergence to a stable state if all voters use the same farsighted radius.

Theorem 1. *Let us assume that all voters in* N *participate in an iterative voting scenario and use a farsighted plurality response function* ρ^{FP} *and all farsight radia are*

[6] Notice that if the best FV responses exist, they necessarily include the current winner w.

equal, $\forall i \in N,\ r_i = r \in \mathbb{N}$. *Denote* $\mathbf{b}^t$ *the complete voting profile at time* t *and* $\mathbf{s}^t = \mathbf{s}_{\mathbf{b}^t}$ *the score profile that corresponds to it. Then there are a finite* τ, $\mathbf{b}^*$ *and* $\mathbf{s}^*$ *so that for all* $t > \tau$ *it holds that* $\mathbf{b}^t = \mathbf{b}^*$ *and* $\mathbf{s}^* = \mathbf{s}^t = \mathbf{s}_{\mathbf{b}^*}$.

Proof. Denote $w^t = f(\mathbf{s}^t)$ the winner at iteration t. Also denote by $A_i^t = A_i(\mathbf{s}^t) \subset M$ the set of all possible farsighted responses of voter i at time $t > 0$. At time $t = 0$ let A_i^t also include w_0 for all $i \in N$. Let $U_i^t \subset M$ be defined as $U_i^t = A_i^t \cup \{b_i^t\}$, i.e., the current ballot and the set of all possible better farsighted ballots.

We will show that two conditions hold simultaneously and lead to convergence as required: a) $U_i^{t+1} \subseteq U_i^t$; and b) $\mathbf{s}^t(w^t) \leq \mathbf{s}^{t+1}(w^{t+1})$ and, furthermore, if $\mathbf{s}^t(w^t) = \mathbf{s}^{t+1}(w^{t+1})$ then either $w^t = w^{t+1}$ or $\widehat{Q}(w^{t+1}) < \widehat{Q}(w^t)$. In other words, for all voters their corresponding sets of better farsighted responses do not grow, and the score of the winner does not decrease neither w.r.t. gathered points nor w.r.t. the tiebreaking order.

Assume the contrary, and let us consider the first iteration t where either condition is violated for the first time.

Assume that $\mathbf{s}^t(w^t) > \mathbf{s}^{t+1}(w^{t+1})$. That is, at step t, some voter $i \in N$ has changed his ballot from $b_i^t = w^t$ to $b_i^{t+1} = c \neq w^t$. Notice that c is not necessarily the winner w^{t+1} of the step $t+1$. It has to hold, however, that $c \succ_i w_t$. Since t is the first instance when either of our conditions fail, it has to holds that $U_i^{t'} \setminus U_i^{t'-1} = \emptyset$ for all $\theta' \leq t$. In particular, for any $\tau \leq t$ $c \in U_i^t$. Consider now, the time $\tau < t$, when voter i has first switched his ballot to w_t. Since $c \in U_i^\tau$ one of the following cases occurs:

- $c \in A_i^\tau$. To switch to w^t, by the definition of the best farsighted plurality response, it has to holds that $w^t \in A_i^\tau$ and is the best w.r.t. Q_i. However, we have already established that $c \succ_i w^t$. Contradiction.
- $c = b_i^\tau$. Again an impossibility, because it has to holds that $w^t = b_i^{\tau+1} \succ_i b_i^\tau = c$.

Notice that the reasoning above applies to the case where the winner score persists, but its tiebreaking order is violated.

Assume that for some $i \in N$ $U_i^{t+1} \setminus U_i^t \neq \emptyset$ *and* $\mathbf{s}^t(w^t) \leq \mathbf{s}^{t+1}(w^{t+1})$. Consider two complimentary sub-cases, $b_i^{t+1} \neq w^{t+1}$ and $b_i^{t+1} = w^{t+1}$.

- Assume that $b_i^{t+1} \neq w^{t+1}$, that is voter i does not vote for the winner. It is easy to see that $U_i^{t+1} \setminus U_i^t \neq \emptyset$ if and only if $A_i^{t+1} \setminus A_i^t \neq \emptyset$. Let $c \in A_i^{t+1} \setminus A_i^t \neq \emptyset$. Since $\mathbf{s}^t(w^t) \leq \mathbf{s}^{t+1}(w^{t+1})$, this can only occur if c has received an additional points at time t so that $\mathbf{s}^t(c) < \mathbf{s}^{t+1}(c)$. Therefore, there is a voter $j \in N$ so that $c \in A_j^t$. However, since farsight radia are equal, it would also entail that $c \in A_i^t$ – contradiction.
- Assume that $b_i^{t+1} = w^{t+1}$. If $w_t = w_{t+1}$ then we can use the same reasoning as the above sub-case. Thus, w.l.o.g., assume that $w_t \neq w_{t+1}$. Given that for all $j \in N$ $U_j^t \setminus U_j^{t-1} = \emptyset$ we can easily see that w_t keeps his points at step $t+1$. As a result, for any $c \notin A_i^t$ is has to hold that $c \notin A_i^{t+1}$ – a contradiction.

Notice that sequences $\{U_i^t\}$ and $\{\mathbf{s}^t(w^t)\}$ are bounded by the empty set and the number of voters, respectively. Hence, there is a point τ' after which neither the set nor the score (including a shift along the tiebreaking order) do not change. In particular, exist U_i so that for all $t > \tau'$, $U_i^t = U_i$. Therefore, there is $\tau > \tau'$, where all voters have updated their ballots to the best farsight response in U_i or their response can not change because U_i only includes their current vote. As a result $\mathbf{b}^t = \mathbf{b}^\tau$ and $\mathbf{s}^t = \mathbf{s}^\tau$ for all $t > \tau$.

Our proof relies on the fact that farsight radia are equal, however this is not just a property of this specific proof. Rather, homogeneity of farsight radia is a necessary condition for convergence to a stable state.

Theorem 2. *There is an iterative voting scenario with farsighted voters with heterogeneous farsight radia where the iteration does not converge.*

Proof. We will prove this by constructing a scenario where a cycle of farsighted response exists. Let there be 5 candidates, named named for convenience a through d, rather than just numbered, and let the preference profile of the first three voters be:

$$\text{voter 1} : a \succ_1 c \succ_1 b \succ_1 d \succ_1 e$$
$$\text{voter 2} : b \succ_2 c \succ_2 a \succ_2 d \succ_2 e$$
$$\text{voter 3} : c \succ_3 d \succ_3 e \succ_3 a \succ_3 b$$

All remaining voters of the profile would not participate in the cycle, but their profiles are chosen so that that the initial score profile, $\mathbf{s}^0$ is given by: $\mathbf{s}^0(a) = \mathbf{s}^0(b) = 4$, $\mathbf{s}^0(c) = 6$, $\mathbf{s}^0(d) = \mathbf{s}^0(e) = 10$. In addition, assume that $r_1 = r_2 = 6$ and $r_3 = 2$, while the tie-breaking preference order, $\widehat{Q}$, is $a \succ b \succ c \succ e \succ d$. Then the following cycle of farsighted votes exists starting from $\mathbf{s}^t = \mathbf{s}^0$:

- Voter 3 is changing his vote from c to d, leading to the scoring profile $\mathbf{s}^{t+1} = (4, 4, 5, 11, 10)$.
- Voters 1 and 2 change their votes in favour of c one after the other. At both of these changes $w^{t+1} = w^{t+2} = d$ and has 11 points, so neither a is among the frigged responses of voters 1, nor b is among the farsighted responses of voter 2. Hence, the score profiles becomes $\mathbf{s}^{t+3} = (3, 3, 7, 11, 10)$.
- c now becomes a farsighted better response for voter 3, since only two more votes besides his own would be necessary to make c the winner of the election. In fact, this is the best farsighted response for voter-3, which he makes, turning the score profile into $\mathbf{s}^{t+4} = (3, 3, 8, 10, 10)$.
- Candidates a and b can now win by tiebreaking if they gain 6 more votes in addition to those granted by voters 1 and 2 reverting to their original ballots. In other words a and b are not best farsight responses of voters 1 and 2 respectively. As a result of adopting these ballot modifications the score again becomes $\mathbf{s}^{t+6} = (4, 4, 6, 10, 10) = \mathbf{s}^t$. The cycle is complete.

4 Veto

Due to the nature of veto, a farsighted best response strategy is, in fact, similar to the myopic version, as vetoing the current winner is always the right strategy. This does not mean that convergence stems from previous results, as a stable state in the myopic scenario might not be a stable state in the farsighted version. However, our proof is similar in many respects to the myopic version found in [9], and we try to use similar notations, where applicable.

Surprisingly, our veto result is stronger than our plurality one, as we show iterative farsighted veto converges even for voters with different radii.

Theorem 3. *Let us assume that all voters in N participate in an iterative voting scenario and use a farsighted plurality response function ρ^{FP}. Denote $\mathbf{b}^t$ the complete voting profile at time t and $\mathbf{s}^t = \mathbf{s}_{\mathbf{b}^t}$ the score profile that corresponds to it. Then there are a finite τ, $\mathbf{b}^*$ and $\mathbf{s}^*$ so that for all $t > \tau$ it holds that $\mathbf{b}^t = \mathbf{b}^*$ and $\mathbf{s}^* = \mathbf{s}^t = \mathbf{s}_{\mathbf{b}^*}$.*

Proof. We begin by assuming that the theorem is false, and there is a profile $\mathbf{b}$ that does not converge to a stable state. Therefore, we know there is a cycle $\mathbf{b}^t, \mathbf{b}^{t+1}, \ldots, \mathbf{b}^{t+k}$ for some $t, k \in \mathbb{N}$ that repeats ad infinitum. We shall focus on these states and mark them as $G_0, \ldots, G_k$. We shall use the notation $max(G_i)$ to indicate the maximal score in a particular profile (i.e., $max(G_i) = \max_{c \in N} \mathbf{s}_{G_i}(c)$). Notice that the choice of which state is G_0 is arbitrary, and the numbering can begin at every point in the cycle.

Lemma 1. *For every G_i, if $j < i$, $max(G_I) \leq max(G_j) + 1$, and if the inequality is tight, there is only one candidate with the score $max(G_i)$.*

Proof. Proving by induction, the base case is trivial. Assuming it is true after $h - 1$ steps, proving it for step h. Examining G_{h-1}, if there was a $j < h - 1$ for which $max(G_{h-1}) = max(G_j) + 1$, there is a single winner, which looses a point, and therefore the winner in G_h will have, at most, $max(G_{h-1})$ points. Thus, $max(G_h) \leq max(G_{h-1})$, and the claim stems from its truth for G_{h-1}. If for every $j < h - 1$ $max(G_{h-1}) \leq max(G_j)$, the maximal score in G_h will rise by at most one point, i.e., $max(G_h) \leq max(G_j) + 1$ for all $j < h$. Furthermore, if it indeed grows, there is only a single candidate with that number of points (as only the candidate that got an extra point has this score). If the maximal score in G_h didn't grow, $max(G_h) \leq max(G_{h-1})$, and claim is true from its correctness for G_{h-1}.

Corollary 1 *For every $1 \leq i \leq k$, $max(G_0) + 1 \geq max(G_i) \geq max(G_0) - 1$.*

Proof. Since G_0 can be any state in the cycle, Lemma 1 can be at any order – for any i, j, G_i can be before or after G_j.

Lemma 2. *There are, at most, two different values for $max(G_i)$.*

Proof. Suppose there are 3. Due to Lemma 1, there is a single winner in the two top values. However, once there is a single candidate in the lower of these, there is no way for this candidate to receive another point, as the sole step is to remove a point from the winner.

If for every G_i there is only a single winner with the maximal winner, this means the candidate granted a point in the move is the one becoming the winner. Hence, each voter move is making a candidate it previously vetoed the winner. I.e, its situation is slowly deteriorating, as candidates it ranks low become winners — this is a finite process, with at most $n \cdot (m - 1)$ steps, contradicting the cycle existence. Therefore, we must assume there is at least one state G_i in which there is more than one candidate with the maximal score. We shall term one of these states G_0 (obviously, any state with a different maximal score than $max(G_0)$ has a maximal score of $max(G_0)+1$ and a single candidate with that score, thanks to Lemma 1).

Lemma 3. *For any state G_i where $max(G_i) = max(G_0)$, $\{c \in M | \mathbf{s}_{G_i}(c) \geq max(G_i) - 1\} = \{c \in M | \mathbf{s}_{G_0}(c) \geq max(G_0) - 1\}$ and $|\{c \in M | \mathbf{s}_{G_i}(c) = max(G_i) - 1\}| = |\{c \in M | \mathbf{s}_{G_0}(c) = max(G_0) - 1\}|$, $|\{c \in M | \mathbf{s}_{G_i}(c) = max(G_i)\}| = |\{c \in M | \mathbf{s}_{G_0}(c) = max(G_0)\}|$ (i.e., the set of candidates with score $max(G_0)$ or $max(G_0) - 1$ is the same, as well as number of candidates with each of these scores).*

Proof. Suppose there is a G_i which has a candidate c such that $\mathbf{s}_{G_i}(c) \geq max(G_i) - 1$ but $\mathbf{s}_{G_0}(c) < max(G_0) - 1$. This means at some point c had a score of $max(G_0) - 1$ and lost a point. But a candidate can only lose a point when it is the winner, and there is no state where the maximal score is $max(G_0) - 1$ (if G_0 has an "extra" candidate in comparison to G_i the same argument applies, switching between G_i and G_0).

This means the only candidates changing scores are those in $\{c \in M | \mathbf{s}_{G_0}(c) \geq max(G_0) - 1\}$, and as this is a zero-sum game and the maximal score does not change, the number of candidates with the same score does not change.

Examining the set B of candidates for which there is a G_i where they have a score of $max(G_0)$ and there is a G_j where they have a score of $mac(G_0) - 1$ (this is a non-empty set, as some candidate is vetoed between G_0 and G_1), we mark as z the candidate ranked lowest in $\widehat{Q}$ (i.e., $\widehat{Q}(z) \geq \widehat{Q}(z')$ for all $z' \in B$).

Since z changes its score, there is a state G_i where z has the score $max(G_0)$ and is vetoed, i.e., z is the winner if G_i. This means there is no other candidate from B with the score $max(G_0)$. As the number of candidates with $max(G_0)$ doesn't change (according to Lemma 3), this means that at every state G_j in which $max(G_j) = max(G_0)$, there is only a single candidate from B with $max(G_0)$ points, and it always wins. This means the candidate getting the point at every stage is the one that becomes the winner — which, as noted above, is a finite process, contradicting the endless cycle.

5 Single Farsighted Voter

Previous sections have shown that heterogeneity of farsighted radii has no effect under the Veto voting rule, but is detrimental to Plurality. In this section we investigate this discrepancy more closely by looking at a special case. Specifically, both under the Plurality and Veto voting rules, we will analyse the situation where only one voter has a non-zero farsighted radius.

Theorem 4. *The iterative process with one farsighted voter under Plurality converges.*

Proof (Proof Sketch). Here we only provide an outline of the proof, and leave the simpler details completion as an exerice for the reader. For the single farsighted voter the set U_i^t will behave similarly to that from the proof of Theorem 1. On the other hand, between the iterations where the farsighted voter changes his ballot, all other voters will behave myopically and their combined ballot will tend to stabilise (see [11]). This overall will lead to the convergence of the complete voter profile.

Another important observation is that, even if the farsighted voting dynamic converges in spite of the farsighted radii heterogeneity, it may be detrimental, i.e., lead to a sub-optimal stable state from the farsighted voter's perspective. In particular, even though for the Veto voting rule radia heterogeneity does not break convergence, its negative effect is more universal.

Example 3. Assume that the Plurality voting rule is used, and that we have 4 voters and 4 candidates (once again conveniently named a through d). Let the truthful preference profile be as follows:

$$\begin{aligned} \text{voter 1} &: a \succ_1 c \succ_1 b \succ_1 d \\ \text{voter 2} &: c \succ_2 a \succ_2 b \succ_2 d \\ \text{voter 3} &: d \succ_3 a \succ_3 b \succ_3 c \\ \text{voter 4} &: b \succ_3 d \succ_3 c \succ_3 a \end{aligned}$$

Let voter 4 be the single farsighted voter with farsighted radius $r_4 = 2$, and the tiebreaking order $\widehat{Q} = a \succ b \succ c \succ d$. It is easy to see that the truthful state is stable under the farsighted response dynamics, and the winner is a. On the other hand, if we consider the standard myopic dynamic ρ^{BR}, it will lead to the equilibrium voting profile $\mathbf{b} = (c, d, d, d)$. This equilibrium state is better from the perspective of voter-4, since $c \succ_4 a$.

Furthermore, assume that we introduce a myopic bias into ρ^{FP}, i.e. when voter has both myopic and farsighted move he always prefers the myopic move. Even in this case, the example would hold. After 2 myopic steps, when voter 4 has switched to candidate d and voter-1 has switched to candidate c, it will be Nash equilibrium in the ρ^{BR} based, standard game. In particular there will be no more myopic moves available. However, voter-4 will still have a farsighted move: to return to his truthful vote. As a result of voter-4 reverting to his truthful ballot, voter-1 will now has a new myopic move to return to his truthful vote. The initial (truthful) state will be restored. That is a cycle is formed, and a stable state under the myopic biased ρ^{FP} does not exist.[7]

Example 4. Assume that the Veto voting rule is used, and that we have 3 voters and 4 candidates (once again conveniently named a through d). Let the truthful preference profile be as follows:

$$\begin{aligned} \text{voter 1} &: c \succ_1 a \succ_1 b \succ_1 d \\ \text{voter 2} &: a \succ_2 b \succ_2 c \succ_2 d \\ \text{voter 3} &: a \succ_3 b \succ_3 c \succ_3 d \end{aligned}$$

Let voter 1 be the single farsighted voter with farsighted radius $r_1 = 1$, and the tiebreaking order is $\widehat{Q} = (a \succ b \succ c \succ d)$. The truth ballot, (d, d, d), is a Nash Equilibrium under the best response dynamics ρ^{BR} with the winner being a. However, voter-1 has a farsighted move: to switch from vetoing c to vetoing a. This is because it is enough that one other voter vetoes b and the best candidate of voter-1 (candidate c) would become the new winner. Under ρ^{FV}, voter-1 adopts the farsighted ballot leading to a voting profile of $\mathbf{b} = (c, d, d)$, and the new winner becomes b. However, now voter-2

[7] Notice, that this situation is different from the one in the theorem 4, because myopic moves are performed *before* any farsighted moves.

and voter-3 have no incentive to change their ballot and voter-1 has no farsighted ballots, i.e. $\mathbf{b} = (c, d, d)$ is a stable stable voting profile. Alas, $a \succ_1 b$, so that voter-1 is worse-off.

The above examples also suggest another observation, although we omit its formal analysis due to space limitations. Under the Plurality rule, increasing the farsighted radius beyond $r = 2$ does not effect the overall iterative behaviour of a system with a single farsighted voter. On the other hand, under the Veto rule, the increase of the farsighted radius will have behavioural impact.

6 Discussion

In this paper we tackled the last, and most complex, condition for iterative voting convergence — myopicness. After previous work has dealt with tie-breaking rules, voting rules, and best response, we have examined the conditions for convergence when there is no longer the assumption that players will only look at the current state. We have explored farsightedness while keeping the other elements of the model unchanged (apart from the voting rule), in order to understand its effect, on its own, on the known properties of iterative voting.

Beyond examining convergence in both plurality and veto (with the surprising fact that the veto results are stronger than the plurality ones), we also begin the study of the effect of farsightedness on the result of the iterative process, by finding that farsightedness is not necessarily beneficial. The effects of various numbers of farsighted voters, and the effect of the radius on the result, are interesting areas for further research.

Another area open to further research is slightly relaxing the opacity of information to the various players. While limited knowledge, in which voters only know total scores of each state, makes sense in many scenarios, there are scenarios where this can be changed. For example, in setting with few voters, the iterative process can be more exposed, and while participants will not necessarily share their truthful preferences, voters can learn from the changes to votes made by other players something regarding their preferences. Such knowledge may enable them to make more subtle farsighted moves.

Another potential extension to this new farsighted research is to endow the players with utilities for each candidate. This allows them a greater degree of nuance in their farsightedness, and they can decide between supporting candidates that require fewer other voters to become victors, vs. more preferable candidates (to them) that require a longer horizon to win. Similar settings have been studied in general games, which may aid such research.

7 Acknowledgements

This work was supported, in part, by the COST Action IC1205 on Computational Social Choice. It has been co-financed by the UK Research Council for project ORCHID—grant EP/I011587/1; by RFFI 14-01-00156-a; by Israel Science Foundation grant #1227/12, Israel Ministry of Science and Technology grant #3-6797, and by Microsoft Research through its PhD Scholarship Programme. Svetlana Obratsova also gratefully acknowledges funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)—ERC grant agreement number 337122.

References

1. Airiau, S., Endriss, U.: Iterated majority voting. In: Rossi, F., Tsoukias, A. (eds.) Algorithmic Decision Theory, Lecture Notes in Computer Science, vol. 5783, pp. 38–49. Springer Berlin / Heidelberg (2009)
2. Bartholdi III, J.J., Tovey, C.A., Trick, M.A.: The computational difficulty of manipulating an election. Social Choice and Welfare 6(3), 227–241 (1989)
3. Brânzei, S., Caragiannis, I., Morgenstern, J., Procaccia, A.D.: How bad is selfish voting? In: Proceedings of the 27th National Conference on Artificial Intelligence (AAAI). pp. 138–144. Bellevue, Washington (July 2013)
4. Conitzer, V., Walsh, T., Xia, L.: Dominating manipulations in voting wih partial information. In: Proceeings of the AAAI Conference (2011)
5. Desmedt, Y., Elkind, E.: Equilibria of plurality voting with abstentions. In: Proceedings of the 11th ACM conference on Electronic commerce (EC). pp. 347–356. Cambridge, Massachusetts (June 2010)
6. Dutta, B., Laslier, J.F.: Costless honesty in voting (2010), in 10th International Meeting of the Society for Social Choice and Welfare, Moscow
7. Gibbard, A.: Manipulation of voting schemes. Econometrica 41(4), 587–602 (July 1973)
8. Grandi, U., Loreggia, A., Rossi, F., Venable, K.B., Walsh, T.: Restricted manipulation in iterative voting: Condorcet efficiency and borda score. In: Proceedings of 3rd International Conference of Algorithmic Decision Theory (ADT). pp. 181–192. Brussels, Belguim (November 2013)
9. Lev, O., Rosenschein, J.S.: Convergence of iterative voting. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS). vol. 2, pp. 611–618. Valencia, Spain (June 2012)
10. Meir, R., Lev, O., Rosenschein, J.S.: A local-dominance theory of voting equilibria. In: Proceedings of the 15th ACM conference on Economics and computation (EC). pp. 313–330. Palo Alto, Califronia (June 2014)
11. Meir, R., Polukarov, M., Rosenschein, J.S., Jennings, N.R.: Convergence to equilibria of plurality voting. In: Proceedings of the 24th National Conference on Artificial Intelligence (AAAI). pp. 823–828. Atlanta (July 2010)
12. Obraztsova, S., Markakis, E., Polukarov, M., Rabinovich, Z., Jennings, N.R.: On the convergence of iterative voting: How restrictive should restricted dynamics be? In: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI). Austin, Texas (January 2015)
13. Obraztsova, S., Markakis, E., Thompson, D.R.M.: Plurality voting with truth-biased agents. In: Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT). pp. 26–37. Aachen, Germany (October 2013)
14. Rabinovich, Z., Obraztsova, S., Lev, O., Markakis, E., Rosenschein, J.S.: Analysis of equilibria in iterative voting schemes. In: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI). Austin, Texas (January 2015)
15. Reijngoud, A., Endriss, U.: Voter response to iterated poll information. In: AAMAS (2012)
16. Reyhani, R., Wilson, M.C.: Best-reply dynamics for scoring rules. In: Proceedings of the 20th European Conference on Artificial Intelligence (ECAI). vol. 2, pp. 672–677. Montpellier, France (June 2012)

17. Satterthwaite, M.A.: 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 (April 1975)
18. Thompson, D.R.M., Lev, O., Leyton-Brown, K., Rosenschein, J.S.: Empirical aspects of plurality election equilibria. In: Proceedings of the 12th International Coference on Autonomous Agents and Multiagent Systems (AAMAS). pp. 391–398. St. Paul, Minnesota (May 2013)
19. Xia, L., Zuckerman, M., Procaccia, A.D., Conitzer, V., Rosenschein, J.S.: Complexity of unweighted coalitional manipulation under some common voting rules. In: Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI). pp. 348–353. Pasadena, California (July 2009), ftp://ftp.cs.huji.ac.il/users/jeff/ijcai09xia.pdf