Non-Myopic Voting Dynamics: An Optimistic Approach

Svetlana Obraztsova
Hebrew University of Jerusalem
svetlana.obraztsova@gmail.com

Omer Lev
University of Toronto
omerl@cs.toronto.edu

Maria Polukarov
University of Southampton
mp3@ecs.soton.ac.uk

Zinovi Rabinovich
Mobileye Vision Technologies Ltd.
zr@zinovi.net

Jeffrey S. Rosenschein
Hebrew University of Jerusalem
jeff@cs.huji.ac.il

Abstract

Iterative voting has presented, in the past few years, a voting model in which a player participates in an election poll, and can change his vote at any time to influence the result. Several extensions for this model have been considered, including some attempts to handle the uncertainty that players may face. However, all those extensions retained the myopic assumption—that is, players change their vote only when they believe that their move will have an immediate effect on the outcome.

In this paper, we address this assumption by allowing for certain non-myopic dynamics. Specifically, the outlook is optimistic to a certain extent, a *horizon*, as players change their vote if they believe that if some other players also move, the outcome can change. We show that players with the same horizon of optimism would converge to a Nash equilibrium under Plurality, and for Veto, even players with varying horizons of optimism always converge. However, such non-myopic behavior is not necessarily a positive feature—as we demonstrate, in some cases it is better for the player to stick to myopic moves.

1 Introduction

The study of processes by which multiple agents with varying individual interests can reach a collective decision, has long been an active area of research in AI. Voting mechanisms are a natural and popular tool that allows a group of agents to make collective choices, despite the differences in their individual preferences over available alternatives. However, one of the main stumbling blocks in this approach has been that voters may choose to hide their truthful preferences from the voting mechanism and seek to influence the outcome in their favor by strategically casting their ballots. Sadly, the famous Gibbard-Satterthwaite theorem [Gibbard, 1973; Satterthwaite, 1975] shows that every voting system is susceptible to this behavior.

Several methods have been proposed to deal with this issue. One of the common approaches has been to explore the complexity of voting mechanisms (or rules) and their manipulation [Bartholdi III *et al.*, 1989; Xia *et al.*, 2009]. However, many common rules, like Plurality and Veto, are easily manipulable, yet are widely used.

Therefore, instead of seeking to avoid manipulations, an alternative approach has been to try and understand the ultimate outcome of voters' strategic behavior in practical scenarios. [Meir *et al.*, 2010] suggested examining the model of iterative voting, in which the voters update their ballots one after another (in an arbitrary order), if this can change the result to their benefit. Crucially, the voters are myopic—that is, only look for immediate effects of their moves. This work has inspired a stream of research that further developed and extended the basic iterative voting model. However, the myopic assumption has been left untouched so far—due, in large part, to the difficulty of effectively modeling non-myopic behavior.

In this paper, we address this assumption by considering a bounded, non-myopic behavior where the voters may change their ballots even if the effect of their move may not follow immediately, but only in a few steps, after some other similar-minded voters have made their appropriate moves. As in the real world, our model of optimistic outlook does not assume that the voters are omnipotent: they can see just a few steps ahead, and they do not know the inner preferences of other players. In a sense, this can be viewed as a sort of bounded rationality, where the players are constrained by what we term the "horizon of optimism" (that might be different for each voter), limiting their tendency to look ahead.

We study the convergence of such non-myopic dynamics to a stable state. We particularly focus on Plurality and Veto, and assume the players use a (non-myopic optimistic) best response function. Our choice is due to previous results showing that no other scoring rule converges with best responses even in the myopic case [Lev and Rosenschein, 2012], and that the scope for other convergent strategies is quite limited [Obraztsova *et al.*, 2015]. Furthermore, we explore whether being a non-myopic optimist is a clear advantage, and find, quite surprisingly, that the answer is negative—such manipulations might result in a less desirable final outcome for the manipulator.

1.1 Related Work

There is an abundant literature on the analysis of voting mechanisms, particularly for Plurality, an overview of which can be found in [Meir *et al.*, 2014]. Here, we shall focus on the iterative voting model presented in [Meir *et al.*, 2010] and its various extensions. This model is the most relevant to our work, although alternative iterative models, such as [Airiau and Endriss, 2009], have also been previously considered.

The main result in [Meir *et al.*, 2010] shows that myopic, best response dynamics always converge to a Nash equilibrium under Plurality and linear tie-breaking. The necessity of linearity of the tie-breaking was shown in [Lev and Rosenschein, 2012], which also extended the convergence results to Veto, while showing that no convergence can be achieved for other scoring rules. Independently, the same results (with alternative proofs) were demonstrated in [Reyhani and Wilson, 2012]. Based on these negative findings, [Grandi *et al.*, 2013; Reijngoud and Endriss, 2012] explored the possibility of designing restricted response functions that would guarantee convergence to a stable state. More recently, [Obraztsova *et al.*, 2015] provided a general characterization of convergent dynamics, which, in particular, enabled the extension of positive results beyond Plurality and Veto.

Other research added different elements to the model. Thus, [Rabinovich *et al.*, 2015] considered iterative processes with truth-biased (as modeled in [Dutta and Laslier, 2010; Thompson *et al.*, 2013; Obraztsova *et al.*, 2013]) and lazy-biased (as modeled in [Desmedt and Elkind, 2010]) voters, and demonstrated quite different convergence results. The quality of the final outcome of iterative voting was studied in [Brânzei *et al.*, 2013], as measured by what they termed the "dynamic price of anarchy". Finally, [Meir *et al.*, 2014] (and later, some parts expanded in [Meir, 2015]) incorporated into the setting uncertainty about the current state, by giving voters only an estimation of the score for each candidate and having the voters assume that scores may vary within some radius (which is an individual voter's parameter). Although this bound is different by nature from our optimism horizon, some of the techniques used in [Meir *et al.*, 2014] provided insights for our work.

[Reyhani *et al.*, 2012], in a slightly different model of uncertainty, assume voters vote optimistically, but they handle a very limited scenario and mostly three candidates.

2 Preliminaries

In this section, we present the notation for iterative voting (partly adopted from [Meir *et al.*, 2014]), and then define the model with non-myopic players.

2.1 Notation

We denote a discrete set of x elements by $[x] = \{1, ..., x\}$. In particular, the set of candidates is $M = [m]$ and the set of voters is $N = [n]$.

Let $\pi(M)$ be the set of all strict linear orders over M. A *preference profile* $\mathbf{Q} \in (\pi(M))^n$ is the vector of all voter preferences over the candidates, so that $Q_i \in \pi(M)$ is the (true) preference order for voter $i \in N$. In particular, $Q_i(c) \in [m]$ is the rank of candidate $c \in M$ in Q_i, and we say that voter $i \in N$ prefers candidate $c \in M$ to candidate $c' \in M$, denoted by $c \succ_i c'$, if $Q_i(c) < Q_i(c')$. Thus, in particular, $q_i = Q_i^{-1}(1) \in M$ is the (true) top choice of voter $i \in N$. Also, a special preference order $\widehat{Q} \in \pi(M)$ will be used to define a lexicographic tie-breaking when necessary.

For a subset $V \subseteq N$, a *partial strategic* (or *voting*) *profile*, $\mathbf{b} \in \pi(M)^{|V|}$, is a vector of ballots submitted by voters in V. A voting profile is *complete* if $V = N$. If $b_i = Q_i$ we say that voter $i \in V$ is *truthful* in $\mathbf{b}$.

Generally, a *voting rule* is a function $f : (\pi(M))^n \to M$, that determines a winner of an election based on a voting profile. In this paper, we will be dealing solely with a class of *scoring* voting rules, characterized by a vector of parameters $(\alpha_1, \dots, \alpha_{m-1}, 0)$ where $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_{m-1} \geq 0$. These parameters essentially give a numerical 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 receives α_1 points, the second highest α_2 points, and so on. By aggregating these points, a winner of an election is defined as one with the highest score.

Formally, we define a *score profile* or *state* as a partial statistic, $\mathbf{s_b}$, of a strategic profile $\mathbf{b}$, that 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 $\mathbf{s_Q}(c) = \sum_{i \in N} \alpha_{Q_i(c)}$.

We can view $\mathbf{s_b}$ as a vector in $\mathbb{N}^m$. Given such a score vector $\mathbf{s}$, the winner is the candidate with the maximal number of points in $\mathbf{s}$ (where some tie-breaking order $\widehat{Q}$ is used in case there are several candidates with the maximal score). In a slight abuse of notation, if voters vote $\mathbf{b}$ and $\mathbf{s}$ is the resulting score vector, we will use $f(\mathbf{s})$ to denote the winner, and $\mathbf{s}(c)$ for $c \in M$ to denote the score of candidate c in the scoring profile $\mathbf{s}$. Thus, for any $c \neq f(\mathbf{s})$ we have $\mathbf{s}(c) \leq \mathbf{s}(f(\mathbf{s}))$ and if $\mathbf{s}(c) = \mathbf{s}(f(\mathbf{s}))$ then $\widehat{Q}(f(\mathbf{s})) < \widehat{Q}(c)$.

We will concentrate on two of the most popular scoring rules: Plurality, defined by the vector $(1, 0, \dots, 0)$, and Veto, defined by the vector $(1, \dots, 1, 0)$. 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 "summarize" 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 $\mathbf{b} \in M^{|V|}$ for a (partial) voting profile.

In the iterative voting model, each voter $i \in N$ has access to a partial score profile $\mathbf{s}_{\mathbf{b}_{-i}}$, which is created from a voting profile $\mathbf{b}$ ($\mathbf{b}_{-i}$ are all votes except that of voter i). Each voter strategically decides whether to change its current vote from b_i. Note that each voter i can calculate what would be the score profile if it changed its vote to something else, as it knows what $s_{\mathbf{b}_{-i}}$ is. We say that $c_i \in M$ is a *better response* of voter i to a score profile $\mathbf{s_b}$ if changing i's vote from b_i to c_i creates a scoring profile $\mathbf{s}_{(\mathbf{b}_{-i}, c_i)}$, and $f(\mathbf{s}_{(\mathbf{b}_{-i}, c_i)}) \succ_i f(\mathbf{s_b})$. A better response that results in the most preferable possible outcome for voter i is called a *best response*. Note that under Veto, vetoing the current winner is always a best response: i.e., vetoing c_i is a better (and best) response to $\mathbf{s_b}$ if $c_i = f(\mathbf{s_b})$, and the new winner is preferred by voter i over c_i. Under Plurality, there may be several votes that result in the same (best) outcome for voter i; however, if i can change the outcome for the better, they can also do so by voting directly for the new winner, so the set of best responses is further restricted to a single strategy c_i such that $c_i = f(\mathbf{s}_{(\mathbf{b}_{-i}, c_i)})$

and for any scoring profile $\mathbf{s}_{(\mathbf{b}_{-i},c'_i)}$ created by replacing voter i's vote with c'_i, $c_i \succ_i f(\mathbf{s}_{(\mathbf{b}_{-i},c'_i)})$ for all $c'_i \in M$.

It is convenient to generalize the above definitions as a *voter response function* $\rho : N \times \mathbb{N}^m \to M$ that maps a score profile $\mathbf{s}$ into a ballot of each voter i. We naturally shorthand $\rho_i(\mathbf{s}) = \rho(i, \mathbf{s})$. For instance, for Plurality, the best response function, ρ^{BR}, is such that $\rho_i^{BR}(\mathbf{s}) = f\left(\mathbf{s}_{\mathbf{b}_{-i}}, \rho_i^{BR}(\mathbf{s})\right)$ and $\rho_i^{BR}(\mathbf{s}_\mathbf{b}) \succ_i f(\mathbf{s}_{\mathbf{b}_{-i}}, c)$ for all $c \in M$, $c \neq \rho_i^{BR}(\mathbf{s}_\mathbf{b})$.

We also can define stability of voting profiles and their corresponding states in terms of a voter response function. Let $\mathbf{b}$ be a partial voting profile for $V \subseteq N$, and $\mathbf{s}$ its corresponding score profile. Then, $\mathbf{b}$ is a *stable voting profile* and $\mathbf{s}$ is a *stable state* (w.r.t. the voter response function ρ) if for all $i \in V$ we have $b_i = \rho_i(\mathbf{s})$.

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 score 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.*, 2014]. Or, it can be non-myopic in nature, as we present next.

2.2 Non-Myopic (Optimistic) Voting Model

An *optimistic non-myopic* voter considers the possibility that their vote might prompt other players to make a similar move (as they believe there are other similarly minded voters), thus resulting in a better outcome in a few steps. However, such a voter is not omnipotent, and is still limited by the common iterative voting constraints. Thus, the underlying votes are still *opaque*—the players only get to observe the score profile at each stage, but do not know which voter has changed their ballot. Furthermore, we use a simple metric to define how optimistic each player is about the chance that others will follow it: each voter i assumes that up to $r_i \in \mathbb{N}$ voters might make a move in support of their vote.

In this paper, we deal only with Plurality and Veto voting rules, as only these are known to converge even with myopic best response strategies (which can be considered as a particular case of non-myopic ones), 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 whom to veto), and not a more intricate division of points. Thus, we make the following definitions of a *Non-Myopic Plurality (NMP) Response* and a *Non-Myopic Veto (NMV) Response*.

Definition 1 (NM-Plurality Response). *Let* $\mathbf{b}$ *be a strategic profile with its corresponding score profile* $\mathbf{s}$*, and* $w = f(\mathbf{s})$ *denote the winner under* $\mathbf{s}$*. For a given voter* $i \in N$*, a ballot* $c_i \in M$ *is a* better non-myopic plurality response *of optimism horizon* r_i *if the following two conditions hold:*

- $c_i \succ_i w$;
- *For* $\mathbf{s}' = \mathbf{s}_{(\mathbf{b}_{-i},c_i)}$ *and* $w' = f(\mathbf{s}')$*, if* $\widehat{Q}(w') > \widehat{Q}(c_i)$*, then* $\mathbf{s}'(w') - \mathbf{s}'(c_i) \leq r_i$*, and if* $\widehat{Q}(w') < \widehat{Q}(c_i)$*, then* $\mathbf{s}'(w') - \mathbf{s}'(c_i) \leq r_i - 1$.

A ballot c_i *is a* best NMP response *for voter* $i \in N$ *of optimism horizon* r_i *if* c_i *is a better NMP response for this voter and there is no ballot* $c'_i \in M$ *that is also a better NMP response and* $c'_i \succ_i c_i$.

Definition 2 (NM-Veto Response). *Let* $\mathbf{s}$ *be a score profile with winner* $w = f(\mathbf{s})$*, and denote by* $C_i^w \subseteq M$ *the subset of all candidates that voter* i *prefers over* w.

A ballot c_i *is a* better non-myopic Veto response *of optimism horizon* r_i *if there is* $c \in C_i^w$ *and a set of voters* $V \subseteq N \setminus \{i\}$, $|V| \leq r_i$*, who can veto candidates in* $M \setminus \{c\}$ *without changing the score of* c*, and make it a new winner.*

A ballot c_i *is a* best NMV response *of optimism horizon* r_i *if it is a better non-myopic Veto response to make* $c \in C_i^w$ *the winner, and there is no better non-myopic Veto response making* $c' \succ_i c$ *the winner.*

From previous results on myopic dynamics [Meir *et al.*, 2010; Lev and Rosenschein, 2012], which are a particular instance of Definitions 1 and 2 with all voters' optimism horizons being 0, it follows that better response strategies do not converge to a Nash equilibrium, unless, in Plurality, they are in a very particular structure [Meir, 2015]. Hence, we focus on best response functions. Also, as often in previous work, we make a natural assumption that the initial profile is the truthful one. We do not make any restrictions on the order in which the agents apply their moves.

We denote the policy of optimistic non-myopic best response under Plurality by ρ^{NMP}, and under Veto by ρ^{NMV}. For the latter, we require that if the current ballot is a best NMV, then ρ^{NMV} favors it over other best responses. Note that this response is coherent with the standard definition of better/best responses, in the sense that when the optimism horizon is zero, the non-myopic and the standard responses coincide. In addition, we must note that a stable state with respect to an optimistic non-myopic best response function is not necessarily a Nash equilibrium, nor the other way around. This is simply because ρ^{NMP} and ρ^{NMV}, unlike ρ^{BR}, are not myopic.

Example 1. *Assume that the Plurality voting rule is used, and that we have 4 voters and 4 candidates (named* a *through* d*), with the tie-breaking order* $a \succ b \succ c \succ d$*. Let the truthful preference profile be as follows:*

$$\begin{aligned} \textit{voter 1} &: a \succ_1 c \succ_1 b \succ_1 d \\ \textit{voter 2} &: c \succ_2 a \succ_2 b \succ_2 d \\ \textit{voter 3} &: d \succ_3 a \succ_3 b \succ_3 c \\ \textit{voter 4} &: b \succ_4 d \succ_4 c \succ_4 a \end{aligned}$$

In the myopic version, voter 4 deviates from the truthful profile by voting for d*. Depending on which between voters 1 and 2 makes the next move, the winner will be* a *or* c*.*

Now, suppose all voters have an optimism horizon of 2*. Since every candidate needs only one additional vote to win the election, all the voters will stick to their strategies (expecting other voters to follow them), and the winner will remain* a *(which is a good result in this case, as it ensures that a Condorcet winner is elected).*

In the following sections, we present our results. Some proofs are omitted due to space limitations.

3 Non-Myopic Plurality

In this section, we investigate the convergence of iterative voting dynamics guided by non-myopic response functions under the Plurality rule. In particular, we show that non-myopic response leads to convergence to a stable state if all voters use the same optimism horizon (Theorem 2). Importantly, as we observe in Theorem 1 below, ρ^{NMP} is essentially different from other, seemingly similar, dynamics, earlier considered in the literature, and hence, our convergence results do not follow from previous work.

Theorem 1. *The non-myopic optimistic dynamic ρ^{NMP} is not equivalent to myopic iterative plurality [Meir* et al.*, 2010], local dominance plurality [Meir* et al.*, 2014], and local regret minimization [Meir, 2015] dynamics.*

Theorem 2. *Assume all voters in N participate in an iterative voting scenario and use a non-myopic Plurality response function ρ^{NMP} with the same optimism horizon: $\forall i \in N$, $r_i = r \in \mathbb{N}$. Let $\mathbf{b}^t$ denote the complete voting profile at time t and $\mathbf{s}^t = \mathbf{s}_{\mathbf{b}^t}$ its corresponding score profile. Then, if $\mathbf{b}^0 = \mathbf{Q}$, there are $\tau \in \mathbb{N}$, $\mathbf{b}^*$ so that $\forall t > \tau$, $\mathbf{b}^t = \mathbf{b}^*$.*

Proof. Let $w^t = f(\mathbf{s}^t)$ be the winner at iteration t. Also denote by $A_i^t = A_i(\mathbf{s}^t) \subseteq M$ the set of all possible non-myopic 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 $U_i^t \subseteq M$ 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 NMP ballots.

We will show that two conditions hold simultaneously and lead to convergence as required:

1. $U_i^{t+1} \subseteq U_i^t$;
2. $\mathbf{s}^t(w^t) \leq \mathbf{s}^{t+1}(w^{t+1})$, and 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 non-myopic responses do not grow, and the score of the winner does not decrease, neither with regard to gathered points nor with regard to the tie-breaking order.

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

Case I 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 hold that $U_i^{t'} \setminus U_i^{t'-1} = \emptyset$ for all $t' \leq t$. In particular, for any $\tau \leq t$ $c \in U_i^\tau$. 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 non-myopic Plurality response, it has to hold 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 hold that $w^t = b_i^{\tau+1} \succ_i b_i^\tau = c$ (or $c \notin U_i^{\tau+1}$, reaching a contradiction).

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

Case II 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 complementary subcases, $b_i^{t+1} \neq w^{t+1}$ 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 point 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 optimism horizons are equal, it would also entail that $c \in A_i^t$—a 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 $c \notin A_i^t$ it 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 tie-breaking order) change. In particular, there exists 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 non-myopic 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 the optimism horizons are equal, which, as we next demonstrate, is a necessary condition for convergence to a stable state.

Theorem 3. *Iterative voting dynamics with non-myopic voters with different optimism horizons may not converge.*

Proof. We construct a scenario with a cycle of non-myopic responses. Let there be 5 candidates, named a through e, and let the preferences 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 NMP votes exists starting from $\mathbf{s}^t = \mathbf{s}^0$:

- Voter 3 changes 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 favor of c one after the other. At both of these changes $w^{t+1} = w^{t+2} = d$ and has 11 points, so neither is a among the NMP responses of voter 1, nor is b among the NMP responses of voter 2. Hence, the score profile becomes $\mathbf{s}^{t+3} = (3, 3, 7, 11, 10)$.

- c now becomes an NMP 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 NMP 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 tie-breaking 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 NMP responses of 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 Non-Myopic Veto

In this section, we prove convergence of non-myopic dynamics for the Veto voting rule. Surprisingly, this result is stronger than that for Plurality, showing that iterative non-myopic Veto converges even for voters with different optimism horizons.

Due to the nature of this rule, a non-myopic Veto response is, in fact, similar to a myopic best response, requiring the veto of the currently winning candidate. This does not mean though that convergence stems from previous results, as a stable state in the myopic scenario is not necessarily stable in the non-myopic model. However, our proof is similar in many respects to the myopic case found in [Lev and Rosenschein, 2012], and we keep similar notation, where applicable.

Theorem 4. *Assume all voters in N participate in an iterative voting scenario and use a non-myopic Veto response function ρ^{NMV} with individual optimism horizons. Let $\mathbf{b}^t$ denote the complete voting profile at time t and $\mathbf{s}^t = \mathbf{s}_{\mathbf{b}^t}$ its corresponding score profile. Then, there are $\tau \in \mathbb{N}$, $\mathbf{b}^*$ so that $\forall t > \tau$, $\mathbf{b}^t = \mathbf{b}^*$.*

Proof. We begin by assuming that the theorem is false, and that 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 $max(G_i)$.*

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

Proof. This follows as a special case of Lemma 1. □

Lemma 2. *$max(G_i)$ has at most two different values.*

If for every G_i there is only a single winner with the maximal score, 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. That is, 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 existence of a cycle. Therefore, we must assume that there is at least one state G_i in which there is more than one candidate with the maximal score. We shall call 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 $max(G_i) = max(G_0)$, we have that $\{c \in M | \mathbf{s}_{G_i}(c) \geq max(G_i) - 1\} = \{c \in M | \mathbf{s}_{G_0}(c) \geq max(G_0) - 1\}$ and $|\{c \in M | \mathbf{s}_{G_i}(c) = max(G_i) - 1\}| = |\{c \in M | \mathbf{s}_{G_0}(c) = max(G_0) - 1\}|$, $|\{c \in M | \mathbf{s}_{G_i}(c) = max(G_i)\}| = |\{c \in M | \mathbf{s}_{G_0}(c) = max(G_0)\}|$. That is, the set of candidates with score $max(G_0)$ or $max(G_0)-1$ is the same, as well as the number of candidates with each of these scores.*

Examining the set B of candidates for which there is a G_i where their score is $max(G_0)$ and there is a G_j where their score is $max(G_0)-1$ (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 $max(G_0)$ does not 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 (this is trivially true when $max(G_j) = max(G_0) + 1$, as there is a single winner there, and they just got the "bump" to that score)—which, as noted above, is a finite process, contradicting the endless cycle. □

5 Single Non-Myopic Voter

Previous sections have shown that the heterogeneity of optimism horizons has no effect on convergence 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 where only one voter has a non-zero optimism horizon. We show that this case does converge even under Plurality. However, the final result may not necessarily be of benefit to the optimistic player, under both rules.

Theorem 5. *The iterative process with one non-myopic voter under Plurality converges.*

Proof. As the proof is very similar to that in Theorem 2, we will provide the outline of the difference. For the single non-myopic voter the set U_i^t will behave similarly to that from the proof of Theorem 2, steadily growing smaller. On the other hand, between the iterations where the non-myopic voter changes his ballot, all other voters will behave myopically and their combined ballot will stabilize (see [Meir *et al.*, 2010]). Hence, all causes of a cycle can be the non-myopic's voter choosing the same strategy over and over, but the never-enlarging set of options precludes this. This overall will lead to the convergence of the complete voter profile. □

However, as we demonstrate next, non-myopicness is not necessarily an advantageous feature, and may lead to a sub-optimal stable state for the non-myopic player.

Example 2. *Assume that the Plurality voting rule is used, and that we have 4 voters and 4 candidates (named a through d), with tie-breaking order $\widehat{Q} = a \succ b \succ c \succ d$. Let the truthful preference profile be as follows:*

$$\begin{array}{ll} \textit{voter 1}: & a \succ_1 c \succ_1 b \succ_1 d \\ \textit{voter 2}: & c \succ_2 a \succ_2 b \succ_2 d \\ \textit{voter 3}: & d \succ_3 a \succ_3 b \succ_3 c \\ \textit{voter 4}: & b \succ_4 d \succ_4 c \succ_4 a \end{array}$$

Assume voter 4 is the only non-myopic voter with an optimism horizon $r_4 = 2$. It is easy to see that the truthful state is stable under non-myopic 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 $\mathbf{b} = (c, c, d, d)$. 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^{NMP}: when a voter has both a myopic and a non-myopic move, it 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 a Nash equilibrium in the ρ^{BR} based, standard game. That is, there will be no more myopic moves available. However, voter 4 will still have a non-myopic move—to return to their truthful vote. As a result of voter 4 reverting to his truthful ballot, voter 1 will now have a new myopic move to return to its own truthful vote. The initial (truthful) state will be restored. This forms a cycle under the myopic biased ρ^{NMP} (note though, that this is not in contradiction with Theorem 5 that deals with unbiased ρ^{NMP}).

This negative effect of non-myopic dynamics is more universal and also holds for the Veto voting rule that was shown to have stronger convergence properties than Plurality.

Example 3. *Assume that the Veto voting rule is used, and that we have 3 voters and 4 candidates (named a through d), with the tie-breaking order $\widehat{Q} = (a \succ b \succ c \succ d)$. Let the truthful preference profile be as follows:*

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

Assume voter 1 is the only non-myopic voter, with optimism horizon $r_1 = 1$. The truthful 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 non-myopic optimistic move: to switch from vetoing d 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^{NMV}, voter 1 adopts the non-myopic optimistic ballot leading to a voting profile of $\mathbf{b} = (a, d, d)$, and the new winner becomes b. However, now voters 2 and 3 have no incentive to change their ballot and voter 1 has no non-myopic optimistic ballots, i.e., $\mathbf{b} = (a, d, d)$ is a stable voting profile. Alas, $a \succ_1 b$, so voter 1 is worse off.

The above examples also provide another interesting observation: Under the Plurality rule, increasing the optimism horizon beyond $r = 2$ does not effect the overall iterative behavior of a system with a single non-myopic voter. This is because if a move is only viable for a larger optimism horizon, no myopic voter will change accordingly, rendering it futile. On the other hand, under the Veto rule, the increase of the optimism horizon will have a behavioral impact.

6 Discussion

In this paper, we addressed another, complex, condition for the convergence of iterative voting—myopic voting. While previous work dealt with different types of tie-breaking schemes, voting rules, best response dynamics and its restrictions, we have relaxed the crucial assumption that players will only look for an immediate change to the outcome. We have explored non-myopic dynamics 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 result that the results for Veto are stronger than those for Plurality), we also considered the effects of non-myopic dynamics on the outcome of the iterative process, and found that it is not necessarily beneficial for non-myopic voters. In this direction, it would also be interesting to study the effects of various numbers of non-myopic voters and of their horizons of optimism on the final outcome.

Moreover, as we explored a fundamentally optimistic outlook, one can consider other approaches to non-myopic strategies. While all non-myopic strategies have an optimistic component, as they make an assumption about the future, some may have a different goal. For example, voters may try to prevent their least favorite option within the horizon of optimism from winning.

Another interesting direction to explore in further research is the possibility and effects of relaxing the opacity of information to the players. While limited knowledge, when voters are only aware of the candidates' total scores at each state, makes sense in many scenarios, there are also situations where more information may be available. For example, it may be possible to estimate the distribution of preferences in the population of voters using pre-election polls, or, in settings with few voters, the iterative process may allow the voters to learn about others' preferences from the moves they made in previous iterations. Such knowledge may enable the players to make more subtly optimistic, even farsighted, moves.

Finally, a potential extension to our model is to endow the players with cardinal utilities for each candidate. This would impose a greater degree of nuance in voters' moves, as they would have to decide between supporting candidates requiring fewer additional vote changes to become victors, and more preferable candidates that have a longer horizon to win. Similar settings have been studied in general games, which may aid research in this direction.

References

[Airiau and Endriss, 2009] Stéphane Airiau and Ulle Endriss. Iterated majority voting. In Francesca Rossi and Alexis Tsoukias, editors, *Algorithmic Decision Theory*, volume 5783 of *Lecture Notes in Computer Science*, pages 38–49. Springer Berlin / Heidelberg, 2009.

[Bartholdi III *et al.*, 1989] John J. Bartholdi III, Craig A. Tovey, and Michael A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[Brânzei *et al.*, 2013] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. How bad is selfish voting? In *Proceedings of the 27th National Conference on Artificial Intelligence (AAAI)*, pages 138–144, Bellevue, Washington, July 2013.

[Desmedt and Elkind, 2010] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic commerce (EC)*, pages 347–356, Cambridge, Massachusetts, June 2010.

[Dutta and Laslier, 2010] Bhaskar Dutta and Jean-Franccois Laslier. Costless honesty in voting. in 10th International Meeting of the Society for Social Choice and Welfare, Moscow, 2010.

[Gibbard, 1973] Allan Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–602, July 1973.

[Grandi *et al.*, 2013] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and borda score. In *Proceedings of 3rd International Conference of Algorithmic Decision Theory (ADT)*, pages 181–192, Brussels, Belguim, November 2013.

[Lev and Rosenschein, 2012] Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 2, pages 611–618, Valencia, Spain, June 2012.

[Meir *et al.*, 2010] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Convergence to equilibria of plurality voting. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 823–828, Atlanta, July 2010.

[Meir *et al.*, 2014] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of the 15th ACM conference on Economics and computation (EC)*, pages 313–330, Palo Alto, Califronia, June 2014.

[Meir, 2015] Reshef Meir. Plurality voting under uncertainty. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*, pages 2103–2109, Austin, Texas, January 2015.

[Obraztsova *et al.*, 2013] Svetlana Obraztsova, Evangelos Markakis, and David Robert Martin Thompson. Plurality voting with truth-biased agents. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, pages 26–37, Aachen, Germany, October 2013.

[Obraztsova *et al.*, 2015] Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R. Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, Austin, Texas, January 2015.

[Rabinovich *et al.*, 2015] Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey S. Rosenschein. Analysis of equilibria in iterative voting schemes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, Austin, Texas, January 2015.

[Reijngoud and Endriss, 2012] A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *AAMAS*, 2012.

[Reyhani and Wilson, 2012] Reyhaneh Reyhani and Mark C. Wilson. Best-reply dynamics for scoring rules. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI)*, volume 2, pages 672–677, Montpellier, France, June 2012.

[Reyhani *et al.*, 2012] Reyhaneh Reyhani, Mark C. Wilson, and Javad Khazaei. Coordination via polling in plurality voting games under inertia. In *Proceedings of COMSOC*, pages 359–370, Krakòw, Poland, September 2012.

[Satterthwaite, 1975] Mark Allen Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, April 1975.

[Thompson *et al.*, 2013] David Robert Martin Thompson, Omer Lev, Kevin Leyton-Brown, and Jeffrey S. Rosenschein. Empirical aspects of plurality election equilibria. In *Proceedings of the 12th International Coference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 391–398, St. Paul, Minnesota, May 2013.

[Xia *et al.*, 2009] Lirong Xia, Michael Zuckerman, Ariel D. Procaccia, Vincent Conitzer, and Jeffrey S. Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 348–353, Pasadena, California, July 2009.