On the Convergence of Iterative Voting: How Restrictive should Restricted Dynamics be?

Svetlana Obraztsova   Evangelos Markakis   Maria Polukarov
Zinovi Rabinovich   Nicholas R. Jennings

Abstract

We study convergence properties of iterative voting procedures. Such a procedure is defined by a voting rule and a (restricted) iterative process, where at each step one agent can modify his vote towards a better outcome for himself. It has already been observed in previous works that if voters are allowed to make arbitrary moves (or even only best responses), such processes may not converge for most common voting rules. It is therefore important to investigate whether and which natural restrictions on the dynamics of iterative voting procedures can guarantee convergence.

To this end, we provide two general conditions on dynamics based on iterative myopic improvements, each of which is sufficient for convergence. We then identify several classes of voting rules, along with their corresponding iterative processes, for which at least one of these conditions hold. Our work generalizes recent results and relaxes a number of restrictive assumptions made in previous research.

1 Introduction

Voting rules constitute a popular tool for preference aggregation and collective decision making in various multi-agent systems, involving entities with possibly diverse preferences. The major concern however of voting as a decision-making process is that voters may have incentives to misreport their real preferences in order to favor certain candidates. The famous Gibbard-Satterthwaite theorem [7, 16] states that under mild assumptions, such scenarios are feasible. Strategic behavior is therefore inherent in most voting rules.

Since manipulation is unavoidable, a natural next step would be to resort to game-theoretic tools in order to model voting behavior and assess the outcome of a voting process. A first attempt is to view elections as games among strategic agents and study the properties of Nash equilibria or other solution concepts. This is not enough however. It has long been observed that Nash equilibria may fail to provide a good prediction of how such voting games may evolve. The reason is that there is usually a multitude of equilibria, some of which are also very unlikely to be observed in practice. A representative example of this is when all voters vote for a commonly least preferred candidate, and a unilateral deviation cannot change the outcome.

Departing from this, we are interested in solution concepts that are more likely to be realized in a voting process. Following [11], in this work we consider iterative voting procedures, where agents start from some initial configuration (e.g., the truthful profile), and subsequently make myopic improvements. Iterative voting has recently received significant attention in the literature, motivated by web services such as Doodle, but also as a way to model the response of an electorate to information polls.

Several results have already been obtained regarding the convergence of best/better response dynamics for various voting rules (see Related Work). One of the main findings is that if players are allowed to make arbitrary moves (or just play only best responses), then convergence of such processes is not guaranteed. Nevertheless, it is often the case that voters will not choose a best response when they update their voting decision. Natural restrictions may apply if the voters are computationally bounded, or if they tend to behave myopically and make only greedy local moves. Voters may also not be interested in plain manipulation but rather in reaching an agreement or compromise with a minimal way of changing their ballot. Restrictions on the dynamics of voting processes as an attempt to model some of the above considerations have already been described recently in [14] and [8]. However, the family of restricted moves that enforces convergence has not been yet characterized.

Contribution: We consider iterative voting procedures, specified by a voting rule along with the allowed dynamics, i.e., a specification of the possible improvement steps that can be made by the voters. Under lexicographic tie-breaking, we provide two general conditions on these processes, each of which is sufficient for convergence. The first one is based essentially on potential function arguments, whereas the second regards the monotonicity of a set sequence defined along an improvement path. We then identify several classes of voting rules, along with restricted dynamics under these rules, for which at least one of these conditions hold. Our work provides a unifying framework for studying convergence issues of such processes. Furthermore, we generalize recent results and relax a number of restrictive assumptions made in previous research.

1.1 Related Work

Some earlier work on iterative voting processes is well summarized in [9], and concerns dynamics for deciding on allocations of public goods. The study of iterative voting in the more recent AI literature was initiated by [11], which focused on improvement dynamics under the Plurality rule. Both positive and negative results were established, depending on the initial voting profile, the tie-breaking rule and the improvement steps allowed (better replies or best responses). In a follow up work, [10] showed that for other voting rules (in fact for the most popular ones), it is often the case that convergence of even best responses cannot be guaranteed. More results along this direction were established in [15], who also improved on the convergence bounds of [11]. An analysis in terms of the quality of equilibria reachable by such processes was obtained in [3]. Finally, a different type of voting processes was studied in [1], where each player is allowed to propose a change in the current state and then a vote is held for its acceptance.

The works most closely related to ours are those of [14] and [8]. Both of them considered procedures where voters may not play a best response but instead move according to certain restricted best responses. Three types of restricted dynamics have been considered (defined in Preliminaries), and convergence results were established for these moves under some families of voting rules. As we exhibit in the following sections, our framework incorporates these positive results and relaxes some of the limitations on the allowed moves.

Finally, there have been many other works applying game-theoretic tools to voting, starting with [6]. More recent research in this line has focused either on studying stronger equilibrium concepts [17] or on different models of voting behavior such as voting with abstentions [4] or truth-biased voting [11, 18, 5, 13]. We do not consider any of these models here.

2 Preliminaries

We first recall the common voting rules studied in this work, and define the setting of iterative voting based on myopic improvement moves by single voters.

2.1 Voting rules

There is a set $V = \{1, \dots, n\}$ of n *voters* (or *agents*) electing a winner from a set $C = \{c_1, \dots, c_m\}$ of m *candidates* (or *alternatives*). Let $\mathcal{L}(C)$ be the set of all strict linear orders on C. Each voter i submits a *vote* (or *ballot*) $b_i \in \mathcal{L}(C)$, which may or may not coincide with his real *preference order*, $\succ_i \in \mathcal{L}(C)$, over the candidates. A *profile* $\mathbf{b} = (b_1, \dots, b_n) \in \mathcal{L}(C)^n$ is a vector of votes, one for each agent. As is common, we denote by $\mathbf{b}_{-i}$ the profile of all votes except that of agent i, so that $\mathbf{b} = (b_i, \mathbf{b}_{-i})$. A *voting rule* $\mathcal{F} : \mathcal{L}(C)^n \to 2^C$ takes a voting profile as input, and produces an outcome—a nonempty subset of candidates, called *the winners* of the election. In this paper, we focus on *resolute* voting rules $\mathcal{F} : \mathcal{L}(C)^n \to C$, which always return a single winner. That is, given their irresolute version, we assume that ties are broken according to a fixed *tie-breaking rule*. Specifically, in this work, we assume lexicographic tie-breaking—i.e., ties are broken in favour of the candidate with the lowest index.

Examples of common voting rules include:

- *Positional scoring rules* (PSRs). Each such rule is associated with a scoring vector $(s_1, ..., s_m)$ where $s_1 > s_m$ and $s_1 \geq s_2 \geq \ldots \geq s_m$. If a voter ranks a candidate at the j-th position, the candidate receives a score of s_j from this vote. The total score of a candidate is the sum of scores over all the votes, and the winner of the election is the candidate with the highest score. This family of rules includes Plurality with the scoring vector $(1, 0, 0, \ldots, 0)$, Veto with $(1, 1, \ldots, 1, 0)$, Borda with $(m-1, m-2, ..., 0)$ and k-approval with $(1, ..., 1, 0, ..., 0)$.

- *Maximin*. Under this rule, the score of a candidate c is the minimum number of voters who prefer c over all pairwise comparisons with the other candidates. The candidate with the highest such score wins the election.

- *Copeland*. The score of a candidate c is the number of pairwise comparisons he wins (i.e., the number of other candidates c', for which the majority of voters prefers c to c'), minus the number of pairwise comparisons he loses. The winner has the highest such score.

- *Bucklin*. In one of its versions, this rule first identifies for each candidate c, the minimum number k for which the majority of voters rank c within their top k choices. Let k_{min} be the minimum such number over all candidates. The election then proceeds as a k_{min}-approval election.

Under the rules defined above, each candidate can be naturally associated with a *score*, derived from a given voting profile. For rules where there is no obvious way to score the candidates, we can also define an artificial score where the winner under a given profile receives 1 point, and other candidates receive 0 points. Thus, w.l.o.g. we can assume that any voting rule corresponds to a scoring algorithm with the property that the candidate with the highest score wins the election (after possibly applying a tie-breaking rule as well). We may also assume that the scores are *integer* numbers. Note that there can be several scoring rules corresponding to a voting rule; in the sequel, whenever we are given a voting rule, we will also assume that it is accompanied by a fixed scoring rule (the natural one when it comes to the voting rules that are defined above). For each candidate $c \in C$, its score at profile $\mathbf{b}$ under voting rule $\mathcal{F}$ is denoted by $s_{\mathcal{F}}(c, \mathbf{b})$ (we drop the indices when clear from the context).

2.2 Iterative voting

Each voting rule $\mathcal{F}$ induces a natural game form, where the strategies available to each voter are given by $\mathcal{L}(C)$, and the outcome of a joint action (i.e., a voting profile) $\mathbf{b}$ is $w_{\mathbf{b}} = \mathcal{F}(\mathbf{b})$. Voter i *prefers* profile $\mathbf{b}'$ over profile $\mathbf{b}$ if $w_{\mathbf{b}'} \succ_i w_{\mathbf{b}}$, and we say that $b_i \xrightarrow{i} b'_i$ is an *improvement move* (or a *better reply*) of agent i w.r.t. $\mathbf{b}$ if he prefers $(b'_i, \mathbf{b}_{-i})$ over $\mathbf{b}$.

A *path* in $\mathcal{L}(C)^n$ is a sequence $(\mathbf{b}^0 \to \mathbf{b}^1 \to \cdots)$ of voting profiles such that for every $k \geq 1$ there exists a unique agent, say voter i, such that $\mathbf{b}^k = (b'_i, \mathbf{b}^{k-1}_{-i})$ for some $b'_i \neq b^{k-1}_i$ in $\mathcal{L}(C)$. It is an *improvement path* if for all $k \geq 1$ it holds that $b^{k-1}_i \xrightarrow{i} b^k_i$ is an improvement move, where i is the unique deviator at step k. The setting of *iterative voting* is based on myopic improvement dynamics as above: the voters start by announcing some initial vote, and then proceed and change their votes in turns, one at a time, up until no one has an objection to the current outcome. As often in previous works, we make a natural assumption that the initial profile is the *truthful* one—that is,

$\mathbf{b}^0 = (\succ_1, \dots, \succ_n)$. We do not make any restrictions on the order in which the agents apply their improvement moves.

Convergence of better replies is not guaranteed though, even for games induced by the simple Plurality rule, hence another natural restriction of *best response* dynamics is usually made: the voters are assumed to make only moves that yield the best possible candidate (w.r.t. the deviator's preferences) among the potential winners at each step. While best responses always converge for Plurality and Veto with linear tie-breaking [11, 10], they may cycle under other rules, such as Copeland [8], Borda, and k-approval [10, 15].

2.2.1 Restricted dynamics

In such settings, convergence can be achieved by restricting the sets of available improvement moves even further. To this end, the following simple dynamics have been previously considered:

- *Second Chance* (SC)[1] [8]: If the current winner is not the deviator's best or second-best choice, he moves his second-best alternative to the top position;
- *k-pragmatist* [2, 14]: The deviator moves his favourite among the k currently highest ranked alternatives to the top position, without changing the relative ranking of the others;
- *Best Upgrade* (BU)[1] [8]: The deviator moves to the top position his favourite alternative among those who can win the election and are currently ranked in the deviator's ballot above the current winner.

We use the term *iterative voting procedure* $(\mathcal{F}, D)$ to define the process based on the improvement dynamics D under the voting rule $\mathcal{F}$. We say that a voting procedure *converges* if every improvement path that contains moves allowed by D is finite under $\mathcal{F}$.

3 Sufficient conditions for convergence

We identify two general conditions on iterative voting procedures, each of which guarantees convergence from the truthful state. These conditions are powerful enough to incorporate all convergence results for the restricted dynamics defined above. Finally, we also propose two natural yet significant relaxations for the BU dynamics and show that they satisfy one of our conditions for most voting rules.

3.1 Function monotonicity

The first condition is based on the potential argument [12]. That is, we define a real-valued function $G : \mathcal{L}(C)^n \to \mathbb{R}$ over the set of voting profiles, and require that it increases along any allowed improvement path. In fact, it is sufficient to require only weak monotonicity, as function $G(\cdot)$ changes its value at every step.

Condition 1 (C1)**.** *Let $(\mathcal{F}, D)$ be an iterative voting procedure. Given a voting profile $\mathbf{b} \in \mathcal{L}(C)^n$, let*

$$G(\mathbf{b}) = s_{\mathcal{F}}(w_{\mathbf{b}}, \mathbf{b}) + \frac{m - index(w_{\mathbf{b}})}{m+1}, \tag{1}$$

where for any candidate c, $index(c)$ indicates its position in the tie-breaking order. Then, for any improvement path $(\mathbf{b}^0 \to \mathbf{b}^1 \to \cdots)$, we have $G(\mathbf{b}^k) \geq G(\mathbf{b}^{k-1})$, $\forall k \geq 1$.

[1]Also referred to as M1 and M2 in [8].

As will become clear from Theorem 1 below, Condition 1 is existential with respect to the scoring function (i.e., for our purposes we only need C1 to hold for some scoring function consistent with a voting rule $\mathcal{F}$). Nevertheless, as explained in Section 2, in this paper we will only consider the natural score function that is associated with each of the rules we are studying.

In what follows, we abuse the notation and write $w_k = w_{\mathbf{b}^k}$ and $s_k(\cdot) = s_{\mathcal{F}}(\cdot, \mathbf{b}^k)$ for a profile $\mathbf{b}^k$, at step k of a path $(\mathbf{b}^0 \rightarrow \mathbf{b}^1 \rightarrow \cdots)$, under a given $(\mathcal{F}, D)$.

Theorem 1. *Any iterative voting procedure $(\mathcal{F}, D)$ that satisfies* C1*, converges in at most $(s^{\max}_{\mathcal{F}} + 1)(m+1)$ steps, where $s^{\max}_{\mathcal{F}}$ is the maximum possible attainable score under $\mathcal{F}$.*

Proof. Let $(\mathbf{b}^0 \rightarrow \mathbf{b}^1 \rightarrow \cdots)$ be any improvement path in $(\mathcal{F}, D)$. Since $\frac{m - index(c)}{m+1} < 1$ for any $c \in C$, and C1 holds, we have that either $s_k(w_k) > s_{k-1}(w_{k-1})$ or $s_k(w_k) = s_{k-1}(w_{k-1})$ and $index(w_k) < index(w_{k-1})$ for any $k \geq 1$. Hence, the function G grows by at least $\frac{1}{m+1}$ at each step, and the bound follows since $G(\mathbf{b}) < s^{\max}_{\mathcal{F}} + 1, \forall\, \mathbf{b}$. □

3.2 Set monotonicity

The second condition follows the idea of [15]. In their work, convergence of best response dynamics for Plurality was (re)proved by showing inclusion monotonicity of the sets of *potential winners* along the improvement path—that is, the sets of candidates for which there exists a voter that can make them win the election by unilaterally applying an improvement move at a given step. The condition we give below is stronger and requires monotone inclusion of individual sets of potentially winning candidates for each voter separately. Moreover, our definition is recursive so that a current winner of the election belongs to the set of potential winners of a voter i, only if it has or could have become a winner due to voter i's move.

Definition 1. *Let $(\mathcal{F}, D)$ be an iterative voting procedure. For $i \in V$ and an improvement path $(\mathbf{b}^0 \rightarrow \mathbf{b}^1 \rightarrow \cdots)$, let*

$$PW_i(\mathbf{b}^0) = \{w_0\} \bigcup \left\{ c \in C \mid \exists b_i^{'} : c = \mathcal{F}\left(b_i^{'}, \mathbf{b}^0_{-i}\right) \wedge c \succ_i w_0 \right\}$$

where $b_i^{'} \in \mathcal{L}(C)$ is consistent with D. For $k \geq 1$, let

$$\begin{aligned} PW_i(\mathbf{b}^k) \quad &= \quad \left\{ c \in C \mid \exists b_i^{'} : c = \mathcal{F}\left(b_i^{'}, \mathbf{b}^k_{-i}\right) \wedge c \succ_i w_k \right\} \\ &\bigcup \quad \begin{cases} \{w_k\} & w_k \in PW_i(\mathbf{b}^{k-1}) \\ \emptyset & \text{otherwise} \end{cases} \end{aligned}$$

Condition 2 (C2)**.** *Let $(\mathcal{F}, D)$ be an iterative voting procedure. Then, for any improvement path $(\mathbf{b}^0 \rightarrow \mathbf{b}^1 \rightarrow \cdots)$ in $(\mathcal{F}, D)$, we have, $PW_i(\mathbf{b}^k) \subseteq PW_i(\mathbf{b}^{k-1})$, $\forall i \in V, k \geq 1$, and at least one of the following holds:*

(a) *at each step k, there exists an agent $i \in V$ for whom the inclusion is strict;*

(b) *there is a finite number q,*[2] *such that for every voter i and candidate c, the maximum possible number of consecutive moves, made by i in favor of c, is bounded by q.*[3]

Theorem 2. *Any iterative voting procedure $(\mathcal{F}, D)$ that satisfies* C2*, converges in at most qmn steps.*

Proof. Consider any improvement path $(\mathbf{b}^0 \rightarrow \mathbf{b}^1 \rightarrow \cdots)$ in $(\mathcal{F}, D)$. If C2(a) holds, then at every improvement step k, the set of potential winners for some voter i, $PW_i(\mathbf{b}^k)$ decreases by at least 1, compared to $PW_i(\mathbf{b}^{k-1})$. This implies that the process converges in at most mn steps.

[2]Usually, a simple funciton of m or n. For instance, as we will show, $q = m - 1$ for many positional scoring rules.

[3]We refer to moves that are consecutive in the subsequence of moves made by voter i only. They need not to be consecutive in the whole improvement path.

For C2(b), fix a voter i, and consider the subsequence of steps where i makes a move. For a candidate c, look at the moves in this subsequence that are in favor of c. We claim that there can be at most one series of consecutive moves in this subsequence where i makes c a winner. Let t be the last step in such a series of moves. Then, at step $t+1$ (regardless of which agent is the deviator), candidate c is excluded from the set of potential winners of voter i, and by the monotonicity of these sets, it will not reappear in $PW_i(\cdot)$ ever again. Hence, each voter makes at most q moves in total in favor of each of the candidates, and the bound of qmn follows. □

3.3 Non-equivalence between C1 and C2

Next, we observe that the function and the set monotonicity conditions do not imply each other. We start with Example 1 of a voting procedure where C1 does not hold, but C2 does.

Example 1. There are 9 voters and m candidates for some large enough m. Figure 1 shows the truthful preference profile, where all the missing candidates within the first 4 positions of each voter are distinct dummy candidates, different from c_1 and c_2. In particular, c_1 appears at position 5 or lower for voters 1, 4 and 5. Ditto for c_2 and voters 6–9. The voting rule is Bucklin, and the iterative

1	2	3	4	5	6	7	8	9
c_2	c_2	c_2	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	c_1
$\dots$	c_1	c_1	c_2	$\dots$	$\dots$	$\dots$	c_1	$\dots$
$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$
$\dots$	$\dots$	$\dots$	$\dots$	c_2	c_1	c_1	$\dots$	$\dots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

Figure 1: Voters' preferences in a game with C2 but not C1

improvement policy restricts each of the voters to a single swap where one candidate is moved one position up in the current ranking, if this candidate becomes a new winner after such a move.

Under the truthful profile $\mathbf{b}^0$, the winner is c_1, with a score of 6 (the Bucklin winning round is 4). Consider now the following sequence ($\mathbf{b}^0 \to \mathbf{b}^1 \to \mathbf{b}^2 \to \mathbf{b}^3 \to \mathbf{b}^4$) of updates under the dynamics specified: first, voter 5 moves c_2 upwards by one position, making c_2 the new winner in $\mathbf{b}^1$. Then, voter 6 changes his vote and ranks c_1 in the 3rd position, making c_1 the winner in $\mathbf{b}^2$. Voter 5 responds by ranking c_2 in the 2nd position. Finally, voter 6 again lifts c_1 by one position. One can check that no voter can change his vote in $\mathbf{b}^4$ to make c_2 or any other candidate a winner.

To see that condition C1 is violated, note that in the truthful profile, the score of c_1 is 6, and observe that in $\mathbf{b}^1$, the winner's score decreases to 5, implying $G(\mathbf{b}^1) < G(\mathbf{b}^0)$.

We show now that condition C2 holds, in particular C2(b). First, look at the sets of potential winners for voters 5 and 6, who are involved in the improvement path. For 5, we have $PW_5(\mathbf{b}^0) = \{c_1, c_2\}$ and $PW_5(\mathbf{b}^1) = \{c_2\}$. Then this set remains unchanged until eventually we get $PW_5(\mathbf{b}^0) = \emptyset$. For voter 6, his initial set, which is $PW_6(\mathbf{b}^0) = \{c_1\}$, remains unchanged until the last step. Now, agent 7 has the same preference order as voter 6 but does not make a move. He has $PW_7(\mathbf{b}^0) = PW_7(\mathbf{b}^1) = PW_7(\mathbf{b}^2) = \{c_1\}$, and then $PW_7(\mathbf{b}^2) = \emptyset$, which remains further unchanged. Similarly, we check the monotonicity of these sets for the rest of the voters. Finally, it is trivial that the number of possible consecutive moves of a voter in favor of a certain candidate is at most $m-1$, hence C2(b) holds. □

Remark 1. *The game in Example 1 satisfies* C2(b)*, but it is easy to see that it does not satisfy* C2(a)*. One can make slight adjustments to this example, so that condition* C1 *still does not hold but* C2(a) *does. We omit the details due to space limitations.*

Next, we give an example where C1 holds but neither version of condition C2 does.

Example 2. The construction is based on the (Borda, BU) procedure. We construct an instance where voter 1 has the preference order $c_1 \succ_1 c_2 \succ_1 c_3 \succ_1 \dots$, and the ranking of voter 2 is $c_1 \succ_2 c_3 \succ_2 c_4 \succ_2 \dots \succ_2 c_2$. The preferences of the other voters are such that in the full truthful profile, $\mathbf{b}^0$, the following conditions hold:

- candidate c_3 is the winner; we denote its score by s;
- candidate c_2 has $s-1$ points in $\mathbf{b}^0$;
- candidates $c_4, c_5, \dots, c_m$ have s points in $\mathbf{b}^0$ (but they all lose due to tie-breaking);
- candidate c_1 has less than $s-1$ points.

We first exhibit that C2 does not hold. Under BU, voter 1 can swap the positions of c_1 and c_2, resulting in candidate c_2 having a score of s and winning the election by the tie-breaking. Let $\mathbf{b}^1$ be this new profile, and consider voter 2. At $\mathbf{b}^0$, he had $PW_2(\mathbf{b}^0) = \{c_2, c_3\}$. But then at $\mathbf{b}^1$ the current winner is the last choice of voter 2, hence under BU there are many candidates that he can turn into a winner. Thus, $PW_2(\mathbf{b}^1) = \{c_2, c_3, ..., c_m\}$, so C2 is violated.

Finally, C1 holds by the results of [8], falling under the first case of Proposition 3 below.

The following Proposition demonstrates that our framework incorporates the convergence results for the restricted dynamics previously proposed in the literature [14, 8].

Proposition 3. *The function monotonicity holds for all following iterative voting procedures: (i) k-pragmatist under PSRs; (ii) SC and BU under PSR, Copeland or Maximin. Furthermore, SC also satisfies* C2.

Proof. We provide a sketch of the proof since most of the main arguments can be extracted by the works of [14, 8].

For BU, essentially the proof of Theorem 2 in [8], establishes C1. The main idea is that since the only candidates who are already ranked higher than the current winner, say c, can be moved, this implies that the score of c remains unchanged. If this score was equal to s, then the new winner must either win with a score higher than s, or he wins with a score equal to s but is favored by the tie-breaking rule, i.e., he has a lower index than c. This implies the monotonicity of the function G, given its dependence on the score and the index, as defined in (1). Hence, C1 holds.

For SC, again, a similar argument applies. By the definition of the move, each voter will change his vote at most once from the truthful profile. And his move is to place on top his second favourite candidate, when the current winner is not his best or second-best alternative. Hence again, the score of the current winner never decreases and the function G is monotone. To see why C2 holds, observe that in the beginning of the process the set of potential winners for a voter contains the current (truthful) winner and possibly the second best alternative. If a voter i moves to make the second best alternative a winner, the truthful winner no longer belongs to $PW_i(\cdot)$. If any move is made afterwards by another voter, then $PW_i(\cdot)$ becomes empty and i cannot make his second best candidate a winner again in the future. Hence, the set monotonicity holds. It is also then trivial to verify that both of its versions C2(a) and C2(b) hold in this case.

Finally, regarding k-pragmatist, by the proof of Lemma 2 in [14], we know that along any improvement path, the set of the k highest ranked candidates does not change. Let S be this set. Now, since a deviator moves his favourite candidate among S to the top and keeps the relative ranking of the others the same, this means that no candidate from S loses any points under a PSR. Thus again, the score of the current winner does not decrease. □

3.4 Relaxations of Best Upgrade dynamics

Here, we observe that the previously discussed BU dynamics can be significantly relaxed, while preserving convergence. In particular, consider the following relaxations:

- BU-1: As in BU, the deviator moves to the top position in his vote his favourite alternative among those who can become a new winner and are currently ranked in the deviator's ballot above the current winner, c. He can also freely shuffle among themselves all the candidates ranked above c. The same applies for all the candidates below c (again among themselves). He is not allowed to change the ranking of c.
- BU-2: As before, the deviator moves to the top position in his vote his favourite alternative among those who can become a new winner and are currently ranked in the deviator's ballot above the current winner, c. He is again restricted to keep the rank of c unchanged, but can now shuffle the remaining candidates absolutely freely.

We can show the following result.

Theorem 4. *Both* BU-1 *and* BU-2 *satisfy* C1 *under PSRs. Furthermore,* BU-1 *satisfies* C1 *under Copeland and Maximin.*

Proof. The main argument is similar to Proposition 3. Let c be the winner at some profile along an improvement path. Under BU-1, the score of c does not change. This implies that the score of the new winner is at least as high. This holds both for PSRs but also for Maximin and Copeland (because the part of the majority graph involving c remains unchanged). Hence, as in Proposition 3, it is easy to see that the function G is monotone.

With BU-2, if we freely shuffle all the candidates apart from c, then the score of c will still remain the same under a PSR. Hence the same arguments apply. We can no longer guarantee convergence for Maximin and Copeland though, since the score of c may change (the part of the majority graph regarding c now changes). □

4 More monotone classes of iterative procedures

In this section, we demonstrate more iterative voting procedures, for which both the function monotonicity and the set monotonicity conditions hold. Our first example is for the Maximin rule with a natural improvement dynamic, termed *upgrade*, where a deviating voter moves a new winner to a higher position, while the relative ranking of the remaining candidates is kept unchanged. Our second family of rules concerns a subclass of integer PSRs, termed *unit gap* scoring rules, where the difference in any two consecuitive scores s_j, s_{j+1} in the integer scoring vector of the rule is bounded by 1 (among others, this class contains all common PSRs such as Plurality, Veto, k-approval and Borda). We show that C1 and C2 hold under the iterative process, called *unit upgrade*, where a new winner is moved by exactly one position higher in the ballot of the deviating voter, without changing the relative ranking of other candidates. Importantly, restricting the voters to not move higher any candidates other than the new winner, is necessary for convergence for both Maximin and unit gap scoring rules.

4.1 Maximin with upgrade

Consider the following policy for improvement moves:

- *upgrade* (U): at each step, the deviator moves his favorite alternative among those who can win the election, to a higher (but not necessarily top) position in his vote, and keeps the relative ranking of the rest of the candidates unchanged. The upgraded candidate is the new winner.

Theorem 5. *The iterative procedure (Maximin, U) is function monotone and set monotone—specifically,* C2(a) *holds.*

The following lemma demonstrates a useful property.

Lemma 1. *Let* $(\mathbf{b}^0 \to \mathbf{b}^1 \to \cdots)$ *be an improvement path under Maximin. For each candidate* $c \in C$*, let* $TO_c(\mathbf{b}^k)$ *be the set of his toughest opponents—i.e., candidates against which* c *has minimal support in all pairwise comparisons:*

$$TO_c(\mathbf{b}^k) = \arg \min_{x \in C \setminus \{c\}} n_k(c, x)$$

where $n_k(c, x)$ *is the number of voters that prefer* c *over a candidate* x *in a profile* $\mathbf{b}^k$*. For any* $k \geq 1$*, if* $s_k(c) > s_{k-1}(c)$ *then* $TO_c(\mathbf{b}^{k-1}) \subseteq TO_c(\mathbf{b}^k)$*.*

Proof. Since $\min_{x \in C \setminus \{c\}} n_k(c, x) = s_k(c) > s_{k-1}(c) = \min_{x \in C \setminus \{c\}} n_{k-1}(c, x)$, at step k the deviating voter awards candidate c an additional point against each of his toughest opponents at step $k - 1$ (by moving c from under $x \in TO_c(\mathbf{b}^{k-1})$ above them). Thus, all of them must remain his toughest opponents at step k. □

Proof of Theorem 5. Assume on the contrary that C1 or C2 do not hold. Let $t \geq 1$ be the first step on the upgrade path $(\mathbf{b}^0 \to \mathbf{b}^1 \to \cdots)$ where monotonicity breaks—that is, $G(\mathbf{b}^k) \geq G(\mathbf{b}^{k-1})$ and $PW_i(\mathbf{b}^k) \subseteq PW_i(\mathbf{b}^{k-1})$, for every $1 \leq k \leq t-1, i \in V$.

<u>*Case 1*</u>*:* Assume first that $G(\mathbf{b}^t) < G(\mathbf{b}^{t-1})$. This is only possible if the Maximin score of the winner at step $t-1$ decreases at step t: that is, $s_t(w_{t-1}) < s_{t-1}(w_{t-1})$. If this was not the case, then by the definition of G it follows that w_{t-1} has at least the same score and a lower index than w_t, a contradiction. By the upgrade policy, this means that the deviator at step t (say, voter i) moves candidate w_t from under w_{t-1} above w_{t-1} in his ballot. Since i prefers w_t to w_{t-1} (as he is improving at step t), there was a step $k < t$ at which voter i made w_{t-1} a winner. This is due to the upgrade policy, the fact that the process starts from the truthful state, and the fact that w_{t-1} was ranked higher than w_t at b_i^{t-1}. That is, w_{t-1} was the most preferable candidate among potential winners of i at step $k-1$, and hence $w_{t-1} \in PW_i(\mathbf{b}^{k-1})$ but $w_t \notin PW_i(\mathbf{b}^{k-1})$. However, $w_t \in PW_i(\mathbf{b}^{t-1})$, in contradiction to the set monotonicity until step $t-1$.

<u>*Case 2*</u>*:* Suppose now that $G(\mathbf{b}^t) \geq G(\mathbf{b}^{t-1})$, but $PW_i(\mathbf{b}^t) \not\subseteq PW_i(\mathbf{b}^{t-1})$ for some $i \in V$. Let $c \in C$ and $i \in V$ such that $c \in PW_i(\mathbf{b}^t) \setminus PW_i(\mathbf{b}^{t-1})$. First assume that the set of c's toughest opponents decreased at step t: that is, $TO_c(\mathbf{b}^{t-1}) \not\subseteq TO_c(\mathbf{b}^t)$, so there is a candidate $c' \in TO_c(\mathbf{b}^{t-1})$ with $c' \notin TO_c(\mathbf{b}^t)$. By Lemma 1, we have that $s_t(c) \leq s_{t-1}(c)$. In fact, equality is not possible here. To see this, note that c is not the winner at step t, otherwise c would belong to $PW_i(\mathbf{b}^{t-1})$, by the definition of $PW_i(\cdot)$. Hence c does not receive any additional points from the deviating voter at step t. But then, the only way that candidate c' can stop being a toughest opponent of c at step t, is if some other toughest opponent is moved by the deviator from beneath c, to a position above c, implying that the Maximin score of c decreases by 1: $s_t(c) = s_{t-1}(c) - 1$. Now, since $G(\mathbf{b}^t) \geq G(\mathbf{b}^{t-1})$, and hence, $s_t(w_t) \geq s_{t-1}(w_{t-1})$ (with equality only if w_t beats w_{t-1} in tie-breaking), it turns out that c cannot belong to the set of potential winners of any of the voters at step t, a contradiction. Thereby, we have established that $TO_c(\mathbf{b}^{t-1}) \subseteq TO_c(\mathbf{b}^t)$.

Since $c \in PW_i(\mathbf{b}^t)$, and c is not the winner at step t, voter i can increase the score of c—that is, in his ballot all toughest opponents of c are ranked above c. However, since $c \notin PW_i(\mathbf{b}^{t-1})$, this was not the case at step $t-1$. This means that one of the toughest opponents of c at step $t-1$ was ranked below c in voter i's ballot—and by moving this candidate above candidate c, the score of candidate c decreased by 1. But this excludes c from the sets of potential winners for all the voters, again a contradiction.

Finally, note that at the first step, all the voters who cannot make the truthful winner, w_0, a winner again (that is, all those who rank w_0 above at least one of its toughest opponents—certainly, there is at least one such vote), lose w_0 from their set of potential winners. Similarly, at each step k, the voter who deviated at the previous step will lose the previous winner, w_{k-1}, from his set of potential winners. Hence, C2(a) holds. □

Next, we argue that the requirement of upgrading (i.e., moving up) *only* the winning candidate is necessary for convergence under many voting rules. For instance, cycles have been shown for

Copeland [8], k-approval and Borda [10, 15], even when lexicographic tie-breaking is used. For Maximin, [10] provide a cycling example with deterministic, but not lexicographic, tie-breaking. We strengthen this negative result, by giving an example where ties are broken lexicographically.

Example 3. There are 2 voters $\{1, 2\}$ and 4 candidates $\{a, b, c, d\}$, with $d \succ b \succ c \succ a$ as the tie-breaking rule. At first step, the agents vote sincerely, seen below, which results in a tie between all the candidates, and d wins. Now, as voter 1 perefers candidate b over candidate d, he deviates from his true preference order $abdc$ and votes $abcd$, which makes b win the election (note that this is a best response for voter 1, and it involves moving a non-winning candidate). Next, voter 2 deviates to make candidate c a winner, and so on. We describe the improvement path below, with a cycle starting at the fourth step:

$$\begin{array}{rcl}
(abdc, cdba)\{d\} \xrightarrow{1} (abcd, cdba)\{b\} & \xrightarrow{2} & (abcd, cadb)\{c\} \\
\xrightarrow{1} (bcda, cadb)\{b\} & \xrightarrow{2} & (bcda, adcb)\{d\} \\
\uparrow_1 & & \downarrow_1 \\
(abdc, cadb)\{c\} & \xleftarrow{2} & (abdc, adcb)\{a\}
\end{array}$$

4.2 Unit gap scoring rules with unit upgrade

Let $\mathcal{F}$ be a PSR given by a scoring vector $(s_1, ..., s_m)$. We say that $\mathcal{F}$ is a *unit gap* scoring rule if $s_j - s_{j+1} \leq 1$ for any $j = 1, \ldots, m-1$. This includes the most common PSRs, such as, Plurality, Veto, k-approval and Borda.

For such rules, we restrict the upgrade policy even further:

- *unit upgrade* (UU): at each step, the deviator moves his favorite alternative among those who can become a winner, by exactly *one* position up, and keeps the relative ranking of the others unchanged. The upgraded alternative wins.

Theorem 6. *Let $\mathcal{F}$ be a unit gap scoring rule. Then, the iterative procedure ($\mathcal{F}$, UU) is both function monotone and set monotone—specifically, it satisfies* C2(b).

Proof. Assume on the contrary that C1 or C2 do not hold. Let $t \geq 1$ be the first step on the UU path $(\mathbf{b}^0 \to \mathbf{b}^1 \to \cdots)$ where monotonicity breaks—i.e., $G(\mathbf{b}^k) \geq G(\mathbf{b}^{k-1})$ and $PW_i(\mathbf{b}^k) \subseteq PW_i(\mathbf{b}^{k-1}), \forall 1 \leq k \leq t-1, i \in V$.

Case 1: Assume first that $G(\mathbf{b}^t) < G(\mathbf{b}^{t-1})$. This implies that $s_t(w_t) < s_{t-1}(w_{t-1})$, and hence: $s_t(w_{t-1}) \leq s_t(w_t) < s_{t-1}(w_{t-1})$, i.e., the score of the winner at step $t-1$ has decreased. By the unit upgrade policy, this means that at step $t-1$ some voter (say, i) ranks candidate w_{t-1} right above candidate w_t, and he swaps them at step t. Since i prefers w_t to w_{t-1} (as he is improving at step t), there was a step $k < t$ at which voter i made w_{t-1} a winner (as the process starts from the truthful state). That is, w_{t-1} was the most preferable candidate among potential winners of i at step $k-1$, and hence $w_{t-1} \in PW_i(\mathbf{b}^{k-1})$ but $w_t \notin PW_i(\mathbf{b}^{k-1})$. However, $w_t \in PW_i(\mathbf{b}^{t-1})$, in contradiction to the set monotonicity until step $t-1$.

Case 2: Suppose now that $G(\mathbf{b}^t) \geq G(\mathbf{b}^{t-1})$, but $PW_i(\mathbf{b}^t) \not\subseteq PW_i(\mathbf{b}^{t-1})$ for some $i \in V$. Let $c \in C$ and $i \in V$ such that $c \in PW_i(\mathbf{b}^t) \setminus PW_i(\mathbf{b}^{t-1})$. Since $G(\mathbf{b}^t) \geq G(\mathbf{b}^{t-1})$, and hence, $s_t(w_t) \geq s_{t-1}(w_{t-1})$, this is only possible if candidate c was in the top position in the ballot of voter i at step $t-1$ and is not in the top position in his ballot at step t. That is, c must have lost a point at step t: $s_t(c) = s_{t-1}(c) - 1$. Now, this means that the candidate c cannot belong to the set of potential winners of any of the voters at step t, a contradiction.

Finally, since by the unit upgrade policy the number of consequtive moves that a voter can make in favour of a particular candidate is naturally bounded by $m-1$, the procedure ($\mathcal{F}$, UU) satisfies C2(b). □

We note that both reducing the class of PSRs to the unit gap rules and the further restriction of the upgrade policy to allow only unit upgrades are necessary for each of the monotonicity conditions to hold. The following Example shows that both C1 and C2 may not hold for (Borda, U).

Example 4. There are 5 candidates $\{a, b, c, d, e\}$ and 13 voters $\{1, 2, 3, \ldots, 13\}$, but only the first three will be involved in the improvement process. The voting rule is Borda with the tie-breaking order $a \succ b \succ c \succ d \succ e$, and the improvement dynamic is upgrade.

Let agents 4–8 vote (sincerely) $cabde$, agents 9–11 vote $bcaed$, agent 12 vote $cabed$ and agent 13 $baced$. Thus, from these voters the candidates get the initial scores of $(27, 28, 35, 5, 5)$. Consider the following improvement path by the voters 1–3, voting truthfully at the first step:

$$\begin{aligned} & (abdec, deabc, debac)\{c\} \xrightarrow{1} (badec, deabc, debac)\{b\} \\ \xrightarrow{2} \quad & (badec, adebc, debac)\{a\} \xrightarrow{3} (badec, adebc, bdeac)\{b\} \\ \xrightarrow{1} \quad & (badec, adebc, bdeac)\{a\} \end{aligned}$$

Note that the score of the winner at the first three steps is 35, it goes up to 37 at step 4 and down to 36 at step 5; hence, C1 does not hold. Consider now the sets of potential winners for voter 1: at the first two steps it contains candidates b and c, when voter b leaves the set at the 3rd step. Now, at the 4th step, candidate a, who never was present in $PW_1(\cdot)$, joins the set, thus violating condition C2.

Remark 2. *The previous example can be easily modified to show that both C1 and C2 can be violated under positional scoring rules with non-unit gap scores, even if the agents apply only unit upgrades. For instance, it is sufficient to consider ($\mathcal{F}$, UU) where $\mathcal{F}$ is given by a scoring vector* $(5, 4, 2, 1, 0)$ *with only one non-unit gap.*

5 Conclusions

We provided a framework for studying convergence properties of iterative voting procedures under restricted dynamics. We established two general sufficient conditions that guarantee convergence of such myopic improvements. We then identified several classes of voting rules, along with their corresponding iterative processes, for which at least one of these conditions hold. Our work puts under the same framework recent results, it generalizes some of them by relaxing some restrictive assumptions, and also provides further positive results for more families of rules and dynamics.

Besides gaining a better understanding of what makes an iterative voting procedure converge, it is also interesting and important to evaluate the quality of outcomes obtained by iterative voting procedures. Which of the restricted dynamics can guarantee convergence to a Nash equilibrium and under what conditions? Analyzing the Dynamic Price of Anarchy and analogous measures for the quality of outcomes of such iterative procedures, along the lines of [3], is certainly a topic worth pursuing.

6 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. It is also co-financed by the UK Research Council for project ORCHID, grant EP/I011587/1; and by the grant RFFI 14-01-00156-a.

References

[1] S. Airiau and U. Endriss. Iterated majority voting. In *Algorithmic Decision Theory*, volume 5783 of *LNCS*, pages 38–49. Springer, 2009.

[2] S. J. Brams and P. C. Fishburn. *Approval Voting*. Birkhäuser, 1983.

[3] S. Branzei, I. Caragiannis, J. Morgenstern, and A. D. Procaccia. How bad is selfish voting? In *AAAI*, 2013.

[4] Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *ACM EC*, pages 347–356, 2010.

[5] B. Dutta and J.-F. Laslier. Costless honesty in voting. in 10th International Meeting of the Society for Social Choice and Welfare, Moscow, 2010.

[6] R. Farquharson. *Theory of Voting*. Yale University Press, 1969.

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

[8] U. Grandi, A. Loreggia, F. Rossi, K. B. Venable, and T. Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and borda score. In *ADT13*, 2013.

[9] J.-J. Laffont. Incentives and the allocation of public goods. In *Handbook of Public Economics*, volume 2, chapter 10, pages 537–569. Elsevier, 1987.

[10] O. Lev and J. S. Rosenschein. Convergence of iterative voting. In *AAMAS*, volume 2, pages 611–618, 2012.

[11] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria of plurality voting. In *AAAI*, pages 823–828, 2010.

[12] D. Monderer and L. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

[13] S. Obraztsova, E. Markakis, and D. R. M. Thompson. Plurality voting with truth-biased agents. In *SAGT*, 2013.

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

[15] R. Reyhani and M. Wilson. Best reply dynamics for scoring rules. In *ECAI*, volume 2, pages 672–677, 2012.

[16] M. A. 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, 1975.

[17] M. R. Sertel and M. Sanver. Strong equilibrium outcomes of voting games are the generalized Condorcet winners. *Social Choice and Welfare*, 22:331–347, 2004.

[18] D. R. M. Thompson, O. Lev, K. Leyton-Brown, and J. S. Rosenschein. Empirical aspects of plurality election equilibria. In *AAMAS*, 2013.

Svetlana Obraztsova
National Technical University of Athens
Athens, Greece
Email: svetlana.obraztsova@gmail.com

Evangelos Markakis
Athens University of Economics and Business
Athens, Greece
Email: markakis@gmail.com

Maria Polukarov
University of Southampton
Southampton, United Kingdom
Email: mp3@ecs.soton.ac.uk

Zinovi Rabinovich
Mobileye Vision Technologies Ltd.
Jerusalem, Israel
Email: zr@zinovi.net

Nicholas R. Jennings
University of Southampton
Southampton, United Kingdom
Email: nrj@ecs.soton.ac.uk