Strong and Weak Acyclicity in Iterative Voting

Reshef Meir
Technion—Israel Institute of Technology
reshefm@ie.technion.ac.il

Abstract

We cast the various different models used for the analysis of iterative voting schemes into a general framework, consistent with the literature on acyclicity in games. More specifically, we classify convergence results based on the underlying assumptions on the agent scheduler (the order of players) and the action scheduler (which better-reply is played).

Our main technical result is proving that Plurality with randomized tie-breaking (which is not guaranteed to converge under arbitrary agent schedulers) is weakly-acyclic. I.e., from any initial state there is *some* path of better-replies to a Nash equilibrium. We thus show a separation between restricted-acyclicity and weak-acyclicity of game forms, thereby settling an open question from [18].

1 Introduction

A strand of papers on what is usually called "iterative voting," started roughly 6 years ago in a AAAI paper (that also featured in COMSOC 2010) by Meir, Polukarov, Rosenschein and Jennings [22]. In iterative voting, the voting rules and the voters' preferences are fixed, but voters are strategic and are allowed to change their vote one at a time after observing the interim outcome. The main questions in the field are regarding which voting rules guarantee convergence of the iterative process to a Nash equilibrium, and under what conditions (e.g.,[19, 31, 28]).

Long before that, researchers in economics and game theory since Cournot [8] had been developing a formal framework to study questions about acyclicity and convergence of local improvement dynamics in games [26, 25, 10, 1, 18, 2]. However, these two lines of work remained largely detached from one another. Since the general game-theoretic framework is much more mature and widely used, it is in the interest of the COMSOC community to bridge this gap, which is the main purpose of this work.

Intuitively put, strong-acyclicity means that the game will converge regardless of the order of players/voters and how they select their action (as long as they are improving their utility), i.e. that there are no cycles of better-replies whatsoever; Weak-acyclicity means that while cycles may occur, from any initial state (voting profile) there is at least one path of better-replies that leads to a Nash equilibrium; Restricted-acyclicity is a middle ground, requiring convergence for any order of players (agent scheduler), but allowing the action scheduler to restrict the way they choose among several available replies (e.g., only allowing best-replies). Most relevant to us is the work of Kukushkin [16, 17, 18], who studied general characterizations of game forms that guarantee various notions of acyclicity

Papers on iterative voting typically focus on a specific voting rule (game form), and study its convergence properties. Most results, both positive and negative, are about restricted acyclicity (under various notions of restriction) and include the work of Meir et al. [22], Reyhani and Wilson [31], and Lev and Rosenschein [19].

More recent work on iterative voting deals with voters who are uncertain, truth-biased, lazy-biased, bounded-rational, non-myopic, or apply some other restrictions and/or heuristics that diverge from the standard notion of better-reply in games [30, 12, 13, 27, 23, 29, 28, 24]. Although the framework is suitable for studying such iterative dynamics as well, this paper deals exclusively with myopic better-reply dynamics.[1]

Building on the formalism of Kukushkin [18] for strong/ restricted/ weak-acyclicity of game forms, we re-interpret most of the known results on convergence of better- and best-reply in voting games.

[1] We do consider however two standard ways to handle ties that slightly relax the better-reply definition. See Section 4.4.

f_1	a	b	c
a	a	a	a
b	b	b	b
c	c	c	c

f_2	a	b	c
a	a	a	a
b	a	b	b
c	a	b	c

f_3	x	y
a	a	b
b	b	c
c	c	a

f_4	x	y	z	w
a	ax	ay	az	aw
b	bx	by	bz	bw
c	cx	cy	cz	cw

Figure 1: Four examples of game forms with two agents. f_1 is a dictatorial game form with 3 candidates (the row agent is the dictator). f_2 is the Plurality voting rule with 3 candidates and lexicographic tie-breaking. f_3 and f_4 are non-standard game forms. In f_3, $A_1 = C = \{a, b, c\}, A_2 = \{x, y\}$. Note that f_4 is completely general (there are 3×4 possible outcomes in C, one for each voting profile) and can represent any 3-by-4 game.

1.1 Contribution and structure

The paper unfolds as follows. In Section 2, we define the iterative voting model within the more general framework of acyclic and weakly acyclic game forms. In Section 3 we consider strong acyclicity, and settle an open question regarding the existence of acyclic non-separable game forms. Section 4 covers most known results from the literature on convergence under variations of the Plurality rule, and fills is some small gaps. Our main technical contribution is in Section 5, where we show a strict separation between restricted acyclicity and weak acyclicity, thereby settling another open question. In Section 6 we review known convergence results for other voting rules, and conclude in Section 7.

2 Preliminaries

We build upon the basic notations and definitions of Meir et al. [22] and Kukushkin [18]. We usually denote sets by uppercase letters and vectors by bold letters, e.g., $\mathbf{a} = (a_1, \dots, a_n)$.

2.1 Voting rules and game forms

There is a set C of m alternatives (or *candidates*), and a set N of n strategic agents, or *voters*. A game form (or *voting rule*) f allows each agent $i \in N$ to submit her preferences over the alternatives by selecting an action from a set A_i. Thus the input to f is a vector $\mathbf{a} = (a_1, \dots, a_n)$ called an *action profile*. We also refer to a_i as the *vote* of agent i in profile $\mathbf{a}$. Then, f chooses a winning alternative—i.e., it is a function $f : \mathcal{A} \to C$, where $\mathcal{A} = \times_{i \in N} A_i$. See Fig. 1 for examples.

A voting rule f is *standard* if $A_i = A$ for all i, and A is either $\pi(C)$ (the set of permutations over C) or a coarsening of $\pi(C)$. Thus most common voting rules except Approval are standard. Mixed strategies are not allowed. The definitions in this section apply to all voting rules unless stated otherwise.

Plurality In the Plurality voting rule we have that $A = C$, and the winner is the candidate with the most votes. We allow for a broader set of "Plurality game forms" by considering both weighted and fixed voters, and varying the tie-breaking method. Each of the strategic voters $i \in N$ has an integer weight $w_i \in \mathbb{N}$. In addition, there are $\hat{n}$ "fixed voters" who do not play strategically or change their vote. The vector $\hat{\mathbf{s}} \in \mathbb{N}^m$ (called "initial score vector") specifies the number of fixed votes for each candidate. Weights and initial scores are part of the game form.[2] These extensions also apply to other positional scoring rules.

The *final score* of c for a given profile $\mathbf{a} \in A^n$ in the game form $f_{\mathbf{w},\hat{\mathbf{s}}}$ is the total weight of voters that vote c. We denote the final score vector by $\mathbf{s}_{\hat{\mathbf{s}},\mathbf{w},\mathbf{a}}$ (often just $\mathbf{s}_{\mathbf{a}}$ or $\mathbf{s}$ when the other parameters are clear from the context), where $s(c) = \hat{s}(c) + \sum_{i \in N: a_i = c} w_i$.

[2] All of our results still hold if there are no fixed voters, but allowing fixed voters enables the introduction of simpler examples. For further discussion on fixed voters see [9].

Thus the Plurality rule is defined as $f^P_{\hat{\mathbf{s}},\mathbf{w}}(\mathbf{a}) = \text{argmax}_{c\in C}\, s_{\hat{\mathbf{s}},\mathbf{w},\mathbf{a}}(c)$, breaking ties according to some specified method (lexicographically unless specified otherwise). As with $\mathbf{s}$, we omit the scripts $\mathbf{w}$ and $\hat{\mathbf{s}}$ when they are clear from the context.

2.2 Incentives

Games are attained by adding either cardinal or ordinal utility to a game form. The linear order relation $Q_i \in \pi(C)$ reflects the preferences of agent i. That is, i prefers c over c' (denoted $c \succ_i c'$) if $(c, c') \in Q_i$. The vector containing the preferences of all n agents is called a *preference profile*, and is denoted by $\mathbf{Q} = (Q_1, \dots, Q_n)$. The game form f, coupled with a preference profile $\mathbf{Q}$, defines an ordinal utility normal form game $G = \langle f, \mathbf{Q}\rangle$ with n agents, where agent i prefers outcome $f(\mathbf{a})$ over outcome $f(\mathbf{a}')$ if $f(\mathbf{a}) \succ_i f(\mathbf{a}')$.

Manipulation and Stability Having defined a normal form game, we can now apply standard solution concepts. Let $G = \langle f, \mathbf{Q}\rangle$ be a voting game, and let $\mathbf{a} = (\mathbf{a}_{-i}, a_i)$ be a joint action in G.

We denote by $\mathbf{a} \xrightarrow{i} \mathbf{a}'$ an *individual improvement step*, if (1) $\mathbf{a}, \mathbf{a}'$ differ only by the action of player i; and (2) $f(a_{-i}, a'_i) \succ_i f(a_{-i}, a_i)$. We sometimes omit the actions of the other voters $\mathbf{a}_{-i}$ when they are clear from the context, only writing $a_i \xrightarrow{i} a'_i$. We denote by $I_i(\mathbf{a}) \subseteq A$ the set of actions a'_i s.t. $a_i \xrightarrow{i} a'_i$ is an improvement step of agent i in $\mathbf{a}$, and $I(\mathbf{a}) = \bigcup_{i\in N} \bigcup_{a'_i \in I_i(\mathbf{a})} (\mathbf{a}_{-i}, a'_i)$. $\mathbf{a} \xrightarrow{i} a'_i$ is called a *best reply* if a'_i is i's most preferred candidate in $I_i(\mathbf{a})$.

A joint action $\mathbf{a}$ is a (pure) *Nash equilibrium* (NE) in G if $I(\mathbf{a}) = \emptyset$. That is, no agent can gain by changing his vote, provided that others keep their strategies unchanged. A priori, a game with pure strategies does not have to admit any NE.

Now, observe that when f is a standard voting rule the preference profile $\mathbf{Q}$ induces a special joint action $\mathbf{a}^* = \mathbf{a}^*(\mathbf{Q})$, termed the *truthful state*, such that $\mathbf{a}^*(\mathbf{Q}) = (a^*_1, \dots, a^*_n)$, where $a^*_i \succ_i c$ for all $c \neq a^*_i$. We refer to $f(\mathbf{a}^*)$ as the *truthful outcome* of the $\langle f, \mathbf{Q}\rangle$. If i has an improvement step in the truthful state, then this is a *manipulation*.

2.3 Iterative Games

We consider natural *dynamics* in iterative games. Assume that agents start by announcing some initial profile $\mathbf{a}^0$, and then proceed as follows: at each step t a single agent i may change his vote to $a'_i \in I_i(\mathbf{a}^{t-1})$, resulting in a new state (joint action) $\mathbf{a}^t = (\mathbf{a}^{t-1}_{-i}, a'_i)$. The process ends when no agent has objections, and the outcome is set by the last state. Such a restriction makes sense in many computerized environments, where voters can log-in and change their vote at any time.

Local improvement graphs and schedulers Any game G induces a directed graph whose vertices are all action profiles (states) A^n, and edges are all local improvement steps [35, 1]. The pure Nash equilibria of G are all states with no outgoing edges. Since a state may have multiple outgoing edges ($|I(\mathbf{a})| > 1$), we need to specify which one is selected in a given play.

A *scheduler* ϕ selects which edge is followed at state $\mathbf{a}$ at any step of the game [2]. The scheduler can be decomposed into two parts, namely selecting an agent i to play (agent scheduler ϕ^N), and selecting an action in $I_i(\mathbf{a})$ (action scheduler ϕ^A), where $\phi = (\phi^N, \phi^A)$.

Convergence and acyclicity Given a voting game G, an initial voting profile $\mathbf{a}^0$ and a scheduler ϕ, we get a unique (possibly infinite) path of steps.[3] Also, it is immediate to see that the path is finite if and only if it reaches a Nash equilibrium (which is the last state in the path). We say that the triple $\langle G, \mathbf{a}^0, \phi\rangle$ *converges* if the induced path is finite.

Following [26, 25], a game G has the *finite individual improvement property* (we say that G *is* FIP), if $\langle G, \mathbf{a}^0, \phi\rangle$ converges for *any* $\mathbf{a}^0$ and scheduler ϕ. Games that are FIP are also known as *acyclic games* and as *generalized ordinal potential games* [26].

[3] By "step" we mean an individual improvement step, unless specified otherwise.

It is quite easy to see that not all Plurality games are FIP (see examples in Section 4.1). However, there are alternative, weaker notions of acyclicity and convergence.

- A game G is *weakly-FIP* if there is *some* scheduler ϕ such that $\langle G, \mathbf{a}^0, \phi\rangle$ converges for any $\mathbf{a}^0$. Such games are known as *weakly acyclic*, or as ϕ-potential games [2].
- A game G is *restricted-FIP* if there is *some action scheduler* ϕ^A such that $\langle G, \mathbf{a}^0, (\phi^N, \phi^A)\rangle$ converges for any $\mathbf{a}^0$ and ϕ^N [18]. We term such games as *order-free acyclic*.

Intuitively, restricted FIP means that there is some restriction players can adopt s.t. convergence is guaranteed regardless of the order in which they play. Kukushkin identifies a particular restriction of interest, namely restriction to best-reply improvements, and defines the *finite best-reply property* (FBRP) and its weak and restricted analogs. We emphasize that an action scheduler *must* select an action in $I_i(\mathbf{a})$, if one exists. Thus restricted dynamics that may disallow all available actions (as in [12, 13]) do not fall under the definition of restricted-FIP (but can be considered as separate dynamics).

We identify a different restriction, namely *direct reply* that is well defined under the Plurality rule, where $A = C$. Formally, a step $\mathbf{a} \xrightarrow{i} \mathbf{a}'$ is a direct reply if $f(\mathbf{a}') = a'_i$, i.e., if the i votes for the new winner. ϕ^A is direct if it always selects a direct reply. We get the following definitions for a Plurality game G:

- G is *FDRP* if $\langle G, \mathbf{a}^0, (\phi^N, \phi^A)\rangle$ converges for any $\mathbf{a}^0$, any ϕ^N, and any direct ϕ^A.
- G is *weakly-FDRP* if there is some direct ϕ such that $\langle G, \mathbf{a}^0, \phi\rangle$ converges for any $\mathbf{a}^0$.
- G is *restricted-FDRP* if there is some direct ϕ^A such that $\langle G, \mathbf{a}^0, (\phi^N, \phi^A)\rangle$ converges for any $\mathbf{a}^0$ and ϕ^N.
- FDBRP means that replies are both best and direct. Note that it is unique and thus cannot be further restricted.

Finally, a game form f has the X property (where X is any of the above versions of finite improvement) if $\langle f, \mathbf{Q}\rangle$ is X for any preference profile $\mathbf{Q}$. We have the following entailments, both for games and for game forms. The third row is only relevant for Plurality (and for Veto, where a direct reply is to veto the current winner).

FBRP				restricted-FBRP	$\Rightarrow$	weak-FBRP		
$\Uparrow$				$\Downarrow$		$\Downarrow$		
FIP	$\Rightarrow$	FDBRP	$\Rightarrow$	restricted-FIP	$\Rightarrow$	weak-FIP	$\Rightarrow$	pure Nash exists
$\Downarrow$				$\Uparrow$		$\Uparrow$		
FDRP				restricted-FDRP	$\Rightarrow$	weak-FDRP		

Kukushkin notes that there are no known examples of game forms that are weak-FIP, but not restricted-FIP. We settle this question later in Section 5.

Convergence from the truth We say that a game G is *FIP from state* $\mathbf{a}$ if $\langle G, \mathbf{a}, \phi\rangle$ converges for any ϕ. Clearly a game is FIP iff it is FIP from $\mathbf{a}$ for any $\mathbf{a} \in A^n$. The definitions for other all other notions of finite improvement properties are analogous.

We are particularly interested in convergence from the truthful state $\mathbf{a}^*$. This is since: a. it is rather plausible to assume that agents will start by voting truthfully, especially when not sure about others' preferences; and b. even with complete information, they may be inclined to start truthfully, as they can always later change their vote.

3 Strong Acyclicity

Kukushkin [18] shows that a voting rule guarantees the existence of ordinal potential, if and only if it is a dictatorship.[4] However this *does not* preclude the existence of voting rules with FIP (*generalized* ordinal potential). Indeed, Kukushkin provides a partial characterization of FIP games forms. For example, a rule where there is a linear order L over C, and the winner is the first candidate according to L that is top-ranked by at least one voter.

[4] A [generalized] ordinal potential is a function that strictly increases [if] iff some agent plays a better-reply.

Kukushkin conjectures that the set of FIP game forms is quite limited, and in particular may only contain "separable rules" where each voter specifies a single desired candidate (as in Plurality). This conjecture was proved for $n = 2$ in [4] (separable game forms are called "assignable" there).

Indeed, for most common voting rules (which are non-separable with the exception of Plurality and dictatorships), it is easy to find examples where some cycles occur. Thus one should focus on the weaker notions of convergence discussed in Section 1, which is what we do in the remainder of the paper.

4 Plurality

4.1 Lexicographic Tie-Breaking

Lemma 1. *Consider a game* $\left\langle f^{PL}_{\mathbf{w},\hat{\mathbf{s}}}, \mathbf{Q} \right\rangle$*. If there exists a better reply for a given agent* i *at state* $\mathbf{a}^{t-1}$*, then* i *has a direct best reply.*

The proof is trivial under lexicographic tie-breaking, by letting i vote for her most preferred candidate among all better replies. In this case the direct best reply is also unique.

One implication of the lemma is that it is justified and natural to restrict our discussion to direct replies and focus on FDRP, as w.l.o.g. a voter always has a direct reply that is at least as good as any other reply.

Unweighted Voters We start with the most simple setting, and state the main result from [22] using our more fine-grained terminology.

Theorem 2 (Theorem 3 in [22])**.** $f^{PL}_{\hat{\mathbf{s}}}$ *is FDRP. Moreover, any path of direct replies will converge after at most* $m^2 n^2$ *steps. In particular, Plurality is order-free acyclic.*

In fact, the theorem in [22] only refers to FDBRP, but the proof can be modified for our more general theorem (details omitted). The bound on the number of direct-best-reply steps was later improved to $O(mn)$ in [31, Theorem 5.4].

Proposition 3 (Prop. 4 in [22])**.** $f^{PL}_{\hat{\mathbf{s}}}$ *is not FBRP, even from the truthful state. There are: (a) a counterexample with two agents and an arbitrary initial state; (b) a counterexample with three agents and a truthful initial vote.*

Weighted Voters The next results still consider Plurality with lexicographic tie-breaking, but where voters have arbitrary integer weights.

Proposition 4 (Prop. 5 in [22])**.** $f^{PL}_{\hat{\mathbf{s}},\mathbf{w}}$ *is not restricted-FDRP, even from the truthful state.*

Example 4. The initial fixed score of candidates $\{a, b, c, d\}$ is $\hat{\mathbf{s}} = (0, 1, 2, 3)$. The weight of each voter $i \in \{1, 2, 3\}$ is i. The preference profile is as follows: $c \succ_1 d \succ_1 b \succ_1 a$, $b \succ_2 c \succ_2 a \succ_2 d$, and $a \succ_3 b \succ_3 c \succ_3 d$. We omit the rest of the proof. $\diamondsuit$

Theorem 5 (Theorem 6(a) in [22])**.** $f^{PL}_{\hat{\mathbf{s}},\mathbf{w}}$ *is FDRP for* $n = 2$*.*

Theorem 6 (Theorem 6(b) in [22])**.** $f^{PL}_{\hat{\mathbf{s}},\mathbf{w}}$ *is FIP from the truth for* $n = 2$*.*

It remains an open question whether there is any restriction on better replies that guarantees order-free acyclicity in weighted games, i.e. if $f^{PL}_{\mathbf{w}}$ is restricted-FIP for $n > 2$. However Prop. 4 shows that if such restricted dynamic exists, it must make use of indirect replies, which is rather unnatural. We thus conjecture that such restricted dynamics does not exist.

4.2 Arbitrary tie-breaking

Lev and Rosenschein [19] showed that for any positional scoring rule (including Plurality), we can assign some (deterministic) tie breaking rule, so that the resulting voting rule may contains cycles. For any positional scoring rule f_α with score vector α, denote by f^{LR}_α the same rule with the Lev-Rosenschein tie-breaking.

Proposition 7 (Theorem 1 in [19]). *f^{LR}_α is not FBRP for any α, even for $n = 2$, and even from the truth. In particular, Plurality with the Lev-Rosenschein tie-breaking (f^{PLR}) is not FBRP.*

In fact, a slight modification of their example (switching a and b in voter 2's preferences) yields the following:

Proposition 8. *f^{PLR} is not restricted-FIP, even for $n = 2$, and even from the truth.*

4.3 Randomized tie-breaking

Formally, the game form $f^{PR}_{\hat{\mathbf{s}},\mathbf{w}}$ maps any state $\mathbf{a} \in A^n$ to the set $\text{argmax}_{c\in C}\, s_{\hat{\mathbf{s}},\mathbf{w},\mathbf{a}}(c)$. Since under randomized tie-breaking there are multiple winners, let $W^t = f^{PR}(\mathbf{a}^t) \subseteq C$ denote the set of winners at time t.[5] We define a direct reply $a_i^{t-1} \xrightarrow{i} a_i^t$ as one where $a_i^t \in W^t$.

If ties are broken randomly, $\succ_i$ does *not* induce a complete order over outcomes. For instance, the order $a \succ_i b \succ_i c$ does not determine if i will prefer $\{b\}$ over $\{a, c\}$. However, we can naturally extend Q_i to a *partial preference order* over subsets. There are several standard extensions, using the following axioms:[6]

K (Kelly [14]): $(\forall a \in X, b \in W, a \succ_i b) \;\Rightarrow\; X \succ_i W$;

G (Gärdenfors [11]): $(\forall b \in W, a \succ_i b) \;\Rightarrow\; \{a\} \succ_i (\{a\} \cup W) \succ_i W$;

R (Responsiveness [32]): $a \succ_i b \iff \forall W \subseteq C \setminus \{a, b\},\ (\{a\} \cup W) \succ_i (\{b\} \cup W)$.

The axioms reflect various beliefs a rational voter may have on the tie-breaking procedure: the K axiom reflects no assumptions whatsoever; The G axiom is consistent with tie-breaking according to a fixed and unknown order; and G+R axioms are consistent with random tie-breaking with equal probabilities (see Lemma 11). In this section we assume both G+R axioms hold (G entails K), however our results do not depend on these interpretations, and we do not specify the voter's preferences in cases not covered by the above axioms.

Theorem 9 (Theorem 8 in [22]). *$f^{PR}_{\hat{\mathbf{s}}}$ is FBRP from the truth.*

Here, too, the result in Meir et al. was restricted to direct best replies, but can be easily extended once we prove that a restriction to direct replies is valid:

Lemma 10. *If there exists a better-reply in f^{PR} for agent i at state $\mathbf{a}^{t-1}$, then i has a direct best-reply.*

Proof. Suppose there is a better reply $a_i^{t-1} \xrightarrow{i} b$ at time $t-1$. As some best reply always exists, denote by b' an arbitrary best reply. Let $W = f^R(\mathbf{a}^{t-1}_{-i}, b')$, and let a' be the most preferred candidate of i in W. Then we argue that $a_i^{t-1} \xrightarrow{i} a'$ is a direct best reply of i (for the lexicographic case this follows immediately from $W = \{a'\}$ and $f^L(\mathbf{a}^{t-1}_{-i}, a') = W = \{a'\}$).

Assume towards a contradiction that $a_i^{t-1} \xrightarrow{i} b'$ is not a direct reply. Then it must be of type 3, meaning $b' \notin W$. By voting for $a' \in W$, we get that $f_R(\mathbf{a}^{t-1}_{-i}, a') = \{a'\}$, i.e., a' remains the unique winner. If $|W| = 1$ then we are done as in the lexicographic case. Otherwise we apply Axiom G with $W' = W \setminus \{a'\}$, and get that $a' \succ_i (\{a'\} \cup W') = W$. That is, $f_R(\mathbf{a}^{t-1}_{-i}, a') \succ_i W$, which is a contradiction to b' being a best reply. $\square$

[5]This is a slight abuse of the notation we introduce in the beginning, where we defined the set of possible outcomes of f to be C. Here we allow any $W \in 2^C \setminus \{\emptyset\}$ as a possible outcome.

[6]We thank an anonymous reviewer for the references.

Cardinal utilities A (cardinal) utility function is a mapping of candidates to real numbers $u : C \to \mathbb{R}$, where $u_i(c) \in \mathbb{R}$ is the utility of candidate c to agent i. We say that u is *consistent* with a preference relation $\succ_i$ if $u(c) > u(c') \Leftrightarrow c \succ_i c'$. The definition of cardinal utility naturally extends to multiple winners by setting $u_i(W) = \frac{1}{|W|}\sum_{c\in W} u_i(c)$ for any subset $W \subseteq C$.[7]

Lemma 11 ([22]). *Any utility function u which is consistent with preference order $\succ_i$, is also consistent with the partial order over subsets induced by Axiom G and Axiom R.*

Finally, in contrast to the lexicographic case, convergence is no longer guaranteed if agents start from an arbitrary profile of votes, or are allowed to use direct-replies that are not best-replies.

Proposition 12 (Prop. 9 in [22]). *f^{PR} is not restricted-FIP.*

Example 12. There are 4 candidates $\{a, b, c, x\}$ and 3 agents with utilities $u_1 = (7, 3, 0, 4)$, $u_2 = (0, 7, 3, 4)$ and $u_3 = (3, 0, 7, 4)$. In particular, the following preference relations hold: $a \succ_1 \{a, b\} \succ_1 x \succ_1 \{a, c\}$; $b \succ_2 \{b, c\} \succ_2 x \succ_2 \{a, b\}$; and $c \succ_3 \{a, c\} \succ_3 x \succ_3 \{b, c\}$.

Consider the initial state $\mathbf{a}_0 = (a, b, x)$ with $\mathbf{s}(\mathbf{a}_0) = (1, 1, 0, 1)$ and the outcome $\{a, b, x\}$. We have the following cycle where every step is the unique reply of the playing agent. ♢

Proposition 13. *f^{PR} is not FDRP even from the truth.*

Example 13. We take the game from Ex. 12, and add for each voter $i \in \{1, 2, 3\}$ a candidate d_i, s.t. $u_i(d_i) = 8, u_i(d_j) = j$ for $j \neq i$. We also add an initial score of 3 to each of the candidates $\{a, b, c, x\}$. Voter 3 moves first to $a^1_3 = x$, which is a direct reply. Then voters 1 and 2 move to their best replies a, b, respectively. Now the cycle continues as in Ex. 12. ♢

4.4 Stochastic Dominance and Local Dominance

Reyhani and Wilson also consider random tie-breaking, but under a more risk-averse model than that of Meir et al. [22]. Specifically, they assume that a voter will only perform a step that *stochastically dominates* (SD) the current winner. Preferences under SD hold both Axioms G and R, and *explicitly forbid* any step that is not implied by these axioms.

Theorem 14 (Theorem 5.7 in [31]). *Plurality with stochastic dominance tie-breaking is FDBRP.*

Since any SD step is also a better-reply under random tie breaking, any strong or restricted convergence result for the latter applies to the former, but not vice-versa.

An even more risk-averse approach than Stochastic Dominance, is Local Dominance. Suppose voters' preferences hold Axiom G, and forbid any step that is not implied by G.

This coincides with the special case of the local-dominance approach in [23, 24] when there is no uncertainty regarding the score (i.e. the only source of uncertainty is the tie-breaking). We get the following corollary:

Theorem 15 (Theorem 11 in the full version of [24]). *Plurality with Local-Dominance tie-breaking is FDRP.*

Note that since Axioms G+R include G, any LD step is also an SD step, so a restriction to LD can only eliminate cycles. Thus FBDRP follows from Theorem 14. We note that with either SD or LD tie-breaking there may be new stable states that are not Nash-equilibria. Even so, an analysis of Ex. 10(a) in [22] shows that all steps are entailed by Axiom G (and thus by Axioms G+R). Thus neither game form is FIP.

What if we assume that voters are even more risk-averse and only follow steps that are better-replies by Axiom K? Then it is easy to see that only moves to a more-preferred candidate can be better-replies (any move to or from a tie cannot follow from K and is thus forbidden), which means there are trivially no cycles.

[7] One interpretation is that we randomize the final winner from the set W, and hence the term randomized tie-breaking. For a thorough discussion of cardinal and ordinal utilities in normal form games, see [3].

5 Weak Acyclicity

Except for Plurality and Veto convergence is not guaranteed even under restrictions and from the truthful state. In contrast, simulations [13, 23, 15] show that iterative voting almost always converges even in situations where this is not guaranteed by theory. We believe that weak acyclicity is an important part of the explanation to this gap.

Fabrikant et al. [10] provide a sufficient condition for weak-acyclicity, namely that any subgame contains a *unique* Nash equilibrium. Unfortunately, this criterion is not very useful for most voting rules, where typically (at least) all unanimous votes form equilibria. Another sufficient condition due to Apt and Simon [2] is by eliminating never-best-reply strategies, and the prospects of applying it to voting games is not yet clear. We thus need to use specialized techniques.

Plurality with random tie-breaking This is our main theorem:

Theorem 16. *$f^{PR}_{\hat{\mathbf{s}}}$ is weak-FDRP.*

The full proof is in the appendix. As an outline, we consider all states reachable form some initial state $\mathbf{a}^0$ via direct-replies. The main part is showing that one of these states must be truthful in some derived subgame. Then, we get from Theorem 9 that the subgame must have some reachable PNE, which must also be reachable from $\mathbf{a}^0$ by construction.

Weighted Plurality We saw in earlier sections that cycles of direct responses can emerge when voters are weighted. We conjecture that such cycles must depend on the order of agents, and that certain orders will break such cycles and reach an equilibrium, at least from the truthful state.

Conjecture 17. *$f^{PL}_{\hat{\mathbf{s}},\mathbf{w}}$ is weak-FDRP (in particular weak-FIP).*

Similar techniques to those used so far appear to be insufficient to prove the conjecture. For example, in contrast to the unweighted case, a voter might return to a candidate she deserted in *any scheduler*, even if only two weight levels are present. We thus leave the proof of the general conjecture for future work.

Yet, we want to demonstrate the power of weak acyclicity over restricted acyclicity, even when there are no restrictions on the utility space. That is, to provide a definite (negative) answer to Kukushkin's question of whether weak acyclicity entails restricted acyclicity.

Theorem 18. *There exist a game form f^* s.t. f^* is weak-FIP but not restricted-FIP.*

Proof. Consider the game defined in Example 4. If we ignore agents' preferences, we get a particular game form $f^{PL}_{\hat{\mathbf{s}},\mathbf{w}}$ where $N = \{1, 2, 3\}$, $M = \{a, b, c, d\}$, $\hat{\mathbf{s}} = (0, 1, 2, 3)$ and $\mathbf{w} = (1, 2, 3)$.

We define f^* by modifying $f^{PL}_{\hat{\mathbf{s}},\mathbf{w}}$ with the following restrictions on agents' actions: $A_1 = \{c, d\}, A_2 = \{b, c\}, A_3 = \{a, b, d\}$. Thus f^* is a $2 \times 2 \times 3$ game form, presented in Figure 2(a).

We first show that f^* is not restricted-FIP. Indeed, consider the game G^* accepted from f^* with the preferences from Example 4 (Figure 2(b)). We can see that there is a cycle of length 6 (in bold). An agent scheduler that always selects the agent with the bold reply guarantees convergence does not occur, since in all 6 relevant states the selected agent has no alternative replies.[8]

Next, we show that f^* is weak-FIP. That is, for any preference profile there is some scheduler that guarantees convergence. We thus divide into cases according to the preferences of agent 3. In each case, we specify a state where the scheduler selects agent 3, the action of the agent, and the new state.

We note that since all thick edges must be oriented in the same direction, $a \succ_3 b$ if and only if $b \succ_3 c$. Thus the following three cases are exhaustive.

	Q_3	state	action	new state
1	$b \succ d$	(d, b, a)	b	(d, b, b)
2	$d \succ b$ & $d \succ a$	(c, b, b)	d	(c, b, d)
3	$a \succ d \succ b \succ c$	(d, c, b)	d	(d, c, d)

[8] In contrast, Example 4 only shows that direct replies do not converge, but allowing indirect replies would break the cycle.

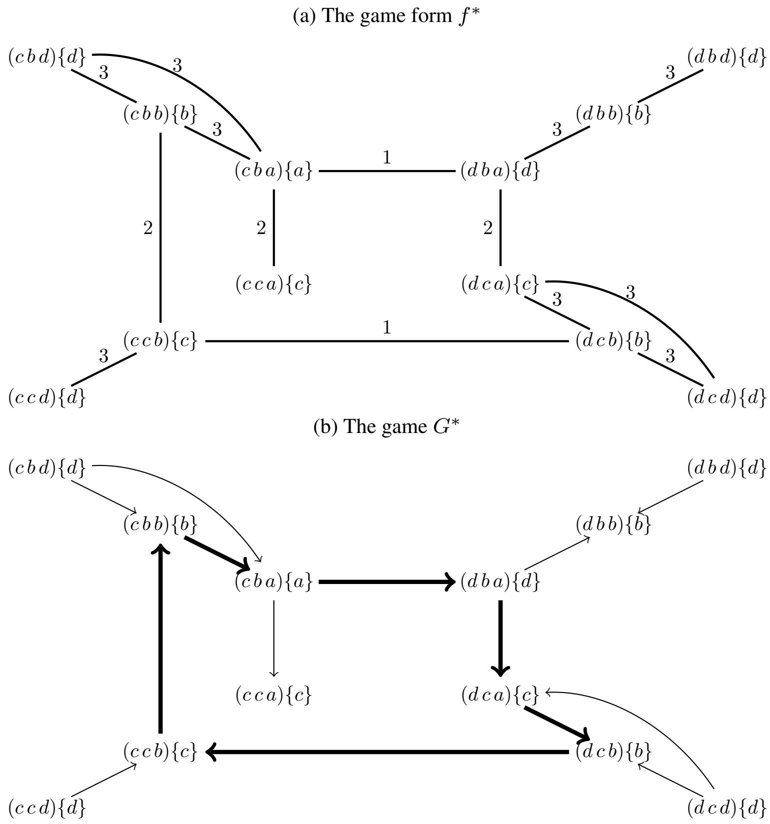

Figure 2: In each state we specify the actions of all 3 agents, and the outcome in curly brackets. Agent 1 controls the horizontal axis, agent 2 the vertical axis, and agent 3 the in/out axis. We omit edges between states with identical outcomes, since such moves are impossible for any transitive preferences. A directed edge in (b) is a better-reply in G^*.

In either case, agent 3 moves from a state on the cycle to a Nash equilibrium. □

6 Other Voting Rules

Borda Lev and Rosenschein claim that Borda fails to converge with best-replies from the truth, regardless of the tie-breaking rule being used. That is, that Borda is not FBRP (from the truth). In their example, there are four candidates $C = \{a, b, c, d\}$, and two voters, where $a \succ_1 b \succ_1 c \succ_1 d; c \succ_2 d \succ_2 b \succ_2 a$.

However, the steps used in their proof are no longer best-replies for some tie-breakings (e.g. the first step which makes b the winner is not a best-reply if a beats c in a tie). The proof still shows that Borda is not FIP from the truth for any tie-breaking (since they describe a cycle of better-replies). Further, if we assume specific linear tie-breaking where $b \succ c \succ a$, then all steps in the proof are indeed best-replies. Since this is just a permuted lexicographic order, we get the following corollary:

Proposition 19 (Roughly Theorem 3 in [19])**.** *Borda with lexicographic tie-breaking is not FBRP.*

Reyhani and Wilson [31] provided an even stronger counter example, with only three candidates ($a \succ_1 b \succ_1 c, b \succ_2 c \succ_2 a$). They remark that they only allow "reasonable" moves, but this remark is redundant, since in each step there is only one valid better-reply.

Proposition 20 (Example 6.1 in [31])**.** *Borda with lexicographic tie-breaking is not weak-FIP.*

Voting rule	FIP	FBRP	FDBRP	restricted-FIP	Weak-FIP
Dictator	V	V	-	V	V
Plurality (lex.)	X	X [22]	V [22]	V	V
Plurality (LD/SD)	X [22]	?	V [31, 24]	V	V
Plurality (rand.)	X	X	X	X[22]	V (Thm. 16)
Weighted Plurality (lex.)	X	X	X [22]	?	?
Veto	X	X (Ex. 23)	V [31, 19]	V	V
k-approval ($k \geq 2$)	X	X [19, 20]	-	X	X (Ex. 24)
Borda	X	X [19, 31]	-	X	X [31]
PSRs (except k-approval)	X	X [19, 20]	-	?	?
Approval	X	X (Thm. 25)	-	V (Thm. 25)	V

Table 1: Positive results carry to the right side, negative to the left side. We assume lexicographic tie breaking in all rules except Plurality. FDBRP is only well-defined for Plurality and Veto.

Veto (AntiPlurality) [31] and [19] independently provided convergence results for the Veto rule. Both considered a natural restriction on better replies by vetoing the current winner. We will thus call a such a vote a "direct reply" under veto.

Theorem 21 (Theorem 4.8 in [31], Theorem 4 in [19])**.** *Veto with lexicographic tie-breaking is FDBRP.*

Theorem 22 (Theorem 4.11 in [31])**.** *Veto with stochastic-dominance tie-breaking is FDBRP.*

We did not check if the proofs still hold for FDRP, as is the case for Plurality.
Since neither paper showed that the restriction to direct responses is necessary, we show it next.

Proposition 23. *Veto with lexicographic tie-breaking is not FBRP.*

Example 23. Denote by $-x$ the action of vetoing x. Consider the preference profile over $C = \{a, b, c\}$: $b \succ_1 c \succ_1 a; a \succ_2 c \succ_2 b$. The initial score is $(-1, 0, 0)$. Then there is a cycle:
$(-c, -b)\{a\} \xrightarrow{1} (-b, -b)\{c\} \xrightarrow{2} (-b, -c)\{a\} \xrightarrow{1} (-c, -c)\{b\} \xrightarrow{2} (-c, -b)\{a\}$. ◊

k-approval and PSRs It has been shown that k-Approval with lexicographic tie-breaking is not FBRP even from the truth, for any $k \geq 2$ (see [31, Example 6.3] for $k = 2$, [19, Theorem 8] for $k \geq 3$).

This result was later extended by Lev [20, Theorem 4.3] to any Positional Scoring Rule (PSR) except Plurality and Veto. We show that at least for k-Approval, results are even more negative.

Proposition 24. *For any fixed $k \geq 2$, k-approval with lexicographic tie-breaking may have no Nash equilibrium. In particular it is not weak-FIP.*

Example 24. We extend the example from [31] for $k = 2$. We use 2 voters, and a linear tie-breaking order where $x \in C$ beats all other candidates. Voter 1 ranks $x \in C$ first, and voter 2 ranks x last. The other candidates are ranked arbitrarily.

To see why there are no equilibria, note that the score of the winner is at most 2. Thus classify into the sets S_1, S_2 all states where the score of the winner is 1 and 2, respectively. In any $\mathbf{a} \in S_1$ candidate x wins, and there is some $z \in a_1 \setminus a_2$ s.t. $z \neq x$ (since $a_1 \neq a_2$). Thus voter 2 has an incentive to add one of her votes to z, and reach $f(\mathbf{a}') = z \succ_2 x$. In all of S_2, $x \in a_1 \cap a_2$, in which case voter 2 has a move to any other candidate, or the winner is $y \neq x$. In the latter case, voter 1 has a move to $C \setminus a_1$, which makes x the winner. ◊

Clearly if there are $n \geq 3$ voters and $m > kn$ candidates, then there is an equilibrium where one candidate (the one leading on tie-breaking) gets 3 votes, where all others get at most 1. However it is unclear whether weak-FIP (or even restricted-FIP) holds for any parameter value, or for other PSRs.

The problem is that unlike Plurality and Veto, there may be a best-replies in k-approval even when there are no direct-replies at all.

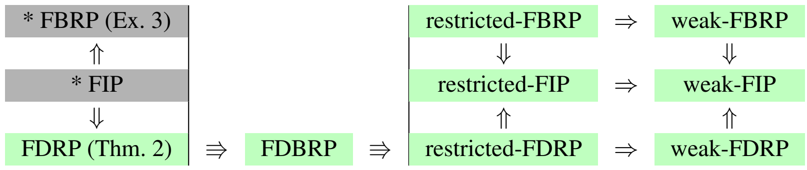

Figure 3: Convergence results for Plurality under lexicographic tie-breaking. Positive results (in light green) carry with the direction of the arrows, whereas negative results (dark gray) carry in the opposite direction. (*) a positive result for $n = 2$ strategic voters from the truth (Thm. 6).

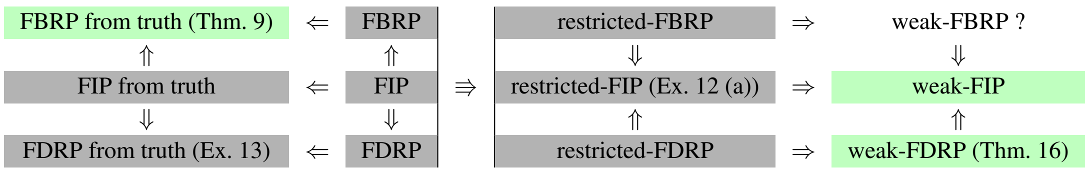

Figure 4: Convergence results for Plurality under random tie-breaking.

Approval Another corollary, is that in the Approval rule the voting policy will determine whether convergence is guaranteed. It is easy to see that the same results holds for $\leq k$-Approval, where each voter marks at most k candidates, and that at least the positive part holds for k-Range voting, since the voter can always concentrate her voting power on the intended winner.

Theorem 25. *Approval with lexicographic tie-breaking is restricted-FBRP but not FBRP (even from the truth).*

Proof. The positive part follows from Theorem 2: in any best-reply $a'_i \subseteq C$ only one candidate x becomes the winner. Thus the voter can restrict her action to $a''_i = \{x\}$, and the game becomes a Plurality game (after the first step of each voter).

For the negative side, we observe that in the example of Lev [20] for 2-Approval, every step is also a valid best-reply under the Approval rule. □

Other rules A new paper by Koolyk et al. [15] suggests that most common voting rules (including Copeland, Maximin, STV and others) with lexicographic tie-breaking, are not FBRP.

7 Conclusions and Future Work

We summarize most known results on iterative voting in Table 1. For Plurality we provide a more detailed picture in Figs. 3,4. Note that in some cases we get positive results if we restrict the initial state or the number of voters (not shown in the table).

Beyond the direct implication of various acyclicity properties on convergence in an interactive setting where agents vote one-by-one, [strong/weak] acyclicity is tightly linked to the convergence properties of more sophisticated learning strategies in repeated games [5, 21], which is another reason to understand them.

Based on the work summarized, and the additional progress made in this paper, we believe that research in this area should focus on three primary directions:

1. Weak-acyclicity seems more indicative than order-free acyclicity to determine convergence in practice. Thus theorists should study which voting rules are weak-FIP, perhaps under reasonable restrictions (as we demonstrated, this property is distinct from restricted-FIP).

2. It is important to experimentally study how people really vote in iterative settings (both in and out of the lab), so that this behavior can be formalized and behavioral models can be improved. The work of [33] is a

preliminary step in this direction, but there is much more to learn. Ideally, we would like to identify a few types of voters, such that for each type we can relatively accurately predict the next action in a particular state. It would be even better if these types are not specific to a particular voting rule or contextual details.

3. We would like to know not only if a voting rule converges under a particular dynamics (always or often), but also what are the properties of the attained outcome—in particular, whether the iterative process improves welfare or fairness, avoids "voting paradoxes" [34] and so on. Towards this end, several researchers (e.g., [30, 7, 23, 6, 15]) have started to explore these questions via theory and simulations. However, a good understanding of how iterative voting shapes the outcome, whether the population of voters consists of humans or artificial agents, is still missing.

A Proofs

Lemma 26. *Consider any game $G = \langle f^{PR}_{\hat{\mathbf{s}}}, \mathbf{Q} \rangle$. Consider some candidate a^*, and suppose that in $\mathbf{a}^0$, there are x, y s.t. $s^0(x) \geq s^0(y) \geq s^0(a^*) + 2$. Then for any sequence of direct replies, $a^* \notin f(\mathbf{a}^t)$.*

Proof. We show that at any time $t \geq 0$ there are x^t, y^t s.t. $s^0(x), s^0(y) \geq s^0(a^*) + 2$. For $t = 0$ this holds for $x^t = x, y^t = y$. Assume by induction that the premise holds for $\mathbf{a}^{t-1}$. Then there are two cases:

1. $|f(\mathbf{a}^{t-1})| \geq 2$. Then since step t must be a direct reply, it must be to some candidate z with $s^{t-1}(z) \geq sw^{t-1} - 1$. Also, either x^{t-1} or y^{t-1} did not lose votes (w.l.o.g. x^{t-1}). Thus $s^t(x), s^t(z) \geq sw^{t-1} \geq s^{t-1}(a^*) + 2 \geq s^t(a^*) + 2$.

2. $|f(\mathbf{a}^{t-1})| = 1$. Then suppose $f(\mathbf{a}^{t-1}) = \{x^{t-1}\}$, and we have that $sw^{t-1} \geq s^{t-1}(a^*) + 3$. The next step is z where either $s^{t-1}(z) = sw^{t-1} - 1$ (and then we conclude as in case 1), or $s^{t-1}(z) = sw^{t-1} - 2$ and x^{t-1} loses 1 vote. In the latter case, $s^t(x^{t-1}) = s^t(z) = sw^{t-1} - 1 \geq s^{t-1}(a^*) + 2 \geq s^t(a^*) + 2$. □

Theorem 16. *$f^{PR}_{\hat{\mathbf{s}}}$ is weak-FDRP.*

Proof. Consider a game $G = \langle f^{PR}_{\hat{\mathbf{s}}}, \mathbf{Q} \rangle$, and an initial state $\mathbf{a}^0$. For a state $\mathbf{a}$, denote by $B(\mathbf{a}) \subseteq A^n$ all states reachable from $\mathbf{a}$ via paths of direct replies. Let $B = B(\mathbf{a}^0)$, and assume towards a contradiction that B does not contain a Nash equilibrium. For every $\mathbf{b} \in B$, let $C(\mathbf{b}) = \{c \in C : \exists \mathbf{a} \in B(\mathbf{b}) \wedge c \in f(\mathbf{a})\}$, i.e. all candidates that are winners in some state reachable from $\mathbf{b}$.

For any $\mathbf{b} \in B(\mathbf{a}^0)$, define a game $G_{\mathbf{b}}$ by taking G and eliminating all candidates *not in* $C(\mathbf{b})$. Since we only consider direct replies, for any $\mathbf{a} \in B(\mathbf{b})$, the set of outgoing edges $I(\mathbf{a})$ is the same in G and in $G_{\mathbf{b}}$ (as any direct reply must be to candidate in $C(\mathbf{b})$). Thus by our assumption, the set $B(\mathbf{b})$ in game $G_{\mathbf{b}}$ does not contain an NE.

For any $\mathbf{b} \in B(\mathbf{a}^0)$, let $\mathbf{b}^*$ be the truthful state of game $G_{\mathbf{b}}$, and let $T(\mathbf{b}) \subseteq N$ be the set of agents who are truthful in $\mathbf{b}$. That is, $i \in T(\mathbf{b})$ if $b_i = b^*_i$.

Let $\mathbf{b}^0$ be some state $\mathbf{b} \in B(\mathbf{a}^0)$ s.t. $|T(\mathbf{b})|$ is maximal, and let $T^0 = T(\mathbf{b}^0)$. If $|T^0| = n$ then $\mathbf{b}^0$ is the truthful state of $G_{\mathbf{b}^0}$, and thus by Theorem 9 all best-reply paths from $\mathbf{b}^0$ in $G_{\mathbf{b}^0}$ lead to an NE, in contradiction to $B(\mathbf{b}^0)$ not containing any NE. Thus $T^0 < n$. We will prove that there is a path from $\mathbf{b}^0$ to a state $\mathbf{b}'$ s.t. $|T(\mathbf{b}')| > |T^0|$.

Let $i \notin T(\mathbf{b}^0)$ (must exist by the previous paragraph). Consider the score of candidate b^*_i at state $\mathbf{b}^0$. We divide into 5 cases. All scores specified below are in the game $G_{\mathbf{b}^0}$.

Case 1. $|f(\mathbf{b}^0)| > 1$ and $b^*_i \in f(\mathbf{b}^0)$ (i.e b^*_i is one of several winners). Then consider the step $\mathbf{b}^0 \xrightarrow{i} b^*_i$. This make b^*_i the unique winner, and thus it is a direct best-reply for i. In the new state $\mathbf{b}' = (\mathbf{b}^0_{-i}, b^*_i)$ we have $T(\mathbf{b}') = T(\mathbf{b}^0) \cup \{i\}$.

Case 2. $s^0(b^*_i) = sw^0 - 1$ (i.e., b^*_i needs one more vote to become a winner). By Axioms G+R, i prefers $f(\mathbf{b}^0_{-i}, b^*_i)$ over $f(\mathbf{b}^0)$. Then similarly to case 1, i has a direct step $\mathbf{b}^0 \xrightarrow{i} b^*_i$, which results in a "more truthful" state $\mathbf{b}'$.

Case 3. $b_i^* = f(\mathbf{b}^0)$ (i.e b_i^* is the unique winner). Then the next step $\mathbf{b}^0 \xrightarrow{j} \mathbf{b}^1$ will bring us to one of the two previous cases. Moreover, it must hold that $j \notin T(\mathbf{b}^0)$ since otherwise $b_j^0 = b_j^* = f(\mathbf{b}^0)$ which means $I_j(\mathbf{b}^0) = \emptyset$. Thus $|T(\mathbf{b}')| = |T(\mathbf{b}^1)| + 1 \geq |T(\mathbf{b}^0)| + 1$.

Case 4. $f(\mathbf{b}^0) = x \neq b_i^*$, and $s^0(x) = s^0(b_i^*) + 2$. We further divide into:

Case 4.1. $s^0(b_i^*) \geq s^0(y)$ for all $y \neq x$. Then the next step by j must be from x, which brings us to one of the two first cases (as in Case 3).

Case 4.2. There is $y \neq x$ s.t. $s^0(x) = s^0(y) + 1 = s^0(b_i^*) + 2$. Then we continue the sequence of steps until the winner's score decreases. Since all steps that maintain sw^t select a more preferred candidate, this most occur at some time t, and $T(\mathbf{b}^0) \subseteq T(\mathbf{b}^t)$. Then at $\mathbf{b}^t$ we are again in Case 1 or 2.

Case 4.3. There is $y \neq x$ s.t. $s^0(x) = s^0(y) = s^0(b_i^*) + 2$. Then by Lemma 26 b_i^* can never be selected, in contradiction to $b_i^* \in C(\mathbf{b}^0)$.

Case 5. $f(\mathbf{b}^0) = x \neq b_i^*$, and $s^0(x) \geq s^0(b_i^*) + 3$. We further divide into:

Case 5.1. For all $y \neq x$, $s^0(y) \leq s^0(x) - 3$. In this case no reply is possible.

Case 5.2. There is some $y \neq x$ s.t. $s^0(y) \geq s^0(b_i^*) + 2$. Then by Lemma 26 b_i^* can never be selected, in contradiction to $b_i^* \in C(\mathbf{b}^0)$.

Case 5.3. There is some $y \neq x$ s.t. $s^0(y) \geq s^0(b_i^*) + 1$ Then the next step must be from x to such y. Which means $s^1(x) = s^1(y) = sw^0 - 1 \geq s^0(b_i^*) + 2 = s^1(b_i^*) + 2$. Thus again by Lemma 26 we reach a contradiction.

Therefore we either construct a path of direct replies to $\mathbf{b}' \in B(\mathbf{b}^0)$ with $|T(\mathbf{b}')| > |T(\mathbf{b}^0)|$ in contradiction to our maximality assumption, or we reach another contradiction. Thus $B(\mathbf{b}^0)$ must contain some NE (both in $G_{\mathbf{b}^0}$ and in G), which means by construction that G is weakly-FDRP from $\mathbf{b}^0$. However since $\mathbf{b}^0 \in B(\mathbf{a}^0)$, we get that G is weakly-FDRP from $\mathbf{a}^0$ as well. $\square$

References

[1] Daniel Andersson, Vladimir Gurvich, and Thomas Dueholm Hansen. On acyclicity of games with cycles. *Discrete Applied Mathematics*, 158(10):1049–1063, 2010.

[2] Krzysztof R. Apt and Sunil Simon. A classification of weakly acyclic games. In *SAGT'12*, pages 1–12. 2012.

[3] Tilman Borgers. Pure strategy dominance. *Econometrica*, 61(2):423–430, 1993.

[4] Endre Boros, Vladimir Gurvich, Kazuhisa Makino, and Dávid Papp. Acyclic, or totally tight, two-person game forms: Characterization and main properties. *Discrete Mathematics*, 310(6):1135–1151, 2010.

[5] Michael Bowling. Convergence and no-regret in multiagent learning. *Advances in neural information processing systems*, 17:209–216, 2005.

[6] Clark Bowman, Jonathan K. Hodge, and Ada Yu. The potential of iterative voting to solve the separability problem in referendum elections. *Theory and decision*, 77(1):111–124, 2014.

[7] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. How bad is selfish voting? In *Proc. of 27th AAAI*, 2013.

[8] Antoine-Augustin Cournot. *Recherches sur les principes mathématiques de la théorie des richesses par Augustin Cournot.* chez L. Hachette, 1838.

[9] Edith Elkind, Umberto Grandi, Francesca Rossi, and Arkadii Slinko. Gibbard–satterthwaite games. In *IJCAI'15*, 2015.

[10] Alex Fabrikant, Aaron D Jaggard, and Michael Schapira. On the structure of weakly acyclic games. In *SAGT'10*, pages 126–137. Springer, 2010.

[11] Peter Gärdenfors. Manipulation of social choice functions. *Journal of Economic Theory*, 13(2):217–228, 1976.

[12] Neelam Gohar. *Manipulative Voting Dynamics*. PhD thesis, University of Liverpool, 2012.

[13] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *ADT'13*, pages 181–192. Springer, 2013.

[14] Jerry S Kelly. Strategy-proofness and social choice functions without singlevaluedness. *Econometrica: Journal of the Econometric Society*, pages 439–446, 1977.

[15] Aaron Koolyk, Omer Lev, and Jeffrey S Rosenschein. Convergence and quality of iterative voting under non-scoring rules (extended abstract). In *AAMAS'16*, 2016.

[16] Nikolai S. Kukushkin. Congestion games: a purely ordinal approach. *Economics Letters*, 64:279–283, 1999.

[17] Nikolai S. Kukushkin. Perfect information and congestion games. *Games and Economic Behavior*, 38:306–317, 2002.

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

[19] Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative voting. In *Proc. of 11th AAMAS*, pages 611–618, 2012.

[20] Omer Lev. *Agent Modeling of Human Interaction: Stability, Dynamics and Cooperation*. PhD thesis, The Hebrew University of Jerusalem, 2015.

[21] Jason R Marden, Gürdal Arslan, and Jeff S Shamma. Regret based dynamics: convergence in weakly acyclic games. In *AAMAS'07*. ACM, 2007.

[22] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nick Jennings. Convergence to equilibria of plurality voting. In *Proc. of 24th AAAI*, pages 823–828, 2010.

[23] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A local-dominance theory of voting equilibria. In *EC'14*, 2014.

[24] Reshef Meir. Plurality voting under uncertainty. In *AAAI'15*, pages 2103–2109, 2015.

[25] Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and economic behavior*, 13(1):111–124, 1996.

[26] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[27] Svetlana Obraztsova, Evangelos Markakis, and David R. M. Thompson. Plurality voting with truth-biased agents. In *SAGT'13*, pages 26–37. Springer, 2013.

[28] 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 *Proc. of 29th AAAI*, 2015.

[29] Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey S. Rosenschein. Analysis of equilibria in iterative voting schemes. In *Proc. of 29th AAAI*, 2015.

[30] Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proc. of 11th AAMAS*, pages 635–644, 2012.

[31] Reyhaneh Reyhani and Mark C Wilson. Best-reply dynamics for scoring rules. In *ECAI'12*. IOS Press, 2012.

[32] Alvin E Roth. The college admissions problem is not equivalent to the marriage problem. *Journal of economic Theory*, 36(2):277–288, 1985.

[33] Maor Tal, Reshef Meir, and Ya'acov Gal. A study of human behavior in voting systems. In *AAMAS 15*, pages 665–673, 2015.

[34] Lirong Xia, Jérôme Lang, and Mingsheng Ying. Sequential voting rules and multiple elections paradoxes. In *TARK'07*, pages 279–288, 2007.

[35] H Peyton Young. The evolution of conventions. *Econometrica: Journal of the Econometric Society*, pages 57–84, 1993.

Reshef Meir
Industrial Engineering and Management
Technion-Israel Institute of Technology
Technion campus, Haifa 32000, Israel
Email: `reshefm@ie.technion.ac.il`