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

Svetlana Obraztsova[1] svetlana.obraztsova@gmail.com
Evangelos Markakis[2] markakis@gmail.com
Maria Polukarov[3] mp3@ecs.soton.ac.uk
Zinovi Rabinovich[4] zr@zinovi.net
Nicholas R. Jennings[3] nrj@ecs.soton.ac.uk

[1]National Technical University of Athens, Greece; Tel-Aviv University, Israel
[2]Athens University of Economics and Business, Greece
[3]University of Southampton, United Kingdom
[4]Mobileye Vision Technologies Ltd., Israel

Abstract

We study convergence properties of iterative voting procedures. Such procedures are 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 is already known that if the iteration dynamics (the manner in which voters are allowed to modify their votes) are unrestricted, then the voting process may not converge. For most common voting rules this may be observed even under the best response dynamics limitation. 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 the dynamics based on iterative myopic improvements, each of which is sufficient for convergence. We then identify several classes of voting rules (including Positional Scoring Rules, Maximin, Copeland and Bucklin), along with their corresponding iterative processes, for which at least one of these conditions hold.

Introduction

Voting mechanisms constitute a popular tool for preference aggregation and collective decision making in multi-agent systems that involve entities with possibly diverse preferences. The major concern, however, with voting as a decision-making process, is that voters may misreport their real preferences in order to favour certain candidates. Indeed, strategic behaviour is inherent in most voting rules, as the famous Gibbard-Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975) asserts. It is then natural to resort to game-theoretic concepts and tools in order to model voting behaviour and assess the outcome of a voting process.

Motivated by web services such as Doodle or Survey Monkey, we investigate strategic behaviour in *iterative voting processes*. Specifically, following (Meir et al. 2010), we consider iterative procedures, where agents start from some initial (most commonly, the truthful) voting configuration, and subsequently make myopic improvements by changing their vote. Iterative voting has recently received significant attention in the literature due to, in part, its potential to provide good predictions for the final outcome of a vote. For instance, an iterative process eliminates low quality Nash equilibria that may arise otherwise, in a one-shot voting game. Also, iterative voting can model the electorate response to poll data, see e.g., (Reijngoud and Endriss 2012).

Several results have 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 agents are allowed to make arbitrary moves (or just play only best responses), then convergence of such processes is not guaranteed, see e.g., (Lev and Rosenschein 2012). 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 make greedy local moves. Voters may also desire a consensus, reached by the smallest possible changes to their ballot, rather than a plain manipulation. Some restrictions on the dynamics of voting processes, as an attempt to model such considerations, were studied recently by (Reijngoud and Endriss 2012) and (Grandi et al. 2013). However, the family of restricted moves that enforces convergence has not been yet characterised, and this is one of the major open directions in the field of iterative voting.

Against this background, in this work we consider iterative voting procedures, specified by a voting rule along with a dynamic process, i.e., a specification of improvement steps that can be made by the voters. We focus on single-winner elections, with lexicographic tie-breaking. Our contribution is twofold. On the conceptual level, we provide two general conditions on these processes, each of which is sufficient for convergence. The first one is based on potential function arguments, whereas the second concerns the monotonicity of a set sequence defined along an improvement path. These conditions provide a unifying framework to study the behaviour of such dynamics. We then utilise this framework to identify dynamic processes that converge for several classes of voting rules (including Positional Scoring Rules, Maximin, and Bucklin). Especially for Bucklin, this is essentially the first positive result regarding convergence of an iterative voting procedure. At the same time, we also generalise recent results and relax a number of restrictive assumptions made in previous research. Finally, our conditions enable us to identify maximal sets of convergent dynamics (to an extent, as explained in the later sections).

Related Work

Earlier work on iterative voting processes is well summarised in (Laffont 1987), and concerns dynamics for deciding on allocations of public goods. The study of iterative voting in the more recent AI literature was initiated by (Meir et al. 2010), who focused on improvement dynamics under the Plurality rule, and provided both positive and negative results, depending on the initial voting profile, the tie-breaking rule and the improvement steps allowed (better replies or best responses). The follow up work (Lev and Rosenschein 2012) showed that for other voting rules, in fact, for most popular ones, it is often the case that convergence of best responses cannot be guaranteed. More results along this direction were established in (Reyhani and Wilson 2012), who also improved on the convergence bounds of (Meir et al. 2010). In the same spirit, (Kukushkin 2011) studied the existence of potential functions in voting games, as a way to prove convergence, and showed that the only voting rule admitting a potential function is the dictatorial rule. An analysis in terms of the quality of equilibria reachable by such processes was given in (Branzei et al. 2013). Finally, some different types of voting processes have been studied in (Airiau and Endriss 2009), where each agent is allowed to propose a change in the current state and then a vote is held for its acceptance, and more recently in (Meir, Lev, and Rosenschein 2014), where voters may have uncertainty about the current state and best responses are defined in terms of local dominance.

The works most closely related to ours are those of (Reijngoud and Endriss 2012) and (Grandi et al. 2013). Both of these consider procedures where voters may not play a best response but instead move according to certain restricted dynamics. Three types of processes have been considered (defined in the next section), and convergence results were established for some families of voting rules. As we show in the following sections, our framework incorporates these positive results and relaxes some of the limitations on the allowed moves. In particular, it demonstrates that our proposed restrictions on voting dynamics are weaker. As a result, a greater variability in voting behaviours can be allowed, while maintaining the stability of the iterative process, i.e., preserving convergence properties.

Finally, there have been other works applying game-theoretic concepts and tools to voting, starting with (Farquharson 1969). More recent research along this line has focused either on studying stronger equilibrium concepts (Sertel and Sanver 2004) or on different models of voting behaviour such as voting with abstentions (Desmedt and Elkind 2010) or truth-biased voting (Meir et al. 2010; Thompson et al. 2013; Dutta and Laslier 2010; Obraztsova et al. 2013). Here we concentrate on the standard voting behaviour model, and do not consider these latter extensions.

Preliminaries

We first recall some of the most common voting rules, and define the setting of iterative voting based on myopic improvement moves by single voters.

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*, $\succ_i \in \mathcal{L}(C)$. A *profile* $\mathbf{b} = (b_1, \dots, b_n) \in \mathcal{L}(C)^n$ is a vector of votes, one for each agent. 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. Specifically, we assume ties are broken according to lexicographic tie-breaking—i.e., 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 $s_1 \geq s_2 \geq \dots \geq s_m$. 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 $(1, 0, 0, \dots, 0)$, Veto with $(1, 1, \dots, 1, 0)$, Borda with $(m-1, m-2, ..., 0)$ and k-approval with $(1, ..., 1, 0, ..., 0)$, i.e., k 1's, followed by 0's.
- *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 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 what follows, 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 $c \in C$, his 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).

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* 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, for which $\mathbf{b}^k = (b_i', \mathbf{b}^{k-1}_{-i})$ for some $b_i' \neq b_i^{k-1}$ in $\mathcal{L}(C)$. It is an *improvement path* if for all $k \geq 1$, the move made by the unique deviator at step k, is an improvement move. 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, $\mathbf{b}^0 = (\succ_1, \ldots, \succ_n)$. 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 the natural restriction of *best response* dynamics is usually made: the deviating voter is assumed to make the best possible move at each step. While best responses always converge for Plurality and Veto with linear tie-breaking (Meir et al. 2010; Lev and Rosenschein 2012), they may cycle under other rules, such as Copeland (Grandi et al. 2013), Borda, and k-approval (Lev and Rosenschein 2012; Reyhani and Wilson 2012).

Restricted dynamics

In such settings, convergence can be achieved by restricting the sets of available improvement moves even further. Such restrictions may potentially arise due to uncertainty or limitations of the voters' computational power. To this end, the following dynamics have been previously considered:

- *Second Chance* (SC)[1] (Grandi et al. 2013): 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* (Brams and Fishburn 1983; Reijngoud and Endriss 2012): 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] (Grandi et al. 2013): 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.

[1]Also referred to as M1 and M2 in (Grandi et al. 2013).

In later sections, we extend some of the above policies and propose alternative voting dynamics. We use the term *iterative voting procedure* $(\mathcal{F}, D)$ to define the process based on the improvement dynamics D under the voting rule $\mathcal{F}$. In this work, we focus on dynamics where the deviator always chooses his best possible action among the allowed ones. We say that a voting procedure *converges* if every improvement path that contains moves allowed by D is finite under $\mathcal{F}$.

Two sufficient conditions for convergence

We present two conditions on iterative voting procedures, each of which guarantees convergence. These conditions are powerful enough to incorporate all convergence results for restricted dynamics in the literature. Moreover, they provide a general framework that lets us identify maximal ranges of converging processes, as we exemplify further.

Function monotonicity

The first condition is based on the potential argument (Monderer and Shapley 1996). 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, weak monotonicity of $G(\cdot)$ will suffice.

Condition 1 (Function monotonicity (FM))**.** *Given an iterative voting procedure* $(\mathcal{F}, D)$ *and a profile* $\mathbf{b}$*, let*

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

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

A weaker variant of this condition has also appeared in (Loreggia 2012). As Theorem 1 below states, Condition 1 guarantees convergence for voting procedures that admit consistent scoring functions, $s_{\mathcal{F}}$, satisfying FM.

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

Set monotonicity

The second condition follows the idea of (Reyhani and Wilson 2012). In their work, convergence of best response dynamics for Plurality was (re)proved by showing inclusion monotonicity for the sets of *potential winners* along an improvement path. These are 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.

In what follows, we slightly 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 $(\mathbf{b}^0 \to \mathbf{b}^1 \to \cdots)$, under a given $(\mathcal{F}, D)$.

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

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

where $b_i^{'}$ above is consistent with D. For $k \geq 1$, let

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

Condition 2 (Set monotonicity (SM)). *Let $(\mathcal{F}, D)$ be an iterative voting procedure. Then, for any improvement path $(\mathbf{b}^0 \to \mathbf{b}^1 \to \cdots)$ 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] so that for every $i \in V$ and $c \in C$, the maximum possible number of consecutive moves that can be made by i in favour of c, is bounded by q.[3]*

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

Non-equivalence between FM and SM

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 FM does not hold, but SM 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.

1	2	3	4	5	6	7	8	9
c_2	c_2	c_2	...	...	...	...	...	c_1
...	c_1	c_1	c_2	...	...	...	c_1	...
...	...	...	...	...	...	...	...	...
...	...	...	...	c_2	c_1	c_1	...	...
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

Figure 1: Voters' preferences in a game with SM but not FM

The voting rule we use is Bucklin, and the iterative improvement policy restricts voters to perform single swaps, where one candidate is moved one position up in the current ranking, if this candidate can become a new winner after such a move. Hence, at each step, the deviator moves his favourite alternative among the potential winners, by exactly *one* position up, and keeps the relative ranking of the others unchanged. We will revisit these dynamics in a later section, where we will establish positive results for Bucklin and a special class of PSRs.

Under the truthful profile $\mathbf{b}^0$, the winner is c_1, with a score of 6 (the Bucklin winning round, k_{min}, 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 FM 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 SM holds, in particular SM(b) holds. 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 $PW_5(\mathbf{b}^4) = \emptyset$. 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 $PW_7(\mathbf{b}^3) = \emptyset$, 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 favour of a certain candidate is at most $m-1$, hence SM(b) holds. □

Remark 1. *The game in Example 1 satisfies* SM(b)*, but not* SM(a)*. One can make slight adjustments, so that condition* FM *still does not hold but* SM(a) *does. We omit the details.*

Next, we give an example where FM holds but neither version of Condition 2 does.

Example 2. The construction is based on the (Borda, BU) procedure. We construct an instance where voter 1 has the preference order $c_1 \succ_1 c_2 \succ_1 c_3 \succ_1 \ldots$, and the ranking of voter 2 is $c_1 \succ_2 c_3 \succ_2 c_4 \succ_2 \ldots \succ_2 c_2$. 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 his score by s);
- candidate c_2 has $s-1$ points in $\mathbf{b}^0$;
- candidates $c_4, c_5, \ldots, c_m$ 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 SM 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 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, $PW_2(\mathbf{b}^1) = \{c_2, c_3, ..., c_m\}$, hence SM is violated.

Finally, FM holds by the results of (Grandi et al. 2013), falling under the first case of Proposition 3 below. □

[2] Usually, a simple function 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 sub-sequence of moves made by voter i only. They need not be consecutive in the whole improvement path.

Known procedures and relaxations

We now turn to explore the power of these monotonicity conditions. We start by showing that they capture the convergence results for previously studied restricted dynamics (Reijngoud and Endriss 2012; Grandi et al. 2013).

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

In fact SC converges under any voting rule. Proposition 3, in particular, demonstrates that the BU dynamic satisfies FM for most common voting rules. However, the FM condition is far more general. To see this, we next observe that BU can be significantly relaxed (i.e., allow a greater range of voter behaviours) within the FM bounds, thus preserving the convergence of the iterative voting process. In particular, consider the following relaxations:

- BU-1: As long as the winner changes to a more favourable candidate, the deviating voter can freely shuffle among themselves all the candidates ranked above c. The same applies to all the candidates ranked below c (again, among themselves). He is not allowed to change the ranking of c.
- BU-2: As before, the deviating voter is restricted to keep the rank of c unchanged, but can now shuffle the remaining candidates absolutely freely to effect the win by some of his more favourable candidates.

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

(Maximal) monotone dynamics under PSRs, Maximin and Bucklin

This section contains our main technical results. We identify several iterative voting procedures, for which either FM or SM hold. Unlike our previous discussion on BU-1 and BU-2, in what follows we will not a priori require that the current winner keeps his absolute position in the deviator's ballot. Indeed, this restriction may be too stringent, and we will seek more natural dynamics. However, as the proof of Theorem 4 suggests, BU-1 and BU-2 are *maximal* for Copeland/Maximin and PSRs, respectively, in the sense that removing any of their restrictions would yield violation of FM. Hence, to preserve monotonicity, we impose different types of limitations.

Our first example is the Maximin rule with a natural improvement dynamic, termed *Upgrade*, where the deviator moves a new winner to a higher position, keeping the relative ranking of the remaining candidates unchanged. As we will show, when the requirement of BU-1 and BU-2 regarding the current winner is removed, restricting the voters to not upgrade any candidate other than the new winner, is necessary for convergence.

We then move to a subclass of integer PSRs, termed *unit gap* scoring rules, where the difference in any two consecutive scores s_j, s_{j+1} is bounded by 1 (this class contains all common PSRs such as Plurality, Veto, k-approval and Borda). We show that FM and SM hold under the iterative process, called *Unit Upgrade*, where a new winner is moved by exactly one position higher in the ballot of the deviator. Importantly, without imposing the restriction of fixing the position of the current winner, this property is implied by the Unit Upgrade policy.

Finally, we demonstrate that the Unit Upgrade dynamics converges for the Bucklin rule, satisfying the set monotonicity condition. This is a particularly interesting result, as it shows the first iterative process that converges for Bucklin (after the SC dynamics that trivially converges for all rules).

Maximin with Upgrade

Consider the following policy for improvement moves:

- *Upgrade* (U): at each step, the deviator moves his favourite 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 other candidates unchanged. The upgraded candidate is the new winner.

Theorem 5. *The iterative procedure (Maximin, U) satisfies* FM *and both* SM(a) *and* SM(b).

We first demonstrate the following 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 declare to prefer c over x in 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 FM or SM does 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 $1 \leq k \leq t-1, i \in V$.

Case 1: 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: $s_t(w_{t-1}) < s_{t-1}(w_{t-1})$. If this was not the case, then by the definition of G, w_{t-1} should have 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 w_t from under w_{t-1}, above w_{t-1} in his ballot. Since i prefers w_t to w_{t-1}, 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 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 $\mathbf{b}^{k-1}$, 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 $c \in PW_i(\mathbf{b}^t) \setminus PW_i(\mathbf{b}^{t-1})$. First assume that the set of c's toughest opponents decreased at step t: $TO_c(\mathbf{b}^{t-1}) \not\subseteq TO_c(\mathbf{b}^t)$, 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 $s_t(c) \leq s_{t-1}(c)$. In fact, equality is not possible. To see this, note that c is not the winner at $\mathbf{b}^t$, 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 one of the toughest opponents of c at $\mathbf{b}^t$, is if some other tough 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), we have $s_t(c) \leq s_t(w_t) - 1$. We claim then that c cannot belong to the set of potential winners of any voter at $\mathbf{b}^t$. This is based on similar arguments as in Case 1, showing that otherwise monotonicity would be violated at an earlier stage. Hence, we reach a contradiction to the fact that $c \in PW_i(\mathbf{b}^t)$. 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 the toughest opponents of c are ranked above c. Since $c \notin PW_i(\mathbf{b}^{t-1})$, this was not the case at $\mathbf{b}^{t-1}$. It is easy to see then that one of the toughest opponents of c at step $t-1$ was ranked below c in voter i's ballot. Since $TO_c(\mathbf{b}^{t-1}) \subseteq TO_c(\mathbf{b}^t)$, this candidate was moved above c at $\mathbf{b}^t$, and the score of c decreased by 1. But this excludes c from the set of potential winners for all the voters, again a contradiction.

Finally, we show that both SM(a) and SM(b) hold. For SM(b), we trivially have $q \leq m-1$. For SM(a), note that at the first step, all the voters who cannot make the truthful winner, w_0, win again (i.e., those who rank w_0 above at least one of his 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, loses the previous winner, w_{k-1}, from his set of potential winners. Hence, SM(a) holds. □

Next, we argue that the requirement of upgrading (i.e., moving up) *only* the winning candidate is necessary for convergence under many rules, when the absolute position of the current winner can change in the ballot. For instance, cycles have been shown for Copeland (Grandi et al. 2013), k-approval and Borda (Lev and Rosenschein 2012; Reyhani and Wilson 2012), even when lexicographic tie-breaking is used. For Maximin, (Lev and Rosenschein 2012) 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 $d \succ b \succ c \succ a$ for tie-breaking. At first step, the agents vote sincerely, seen below, and d wins. As voter 1 prefers b over d, he deviates from his true preference order $abdc$ and votes $abcd$, which makes b win (note that this is a best response for voter 1, and it involves moving a non-winning candidate). Next, voter 2 deviates to make c a winner, and so on. We describe the improvement path below, with a cycle starting at the fourth step:

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

Unit gap scoring rules with Unit Upgrade

Let $\mathcal{F}$ be a PSR with an integer 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 further restrict the Upgrade policy:

- *Unit Upgrade* (UU): at each step, the deviator moves his favourite alternative among the potential winners, 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* SM(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. E.g., for (Borda, U), both FM and SM may not hold. The example can be further modified to show that FM and SM can be violated under positional scoring rules with non-unit gap scores, even if the agents apply only unit upgrades.

The Bucklin rule with Unit Upgrade

Finally, as Example 1 may also suggest, we show that UU converges for Bucklin.

Theorem 7. *The iterative procedure (Bucklin, UU) is set monotone—in particular it satisfies SM(b).*

Theorem 7 demonstrates the power of our technique, which allowed us to prove convergence of the Bucklin rule for the first time. To our knowledge, there was no known reasonable iterative process converging under Bucklin, except the very restrictive SC dynamics that trivially terminates for all voting rules. Indeed, neither BU, nor its relaxed versions BU-1/BU-2 converge for Bucklin, as fixing the position of the current winner in the ballot of the deviator does not guarantee that his score also remains unchanged. However, our results still leave open the question of maximality of the UU dynamics (both under Bucklin and PSRs). Indeed, we observe that monotonicity breaks if we relax the UU restriction—that is, allow the voters to move a new winner higher by more than one position. However, this does not yet imply that a similar dynamics with larger allowed upgrades would not converge (our monotonicity conditions are only sufficient conditions for convergence). Hence it would be interesting to determine how far this policy could be moderated.

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 generalises some of them by relaxing their 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 important to evaluate the quality of outcomes obtained by such procedures. For example, under some instances, voting dynamics may converge to a profile, where voters may still wish to change their vote, but are not allowed to do so due to the policy restriction. Therefore, we need to understand which of the restricted iterative processes can guarantee convergence to a Nash equilibrium, and under what conditions. Analysing the Dynamic Price of Anarchy and analogous measures for the quality of outcomes of iterative procedures, along the lines of (Branzei et al. 2013), would also be equally interesting and have an impact on the use of original motivating applications.

Acknowledgments

This work has been cofinanced by the 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. The authors also gratefully acknowledge funding from the UK Research Council for project ORCHID—grant EP/I011587/1, and from the COST Action IC1205 on Computational Social Choice.

References

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

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

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

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

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

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

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

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

Kukushkin, N. S. 2011. Acyclicity of improvements in finite game forms. *International Journal of Game Theory* 40(1):147–177.

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

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

Loreggia, A. 2012. Iterative voting and multi-mode control in preference aggregation. Master's thesis, University of Padova.

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

Meir, R.; Lev, O.; and Rosenschein, J. S. 2014. A local-dominance theory of voting equilibria. In *ACM Conference on Economics and Computation (EC 2014)*, 313–330.

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

Obraztsova, S.; Markakis, E.; and Thompson, D. R. M. 2013. Plurality voting with truth-biased agents. In *6th Symposium on Algorithmic Game Theory (SAGT 2013)*, 26–37.

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

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

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

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

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