

Convergence to Equilibria in Plurality Voting

Reshef Meir Maria Polukarov Jeffrey S. Rosenschein Nicholas R. Jennings

Abstract

Multi-agent decision problems, in which independent agents have to agree on a joint plan of action or allocation of resources, are central to AI. In such situations, agents' individual preferences over available alternatives may vary, and they may try to reconcile these differences by voting. Based on the fact that agents may have incentives to vote strategically and misreport their real preferences, a number of recent papers have explored different possibilities for avoiding or eliminating such manipulations. In contrast to most prior work, this paper focuses on convergence of strategic behavior to a decision from which no voter will want to deviate.

We consider scenarios where voters cannot coordinate their actions, but are allowed to change their vote after observing the current outcome. We focus on the Plurality voting rule, and study the conditions under which this iterative game is guaranteed to converge to a Nash equilibrium (i.e., to a decision that is stable against further unilateral manipulations).

We show for the first time how convergence depends on the exact attributes of the game, such as the tie-breaking scheme, and on assumptions regarding agents' weights and strategies.

1 Introduction

The notion of strategic voting has been highlighted in research on Social Choice as crucial to understanding the relationship between preferences of a population, and the final outcome of elections. The most widely used voting rule is the Plurality rule, in which each voter has one vote and the winner is the candidate who received the highest number of votes. While it is known that no reasonable voting rule is completely immune to strategic behavior, Plurality has been shown to be particularly susceptible, both in theory and in practice [12, 8]. This makes the analysis of any election campaign—even one where the simple Plurality rule is used—a challenging task. As voters may speculate and counter-speculate, it would be beneficial to have formal tools that would help us understand (and perhaps predict) the final outcome.

Natural tools for this task include the well-studied solution concepts developed for normal form games. While voting games are not commonly presented in this way, several natural formulations have been proposed. Moreover, such formulations are extremely simple in Plurality voting games, where voters only have a few ways available to vote.

While some work has been devoted to the analysis of solution concepts such as *dominant strategies* and *strong equilibria*, this paper concentrates on Nash equilibria (NE). This most prominent solution concept has typically been overlooked, mainly because it appears to be too weak for this problem: there are typically many Nash equilibria in a voting game, but most of them are trivial. For example, if all voters vote for the same candidate, then this is clearly an equilibrium, since any single agent cannot change the result. This means that Plurality is *distorted*, i.e., there can be NE points in which the outcome is not truthful.

The lack of a single prominent solution for the game suggests that in order to fully understand the outcome of the voting procedure, it is not sufficient to consider voters' preferences. The strategies voters' choose to adopt, as well as the information available to them, are necessary for the analysis of possible outcomes. To play an equilibrium strategy for example, voters must know the preferences of others. Partial knowledge is also required in order to eliminate dominated strategies or to collude with other voters.

We consider the other extreme, assuming that voters have initially no knowledge regarding the preferences of the others, and cannot coordinate their actions. Such situations may arise, for example, when voters do not trust one another or have restricted communication abilities. Thus, even if

two voters have exactly the same preferences, they may be reluctant or unable to share this information, and hence they will fail to coordinate their actions. Voters may still try to vote strategically, based on their current information, which may be partial or wrong. The analysis of such settings is of particular interest to AI as it tackles the fundamental problem of multi-agent decision making, where autonomous agents (that may be distant, self-interested and/or unknown to one another) have to choose a joint plan of action or allocate resources or goods. The central questions are (i) whether, (ii) how fast, and (iii) on what alternative the agents will agree.

In our (Plurality) voting model, voters start from some announcement (e.g., the truthful one), but can change their votes after observing the current announcement and outcome.¹ The game proceeds in turns, where a single voter changes his vote at each turn. We study different versions of this game, varying tie-breaking rules, weights and policies of voters, and the initial profile. Our main result shows that in order to guarantee convergence, it is necessary and sufficient that voters restrict their actions to natural best replies.

1.1 Related Work

There have been several studies applying game-theoretic solution concepts to voting games, and to Plurality in particular. [7] model a Plurality voting game where candidates and voters play strategically. They characterize all Nash equilibria in this game under the very restrictive assumption that the preference domain is *single peaked*. Another highly relevant work is that of [5], which concentrates on *dominant strategies* in Plurality voting. Their game formulation is identical to ours, and they prove a necessary and sufficient condition on the profile for the game to be dominance-solvable. Unfortunately, their analysis shows that this rarely occurs, making dominance perhaps a too-strong solution concept for actual situations. A weaker concept, though still stronger than NE, is *Strong Equilibrium*. In strong equilibrium no subset of agents can benefit by making a coordinated diversion. A variation of strong equilibrium was suggested by [10], which characterized its existence and uniqueness in Plurality games. Crucially, all aforementioned papers assume that voters have some prior knowledge regarding the preferences of others.

A more complicated model was suggested by [11], which assumes a non-atomic set of voters and some uncertainty regarding the preferences of other voters. Their main result is that every positional scoring rule (e.g., Veto, Borda, and Plurality) admits at least one voting equilibrium. In contrast, our model applies to a finite number of voters, that possess zero knowledge regarding the distribution of other voters' preferences.

Variations of Plurality and other voting rules have been proposed in order to increase resistance to strategic behavior (e.g., [4]). We focus on achieving a stable outcome *taking such behavior into account*.

Iterative voting procedures have also been investigated in the literature. [3] consider voters with different levels of information, where in the lowest level agents are myopic (as we assume as well). Others assume, in contrast, that voters have sufficient information to forecast the entire game, and show how to solve it with backward induction [6, 9]; most relevant to our work, [1] study conditions for convergence in such a model.

2 Preliminaries

2.1 The Game Form

There is a set C of m candidates, and a set V of n voters. A voting rule f allows each voter to submit his preferences over the candidates by selecting an action from a set A (in Plurality, $A = C$). Then, f chooses a non-empty set of winner candidates—i.e., it is a function $f : A^n \rightarrow 2^C \setminus \{\emptyset\}$.

¹A real-world example of a voting interface that gives rise to a similar procedure is the recently introduced poll gadget for Google Wave. See <http://sites.google.com/site/polloforwave>.

v_1, v_2	a	b	c
a	(14, 9, 3) {a}	(10, 13, 3) {b}	(10, 9, 7) {a}
b	(11, 12, 3) {b}	(7, 16, 3) {b}	(7, 12, 7) {b}
c	(11, 9, 6) {a}	(7, 13, 6) {b}	(7, 9, 10) {c}

Table 1: There is a set $C = \{a, b, c\}$ of candidates with initial scores (7, 9, 3). Voter 1 has weight 3 and voter 2 has weight 4. Thus, $GF_T = \langle \{a, b, c\}, \{1, 2\}, (3, 2), (7, 9, 3) \rangle$. The table shows the outcome vector $\mathbf{s}(a_1, a_2)$ for every joint action of the two voters, as well as the set of winning candidates $GF_T(a_1, a_2)$. In this example there are no ties, and it thus fits both tie-breaking schemes.

Each such voting rule f induces a natural *game form*. In this game form, the strategies available to each voter are A , and the outcome of a joint action is $f(a_1, \dots, a_n)$. Mixed strategies are not allowed. We extend this game form by including the possibility that only k out of the n voters may play strategically. We denote by $K \subseteq V$ the set of k strategic voters (agents) and by $B = V \setminus K$ the set of $n - k$ additional voters who have already cast their votes, and are not participating in the game. Thus, the outcome is $f(a_1, \dots, a_k, b_{k+1}, \dots, b_n)$, where b_{k+1}, \dots, b_n are fixed as part of the game form. This separation of the set of voters does not affect generality, but allows us to encompass situations where only some of the voters behave strategically.

From now on, we restrict our attention to the Plurality rule, unless explicitly stated otherwise. That is, the winner is the candidate (or a set of those) with the most votes; there is no requirement that the winner gain an absolute majority of votes. We assume each of the n voters has a fixed *weight* $w_i \in \mathbb{N}$. The *initial score* $\hat{s}(c)$ of a candidate c is defined as the total weight of the fixed voters who selected c —i.e., $\hat{s}(c) = \sum_{j \in B: b_j=c} w_j$. The *final score* of c for a given joint action $\mathbf{a} \in A^k$ is the total weight of voters that chose c (including the fixed set B): $s(c, \mathbf{a}) = \hat{s}(c) + \sum_{i \in K: a_i=c} w_i$. We sometimes write $s(c)$ if the joint action is clear from the context. We write $s(c) >_p s(c')$ if either $s(c) > s(c')$ or the score is equal and c has a higher priority (lower index). We denote by PL_R the Plurality rule with randomized tie breaking, and by PL_D the Plurality rule with deterministic tie breaking in favor of the candidate with the lower index. We have that $PL_R(\hat{\mathbf{s}}, \mathbf{w}, \mathbf{a}) = \mathop{\text{argmax}}_{c \in C} s(c, \mathbf{a})$, and $PL_D(\hat{\mathbf{s}}, \mathbf{w}, \mathbf{a}) = \{c \in C \text{ s.t. } \forall c' \neq c, s(c, \mathbf{a}) >_p s(c', \mathbf{a})\}$. Note that $PL_D(\hat{\mathbf{s}}, \mathbf{w}, \mathbf{a})$ is always a singleton.

For any joint action, its *outcome vector* $\mathbf{s}(\mathbf{a})$ contains the score of each candidate: $\mathbf{s}(\mathbf{a}) = (s(c_1, \mathbf{a}), \dots, s(c_m, \mathbf{a}))$. For a tie-breaking scheme T ($T = D, R$) the *Game Form* $GF_T = \langle C, K, \mathbf{w}, \hat{\mathbf{s}} \rangle$ specifies the winner for any joint action of the agents—i.e., $GF_T(\mathbf{a}) = PL_T(\hat{\mathbf{s}}, \mathbf{w}, \mathbf{a})$. Table 1 demonstrates a game form with two weighted manipulators.

2.2 Incentives

We now complete the definition of our voting game, by adding incentives to the game form. Let \mathcal{R} be the set of all strict orders over C . The order $\succ_i \in \mathcal{R}$ reflects the preferences of voter i over the candidates. The vector containing the preferences of all k agents is called a *profile*, and is denoted by $\mathbf{r} = (\succ_1, \dots, \succ_k)$. The game form GF_T , coupled with a profile \mathbf{r} , define a normal form game $G_T = \langle GF_T, \mathbf{r} \rangle$ with k players. Player i prefers outcome $GF_T(\mathbf{a})$ over outcome $GF_T(\mathbf{a}')$ if $GF_T(\mathbf{a}) \succ_i GF_T(\mathbf{a}')$.

Note that for deterministic tie-breaking, every pair of outcomes can be compared. If ties are broken randomly, \succ_i does *not* induce a complete order over outcomes, which are *sets* of candidates. A natural solution is to augment agents' preferences with cardinal utilities, where $u_i(c) \in \mathbb{R}$ is the utility of candidate c to agent i . This definition naturally extends to multiple winners by setting $u_i(W) = \frac{1}{|W|} \sum_{c \in W} u_i(c)$.² A utility function u is *consistent* with a preference relation \succ_i if $u(c) > u(c') \Leftrightarrow c \succ_i c'$.

²This makes sense if we randomize the final winner from the set W . For a thorough discussion of cardinal and ordinal utilities in normal form games, see [2].

v_1, v_2	a	b	* c
* a	{a} 3, 2	{b} 2, 1	* {a} 3, 2
b	{b} 2, 1	{b} 2, 1	{b} 2, 1
c	{a} 3, 2	{b} 2, 1	{c} 1, 3

Table 2: A game $G_T = \langle GF_T, \mathbf{r} \rangle$, where GF_T is as in Table 1, and \mathbf{r} is defined by $a \succ_1 b \succ_1 c$ and $c \succ_2 a \succ_2 b$. The table shows the ordinal utility of the outcome to each agent (the final score is not shown). **Bold** outcomes are the NE points. Here the truthful vote (marked with *) is also a NE.

Lemma 1. *For any utility function u which is consistent with preference order \succ_i , the following holds:*

1. $a \succ_i b \Rightarrow \forall W \subseteq C \setminus \{a, b\}, u(\{a\} \cup W) > u(\{b\} \cup W)$;
2. $\forall b \in W, a \succ_i b \Rightarrow u(a) > u(\{a\} \cup W) > u(W)$.

The proof is trivial and is therefore omitted. Lemma 1 induces a partial preference order on the set of outcomes, but it is not yet complete if the cardinal utilities are not specified. For instance, the order $a \succ_i b \succ_i c$ does not determine if i will prefer $\{b\}$ over $\{a, c\}$. When utilities are given explicitly, every pair of outcomes can be compared, and we will slightly abuse the notation by using $GF_R(\mathbf{a}) \succ_i GF_R(\mathbf{a}')$ to note that i prefers the outcome of action \mathbf{a} over that of \mathbf{a}' .

2.3 Manipulation and Stability

Having defined a normal form game, we can now apply standard solution concepts. Let $G_T = \langle GF_T, \mathbf{r} \rangle$ be a Plurality voting game, and let $\mathbf{a} = (a_{-i}, a_i)$ be a joint action in G_T . We say that $a_i \xrightarrow{i} a'_i$ is an *improvement step* of agent i if $GF_T(a_{-i}, a'_i) \succ_i GF_T(a_{-i}, a_i)$. A joint action \mathbf{a} is a *Nash equilibrium* (NE), if no agent has an improvement step from \mathbf{a} in G_T . 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. However, in our voting games there are typically (but not necessarily) many such points.

Now, observe that the preference profile \mathbf{r} induces a special joint action \mathbf{a}^* , termed the *truthful vote*, such that $\mathbf{a}^*(\mathbf{r}) = (a_1^*, \dots, a_k^*)$, where $a_i^* \succ_i c$ for all $c \neq a_i^*$. We also call $\mathbf{a}^*(\mathbf{r})$ the *truthful state* of G_T , and refer to $GF_T(\mathbf{a}^*(\mathbf{r}))$ as the *truthful outcome* of the game. If i has an improvement step in the truthful state, then this is a *manipulation*.³ Thus, \mathbf{r} cannot be manipulated if and only if $\mathbf{a}^*(\mathbf{r})$ is a Nash equilibrium of $G_T = \langle GF_T, \mathbf{r} \rangle$. However, the truthful vote may or may not be included in the NE points of the game, as can be seen from Table 2.

2.4 Game Dynamics

We finally consider natural *dynamics* in Plurality voting games. Assume that players start by announcing some initial vote, and then proceed and change their votes until no one has objections to the current outcome. It is not, however, clear how rational players would act to achieve a stable decision, especially when there are multiple equilibrium points. However, one can make some plausible assumptions about their behavior. First, the agents are likely to only make improvement steps, and to keep their current strategy if such a step is not available. Thus, the game will end when it first reaches a NE. Second, it is often the case that the initial state is truthful, as agents know that they can reconsider and vote differently, if they are not happy with the current outcome.

We start with a simple observation that if the agents may change their votes simultaneously, then convergence is not guaranteed, even if the agents start with the truthful vote and use best replies—that is, vote for their most preferred candidate out of potential winners in the current round.

³This definition of manipulation coincides with the standard definition from social choice theory.

Proposition 2. *If agents are allowed to re-vote simultaneously, the improvement process may never converge.*

Example. The counterexample is the game with 3 candidates $\{a, b, c\}$ with initial scores given by $(0, 0, 2)$. There are 2 voters $\{1, 2\}$ with weights $w_1 = w_2 = 1$ and the following preferences: $a \succ_1 b \succ_1 c$, and $b \succ_2 a \succ_2 c$. The two agents will repeatedly swap their strategies, switching endlessly between the states $\mathbf{a}(\mathbf{r}) = (a, b)$ and (b, a) . Note that this example works for both tie-breaking schemes. \diamond

We therefore restrict our attention to dynamics where simultaneous improvements are not available. That is, given the initial vote \mathbf{a}_0 , the game proceeds in steps, where at each step t , a single player may change his vote, resulting in a new state (joint action) \mathbf{a}_t . 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.

In the remaining sections, we study the conditions under which such iterative games reach an equilibrium point from either an arbitrary or a truthful initial state. We consider variants of the game that differ in tie-breaking schemes or assumptions about the agents' weights or behavior. In cases where convergence is guaranteed, we are also interested in knowing how fast it will occur, and whether we can say anything about the identity of the winner. For example, in Table 2, the game will converge to a NE from any state in at most two steps, and the outcome will be a (which happens to be the truthful outcome), unless the players initially choose the alternative equilibrium (b, b) with outcome b .

3 Results

Let us first provide some useful notation. We denote the outcome at time t by $o_t = PL(\mathbf{a}_t) \subseteq C$, and its score by $s(o_t)$. Suppose that agent i has an improvement step at time t , and as a result the winner switched from o_{t-1} to o_t . The possible steps of i are given by one of the following types (an example of such a step appears in parentheses):

type 1 from $a_{i,t-1} \notin o_{t-1}$ to $a_{i,t} \in o_t$; (step 1 in Ex.4a.)

type 2 from $a_{i,t-1} \in o_{t-1}$ to $a_{i,t} \notin o_t$; (step 2 in Ex.4a.)

type 3 from $a_{i,t-1} \in o_{t-1}$ to $a_{i,t} \in o_t$; (step 1 in Ex.4b.),

where inclusion is replaced with equality for deterministic tie-breaking. We refer to each of these steps as a *better reply* of agent i . If $a_{i,t}$ is i 's most preferred candidate capable of winning, then this is a *best reply*.⁴ Note that there are no best replies of type 2. Finally, we denote by $s_t(c)$ the score of a candidate c *without the vote of the currently playing agent*; thus, it always holds that $s_{t-1}(c) = s_t(c)$.

3.1 Deterministic Tie-Breaking

Our first result shows that under the most simple conditions, the game must converge.

Theorem 3. *Let G_D be a Plurality game with deterministic tie-breaking. If all agents have weight 1 and use best replies, then the game will converge to a NE from any state.*

⁴Any rational move of a myopic agent in the normal form game corresponds to exactly one of the three types of better-reply. In contrast, the definition of best-reply is somewhat different from the traditional one, which allows the agent to choose any strategy that guarantees him a best possible outcome. Here, we assume the improver makes the more natural response by actually voting for o_t . Thus, under our definition, the best reply is always unique.

Proof. We first show that there can be at most $(m - 1) \cdot k$ sequential steps of type 3. Note that at every such step $a \xrightarrow{i} b$ it must hold that $b \succ_i a$. Thus, each voter can only make $m - 1$ such subsequent steps.

Now suppose that a step $a \xrightarrow{i} b$ of type 1 occurs at time t . We claim that at any later time $t' \geq t$: (I) there are at least two candidates whose score is *at least* $s(o_{t-1})$; (II) the score of a will not increase at t' . We use induction on t' to prove both invariants. Right after step t we have that

$$s_t(b) + 1 = s(o_t) >_p s(o_{t-1}) >_p s_t(a) + 1 . \quad (1)$$

Thus, after step t we have at least two candidates with scores of at least $s(o_{t-1})$: $o_t = b$ and $o_{t-1} \neq b$. Also, at step t the score of a has decreased. This proves the base case, $t' = t$.

Assume by induction that both invariants hold until time $t' - 1$, and consider step t' by voter j . Due to (I), we have at least two candidates whose score is at least $s(o_{t-1})$. Due to (II) and Equation (1) we have that $s_{t'}(a) \leq_p s_t(a) <_p s(o_{t-1}) - 1$. Therefore, no single voter can make a a winner and thus a cannot be the best reply for j . This means that (II) still holds after step t' . Also, j has to vote for a candidate c that can beat $o_{t'}$ —i.e., $s_{t'}(c) + 1 >_p s(o_{t'}) >_p s(o_{t-1})$. Therefore, after step t' both c and $o_{t'} \neq c$ will have a score of at least $s(o_{t-1})$ —that is, (I) also holds. \square

The proof also supplies us with a polynomial bound on the rate of convergence. At every step of type 1, at least one candidate is ruled out permanently, and there at most k times a vote can be withdrawn from a candidate. Also, there can be at most mk steps of type 3 between such occurrences. Hence, there are in total at most m^2k^2 steps until convergence. It can be further shown that if all voters start from the truthful state then there are no type 3 steps at all. Thus, the score of the winner never decreases, and convergence occurs in at most mk steps. The proof idea is similar to that of the corresponding randomized case in Theorem 8.

We now show that the restriction to best replies is necessary to guarantee convergence.

Proposition 4. *If agents are not limited to best replies, then: (a) there is a counterexample with two agents; (b) there is a counterexample with an initial truthful vote.*

Example 4a. $C = \{a, b, c\}$. We have a single fixed voter voting for a , thus $\hat{s} = (1, 0, 0)$. The preference profile is defined as $a \succ_1 b \succ_1 c$, $c \succ_2 b \succ_2 a$. The following cycle consists of better replies (the vector denotes the votes (a_1, a_2) at time t , the winner appears in curly brackets):

$$(b, c)\{a\} \xrightarrow{2} (b, b)\{b\} \xrightarrow{1} (c, b)\{a\} \xrightarrow{2} (c, c)\{c\} \xrightarrow{1} (b, c) \quad \diamond$$

Example 4b. $C = \{a, b, c, d\}$. Candidates a, b , and c have 2 fixed voters each, thus $\hat{s} = (2, 2, 2, 0)$. We use 3 agents with the following preferences: $d \succ_1 a \succ_1 b \succ_1 c$, $c \succ_2 b \succ_2 a \succ_2 d$ and $d \succ_3 a \succ_3 b \succ_3 c$. Starting from the truthful state (d, c, d) the agents can make the following two improvement steps (showing only the outcome):

$$(2, 2, 3, 2)\{c\} \xrightarrow{1} (2, 3, 3, 1)\{b\} \xrightarrow{3} (3, 3, 3, 0)\{a\} ,$$

after which agents 1 and 2 repeat the cycle shown in (4a). \diamond

Weighted voters While using the best reply strategies guaranteed convergence for equally weighted agents, this is no longer true for non-identical weights:

Proposition 5. *There is a counterexample with 3 weighted agents that start from the truthful state and use best replies.*

The proof is omitted for the sake of brevity.

However, if there are *only two* weighted voters, either restriction is sufficient:

Theorem 6. Let G_D be a Plurality game with deterministic tie-breaking. If $k = 2$ and both agents (a) use best replies or (b) start from the truthful state, a NE will be reached.

Proof of (6a). Assume there is a cycle, and consider the winners in the first steps: $\{x\} \xrightarrow{1} \{y\} \xrightarrow{2} \{z\}$. Suppose that after step 1 both agents vote for different candidates ($a_{1,2} \neq a_{1,1} = y$). This holds for any later step, as an agent has no reason to vote for the current winner. An agent can never make a step of type 3 (after the first step), since at every step the winner is the candidate that the other agent is voting for. If the first step brings the agents to the same candidate, then in the second step they split again ($a_{2,1} \neq a_{2,2} = z$), and we are back in the previous case. \square

Proof of (6b). We show that the score of the winner can only increase. This clearly holds in the first step, which must be of type 1. Once again, we have that the two agents always vote for different candidates, and thus only steps that increase the score can change the identity of the winner. \square

3.2 Randomized Tie-Breaking

The choice of tie-breaking scheme has a significant impact on the outcome, especially when there are few voters. A randomized tie-breaking rule has the advantage of being neutral —no specific candidate or voter is preferred over another.

In order to prove convergence under randomized tie-breaking, we must show that convergence is guaranteed for *any* utility function which is consistent with the given preference order. That is, we may only use the relations over outcomes that follow directly from Lemma 1. To disprove, it is sufficient to show that for a specific assignment of utilities, the game forms a cycle. In this case, we say that there is a *weak counterexample*. When the existence of a cycle will follow only from the relations induced by Lemma 1, we will say that there is a *strong counterexample*, since it holds for any profile of utility scales that fits the preferences.

In contrast to the deterministic case, the weighted randomized case does not always converge to a Nash equilibrium or possess one at all, even with (only) two agents.

Proposition 7. There is a strong counterexample G_R for two weighted agents with randomized tie-breaking, even if both agents start from the truthful state and use best replies.

Example. $C = \{a, b, c\}$, $\hat{s} = (0, 1, 3)$. There are 2 agents with weights $w_1 = 5$, $w_2 = 3$ and preferences $a \succ_1 b \succ_1 c$, $b \succ_2 c \succ_2 a$ (in particular, $b \succ_2 \{b, c\} \succ_2 c$). The resulting 3×3 normal form game contains no NE states. \diamond

Nevertheless, the conditions mentioned are sufficient for convergence if all agents have the same weight.

Theorem 8. Let G_R be a Plurality game with randomized tie-breaking. If all agents have weight 1 and use best replies, then the game will converge to a NE from the truthful state.

Proof. Our proof shows that in each step, the current agent votes for a *less* preferred candidate. Clearly, the first improvement step of every agent must hold this invariant.

Assume, toward deriving a contradiction, that $b \xrightarrow{i} c$ at time t_2 is the first step s.t. $c \succ_i b$. Let $a \xrightarrow{i} b$ at time $t_1 < t_2$ be the previous step of the same agent i .

We denote by $M_t = o_t$ the set of all winners at time t . Similarly, L_t denotes all candidates whose score is $s(o_t) - 1$.

We claim that for all $t < t_2$, $M_t \cup L_t \subseteq M_{t-1} \cup L_{t-1}$, i.e., the set of “almost winners” can only shrink. Also, the score of the winner cannot decrease. Observe that in order to contradict any of these assertions, there must be a step $x \xrightarrow{j} y$ at time t , where $\{x\} = M_{t-1}$ and $y \notin M_{t-1} \cup L_{t-1}$. In that case, $M_t = L_{t-1} \cup \{x, y\} \succ_j \{x\}$, which means either that $y \succ_j x$ (in contradiction to the minimality of t_2) or that y is not a best reply.

From our last claim we have that $s(o_{t_1-1}) \leq s(o_{t'})$ for any $t_1 \leq t' < t_2$. Now consider the step t_1 . Clearly $b \in M_{t_1-1} \cup L_{t_1-1}$ since otherwise voting for b would not make it a winner. We consider the cases for c separately:

Case 1: $c \notin M_{t_1-1} \cup L_{t_1-1}$. We have that $s_{t_1}(c) \leq s(o_{t_1-1}) - 2$. Let t' be any time s.t. $t_1 \leq t' < t_2$, then $c \notin M_{t'} \cup L_{t'}$. By induction on t' , $s_{t'}(c) \leq s_{t_1}(c) \leq s(o_{t_1-1}) - 2 \leq s(o_{t'}) - 2$, and therefore c cannot become a winner at time $t'+1$, and the improver at time $t'+1$ has no incentive to vote for c . In particular, this holds for $t'+1 = t_2$; hence, agent i will not vote for c .

Case 2: $c \in M_{t_1-1} \cup L_{t_1-1}$. It is not possible that $b \in L_{t_1-1}$ or that $c \in M_{t_1-1}$: since $c \succ_i b$ and i plays best reply, i would have voted for c at step t_1 . Therefore, $b \in M_{t_1-1}$ and $c \in L_{t_1-1}$. After step t_1 , the score of b equals the score of c plus 2; hence, we have that $M_{t_1} = \{b\}$ and $c \notin M_{t_1} \cup L_{t_1}$, and we are back in case 1.

In either case, voting for c at step t_2 leads to a contradiction. Moreover, as agents only vote for a less-preferred candidate, each agent can make at most $m - 1$ steps, hence, at most $(m - 1) \cdot k$ steps in total. \square

However, in contrast to the deterministic case, convergence is no longer guaranteed, if players start from an arbitrary profile of votes. The following example shows that in the randomized tie-breaking setting even best reply dynamics may have cycles, albeit for specific utility scales.

Proposition 9. *If agents start from an arbitrary profile, there is a weak counterexample with 3 agents of weight 1, even if they use best replies.*

Example. There are 4 candidates $\{a, b, c, x\}$ and 3 agents with utilities $u_1 = (5, 4, 0, 3)$, $u_2 = (0, 5, 4, 3)$ and $u_3 = (4, 0, 5, 3)$. In particular, $a \succ_1 \{a, b\} \succ_1 x \succ_1 \{a, c\}$; $b \succ_2 \{b, c\} \succ_2 x \succ_2 \{a, b\}$; and $c \succ_3 \{a, c\} \succ_3 x \succ_3 \{b, c\}$. From the state $\mathbf{a}_0 = (a, b, x)$ with $\mathbf{s}(\mathbf{a}_0) = (1, 1, 0, 1)$ and the outcome $\{a, b, x\}$, the following cycle occurs: $(1, 1, 0, 1)\{a, b, x\} \xrightarrow{2} (1, 0, 0, 2)\{x\} \xrightarrow{3} (1, 0, 1, 1)\{a, x, c\} \xrightarrow{1} (0, 0, 1, 2)\{x\} \xrightarrow{2} (0, 1, 1, 1)\{x, b, c\} \xrightarrow{3} (0, 1, 0, 2)\{x\} \xrightarrow{1} (1, 1, 0, 1)\{a, b, x\}$. \diamond

As in the previous section, if we relax the requirement for best replies, there may be cycles even from the truthful state.

Proposition 10. (a) *If agents use arbitrary better replies, then there is a strong counterexample with 3 agents of weight 1. Moreover, (b) there is a weak counterexample with 2 agents of weight 1, even if they start from the truthful state.*

Example 10a. $C = \{a, b, c\}$ with initial score $\hat{\mathbf{s}} = (0, 1, 0)$. The initial state is $\mathbf{a}_0 = (a, a, b)$ —that is, $\mathbf{s}(\mathbf{a}_0) = (2, 2, 0)$ and the outcome is the winner set $\{a, b\}$. Consider the following cyclic sequence (we write the score vector and the outcome in each step): $(2, 2, 0)\{a, b\} \xrightarrow{2} (1, 2, 1)\{b\} \xrightarrow{1} (0, 2, 2)\{b, c\} \xrightarrow{3} (1, 1, 2)\{c\} \xrightarrow{2} (2, 1, 1)\{a\} \xrightarrow{3} (1, 2, 1)\{b\} \xrightarrow{1} (2, 2, 0)\{a, b\}$. If the preferences are $a \succ_1 c \succ_1 b$, $b \succ_2 a \succ_2 c$ and $c \succ_3 b \succ_3 a$, then each step is indeed an improvement step for the agent whose index is on top of the arrow. \diamond

Example 10b. We use 5 candidates with initial score $(1, 1, 2, 0, 0)$, and 2 agents with utilities $u_1 = (5, 3, 2, 8, 0)$ and $u_2 = (4, 2, 5, 0, 8)$. In particular, $\{b, c\} \succ_1 c$, $\{a, c\} \succ_1 \{a, b, c\}$, and $\{a, b, c\} \succ_2 \{b, c\}$, $c \succ_2 \{a, c\}$, and the following cycle occurs: $(1, 1, 2, 1, 1)\{c\} \xrightarrow{1} (1, 2, 2, 0, 1)\{b, c\} \xrightarrow{2} (2, 2, 2, 0, 0)\{a, b, c\} \xrightarrow{1} (2, 1, 2, 1, 0)\{a, c\} \xrightarrow{2} (1, 1, 2, 1, 1)\{c\}$. \diamond

3.3 Truth-Biased Agents

So far we assumed purely rational behavior on the part of the agents, in the sense that they were indifferent regarding their chosen action (vote), and only cared about the outcome. Thus, for example, if an agent cannot affect the outcome at some round, he simply keeps his *current vote*. This assumption is indeed common when dealing with normal form games, as there is no reason to prefer

Tie breaking	Dynamics Initial state	Best reply from		Any better reply from		Truth biased
		Truth	Anywhere	Truth	Anywhere	
Deterministic	Weighted ($k > 2$)	X (5)	X	X	X	X
	Weighted ($k = 2$)	V	V (6a)	V (6b)	X (4a)	X
	Non-weighted	V	V (3)	X (4b)	X	X (11a)
Randomized	Weighted	X (7)	X	X	X	X
	Non-weighted	V (8)	X (9)	X (10)	X (10)	X (11b)

Table 3: We highlight cases where convergence is guaranteed. The number in brackets refers to the index of the corresponding theorem (marked with **V**) or counterexample (X). Entries with no index follow from other entries in the table.

one strategy over another if outcomes are the same. However, in voting problems it is typically assumed that voters will vote *truthfully* unless they have an incentive to do otherwise. As our model incorporates both settings, it is important to clarify the exact assumptions that are necessary for convergence.

In this section, we consider a variation of our model where agents always prefer their higher-ranked outcomes, but will vote honestly if the outcome remains the same—i.e., the agents are *truth-biased*. Formally, let $W = PL_T(\hat{s}, \mathbf{w}, a_i, \mathbf{a}_{-i})$ and $Z = PL_T(\hat{s}, \mathbf{w}, a'_i, \mathbf{a}_{-i})$ be two possible outcomes of i 's voting. Then, the action a'_i is *better than* a_i if either $Z \succ_i W$, or $Z = W$ and $a'_i \succ_i a_i$. Note that with this definition there is a strict preference order over all possible actions of i at every step. Unfortunately, truth-biased agents may not converge even in the simplest settings:

Proposition 11. *There are strong counterexamples for (a) deterministic tie-breaking, and (b) randomized tie-breaking. This holds even with two non-weighted truth-biased agents that use best reply dynamics and start from the truthful state.*

Example 11a. We use 4 candidates with no initial score. The preferences are defined as $c \succ_1 a \succ_1 b \succ_1 d$ and $d \succ_2 b \succ_2 a \succ_2 c$. The reader can easily verify that in the resulting 4×4 game there are no NE states. \diamond

Example 11b. There are 4 candidates with initial scores $(0, 0, 1, 2)$. The preference profile is given by $a \succ_1 c \succ_1 d \succ_1 b$, $b \succ_2 d \succ_2 c \succ_2 a$. Consider the following cycle, beginning with the truthful state: $(1, 1, 1, 2) \xrightarrow{1} (0, 1, 2, 2) \xrightarrow{2} (0, 0, 2, 3) \xrightarrow{1} (1, 0, 1, 3) \xrightarrow{2} (1, 1, 1, 2)$. \diamond

4 Discussion

We summarize the results in Table 3. We can see that in most cases convergence is not guaranteed unless the agents restrict their strategies to “best replies”—i.e., always select their most-preferred candidate that can win. Also, deterministic tie-breaking seems to encourage convergence more often. This makes sense, as the randomized scheme allows for a richer set of outcomes, and thus agents have more options to “escape” from the current state. Neutrality can be maintained by randomizing a tie-breaking order and publicly announcing it *before* the voters cast their votes.

We saw that if voters are non-weighted, begin from the truthful announcement and use best reply, then they always converge within a polynomial number of steps (in both schemes), but to what outcome? The proofs show that the score of the winner can only increase, and by at most 1 in each iteration. Thus possible winners are only candidates that are either tied with the (truthful) Plurality winner, or fall short by one vote. This means that it is not possible for arbitrarily “bad” candidates to be elected in this process, but does not preclude a competition of more than two candidates. This result suggests that widely observed phenomena such as Duverger’s law only apply in situations where voters have a larger amount of information regarding one another’s preferences, e.g., via public polls.

Our analysis is particularly suitable when the number of voters is small, for two main reasons. First, it is technically easier to perform an iterative voting procedure with few participants. Second, the question of convergence is only relevant when cases of tie or near-tie are common. An analysis in the spirit of [11] would be more suitable when the number of voters increases, as it rarely happens that a single voter would be able to influence the outcome, and almost any outcome is a Nash equilibrium. This limitation of our formulation is due to the fact that the behaviors of voters encompass only myopic improvements. However, it sometimes makes sense for a voter to vote for some candidate, even if this will not immediately change the outcome (but may contribute to such a change if other voters will do the same).

A new voting rule We observe that the improvement steps induced by the best reply policy are unique. If, in addition, the order in which agents play is fixed, we get a *new voting rule*—Iterative Plurality. In this rule, agents submit their full preference profiles, and the center simulates an iterative Plurality game, applying the best replies of the agents according to the predetermined order. It may seem at first glance that Iterative Plurality is somehow resistant to manipulations, as the outcome was shown to be an equilibrium. This is not possible of course, and indeed agents can still manipulate the new rule by submitting false preferences. Such an action can cause the game to converge to a different equilibrium (of the Plurality game), which is better for the manipulator.

Future work It would be interesting to investigate computational and game-theoretic properties of the new, iterative, voting rule. For example, perhaps strategic behavior is scarcer, or computationally harder. Another interesting question arises regarding possible strategic behavior of the election chairperson: can voters be ordered so as to arrange the election of a particular candidate? This is somewhat similar to the idea of manipulating the agenda. Of course, a similar analysis can be carried out on voting rules other than Plurality, or with variations such as voters that join gradually. Such analyses might be restricted to best reply dynamics, as in most interesting rules the voter strategy space is very large. Another key challenge is to modify our best-reply assumption to reflect non-myopic behavior. Finally, even in cases where convergence is not guaranteed, it is worth studying the *proportion* of profiles that contain cycles.

5 Acknowledgments

The authors thank Michael Zuckerman and Ilan Nehama for helpful discussions. R. Meir and J. S. Rosenschein are partially supported by Israel Science Foundation Grant #898/05. M. Polukarov and N. R. Jennings are supported by the ALADDIN project, jointly funded by a BAE Systems and EPSRC strategic partnership (EP/C548051/1).

References

- [1] Stéphane Airiau and Ulle Endriss. Iterated majority voting. In *Proceedings of ADT-09*, pages 38–49. Springer Verlag, 2009.
- [2] Tilman Borgers. Pure strategy dominance. *Econometrica*, 61(2):423–430, 1993.
- [3] Samir Chopra, Eric Pacuit, and Rohit Parikh. Knowledge-theoretic properties of strategic voting. Presented in JELIA-04, Lisbon, Portugal, 2004.
- [4] Vincent Conitzer and Tuomas Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proceedings of IJCAI-03*, pages 781–788. Morgan Kaufmann, 2003.
- [5] Amrita Dhillon and Ben Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46:55–75, 2004.

- [6] Robin Farquharson. *Theory of Voting*. Yale Uni. Press, 1969.
- [7] Timothy J. Feddersen, Itai Sened, and Stephen G. Wright. Rational voting and candidate entry under plurality rule. *American Journal of Political Science*, 34(4):1005–1016, 1990.
- [8] Robert Forsythe, Thomas Rietz, Roger Myerson, and Robert Weber. An experimental study of voting rules and polls in three-candidate elections. *Int. J. of Game Theory*, 25(3):355–83, 1996.
- [9] Richard D. McKelvey and Richard Niemi. A multistage representation of sophisticated voting for binary procedures. *Journal of Economic Theory*, 18:1–22, 1978.
- [10] Matthias Messner and Mattias K. Polborn. Robust political equilibria under plurality and runoff rule, 2002. Mimeo, Bocconi University.
- [11] Roger B. Myerson and Robert J. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, 1993.
- [12] Donald G. Saari. Susceptibility to manipulation. *Public Choice*, 64:21–41, 1990.

Reshef Meir
 School of Engineering and Computer Science
 The Hebrew University of Jerusalem
 Givat Ram, Jerusalem 91904, Israel
 Email: reshef24@cs.huji.ac.il

Maria Polukarov
 School of Electronics and Computer Science
 University of Southampton
 Highfield, Southampton SO17 1BJ, United Kingdom
 Email: mp3@ecs.soton.ac.uk

Jeffrey S. Rosenschein
 School of Engineering and Computer Science
 The Hebrew University of Jerusalem
 Givat Ram, Jerusalem 91904, Israel
 Email: jeff@cs.huji.ac.il

Nicholas R. Jennings
 School of Electronics and Computer Science
 University of Southampton
 Highfield, Southampton SO17 1BJ, United Kingdom
 Email: nrj@ecs.soton.ac.uk