

Convergence of Iterative Voting

Omer Lev
omerl@cs.huji.ac.il

Jeffrey S. Rosenschein
jeff@cs.huji.ac.il

School of Computer Science and Engineering
The Hebrew University of Jerusalem
Jerusalem 91904, Israel

ABSTRACT

In multiagent systems, social choice functions can help aggregate the distinct preferences that agents have over alternatives, enabling them to settle on a single choice. Despite the basic manipulability of all reasonable voting systems, it would still be desirable to find ways to reach a *stable* result, i.e., a situation where no agent would wish to change its vote. One possibility is an iterative process in which, after everyone initially votes, participants may change their votes, one voter at a time. This technique, explored in previous work, converges to a Nash equilibrium when Plurality voting is used, along with a tie-breaking rule that chooses a winner according to a linear order of preferences over candidates.

In this paper, we both consider limitations of the iterative voting method, as well as expanding upon it. We demonstrate the significance of tie-breaking rules, showing that when using a general tie-breaking rule, no scoring rule (nor Maximin) need iteratively converge. However, using a restricted tie-breaking rule (such as the linear order rule used in previous work) does not by itself *ensure* convergence. We demonstrate that many scoring rules (such as Borda) need not converge, regardless of the tie-breaking rule. On a more encouraging note, we prove that Iterative Veto does converge—but that voting rules “between” Plurality and Veto, k -approval rules, do not.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Economics, Theory

Keywords

Social choice theory, Iterative Voting, Nash Equilibrium

1. INTRODUCTION

When multiple agents have independent, perhaps differing, views over a set of alternatives, one way to decide upon an alternative is to use *social choice theory* (i.e., voting) to

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

aggregate their preferences and arrive at a common decision. However, the possibility of strategic manipulation remains a potential pitfall; the well-known Gibbard-Satterthwaite result [7, 13] states that strategic voting is potentially beneficial in any reasonable non-dictatorial voting system. Hence, analyzing elections has been complicated by the possibility that voters would attempt manipulations, and/or speculate on the actions of other players so as to try and counter-manipulate. Further complicating analysis is that effective manipulation is strongly tied to the information each player has of the game and his knowledge of the truthful preferences of other players [15]; many papers dealing with manipulation assume that all players have complete information of the game.

One approach to understanding an election is to treat it as a process, and see if we can reach some point of equilibrium, where all players are satisfied with their votes/manipulations, no longer wishing to change them. The most obvious candidate for such a stable solution would be to find a *Nash equilibrium*. However, as there may be multiple Nash equilibria in a game, many of them trivial (e.g., when all voters vote for any specific candidate), this may seem like an overly-weak option to pursue. It is, in a sense, too broad a tool to use in analyzing an election.

In previous work, Meir et al. [10] suggested the framework of iterative voting: all participants vote, and then—knowing only the result—may change their votes, one at a time, though not in a predetermined order.¹ This iterative process stops when an equilibrium is reached, when no player wishes to change his vote. A similar process can be seen, “in action”, online at various websites used to coordinate dates for an event, such as www.doodle.com; following an initial vote, every participant can change his vote. Obviously, as players change their choices one at a time, iterative voting rules are more naturally suited to a relatively small number of players, or an especially close election.

In their paper, Meir et al. [10] proved that using the simple Plurality voting rule, with a deterministic tie-breaking rule that uses a fixed linear order on candidates to break ties (and further assuming that all voters have equal weight), an iterative vote will converge to a Nash equilibrium when voters always give the best response possible to the current situation (in light of their preferences). They also showed that with weighted voters, or when using better-reply strategies (instead of best-replies), convergence is not guaranteed.

¹If they were allowed to vote simultaneously, it is easy to prove that the result may never converge to an equilibrium; and a predetermined order would just be a new voting rule.

The authors further explored nondeterministic tie-breaking rules, and showed that while they may not always converge, if the starting point is a truthful state and voters are unweighted, the game will converge.

In the current paper, we examine the robustness of this framework, as well as expanding it to encompass, beyond Plurality, an additional voting rule. We discover that when dealing with deterministic tie-breaking rules, the *type* of tie-breaking rule is crucial for a positive result: if we do not restrict the choice of tie-breaking rules, no scoring rule can guarantee convergence, and going beyond scoring rules, the Maximin voting rule is also not guaranteed to converge. Furthermore, regardless of the tie-breaking rule used, iterative voting cannot be generalized to all scoring rules, as the Borda voting rule is not guaranteed to converge under *any* tie-breaking rule. However, when using a linear-order tie-breaking rule, the iterative process with the Veto voting rule does converge when voters use the best-response strategy.² Examining if voting rules “between” Plurality and Veto (k -approval rules) are guaranteed to converge as well, we find that they are not.

1.1 Related Work

While we use the framework established by Meir et al. [10], the notions of an iterative approach to voting, as well as of seeking election equilibria, exist in previous research. An iterative process for reaching decisions was offered for agents in Ephrati and Rosenschein [5], but it uses a mechanism to transfer money-like value among agents, and hence is irrelevant to our voting procedures. Several researchers have considered reaching an equilibrium with an iterative (or dynamic) process, in particular when deciding on an allocation of public goods. A summary of much of that work can be found in Laffont [9], which details various approaches, including different equilibria choices (Nash, local dominant, local maximin) and methods. However, in order to reach an equilibrium, they limit the possible preference choices to single-peaked preferences. Another branch of research deals with a process of having a player propose a change in the current state, and hold a vote on its acceptance. Such a model was used by Shepsle [14], who chose to force an equilibrium by using a combination of preference limitation and organizational limitations. De Trenqualye [3] chose to achieve an equilibrium by using a specific voting rule and Euclidean preferences. More recently, Airiau and Endriss [1] examined—theoretically and experimentally—the possibility of an equilibrium in such games, using Plurality-type voting rules (the threshold can be different than 50% for a change to be accepted).

In searching for equilibria (albeit not iteratively), Feddersen et al. [6] chose (like Laffont) to limit preferences to single-peaked preferences. Others, like Hinich et al. [8], for example, chose to change the single-peak limitation to a specific probabilistic model of voters over a Euclidean space of candidates, while changing other parts of the model (such as allowing for abstentions). A somewhat different approach, taken by Messner and Polborn [11], analyzed equilibria by coalitional manipulation (hence, using a stronger equilib-

²Iterative vetoing is used, in the real world, in various situations, such as elimination decisions in various “reality shows” (e.g., American Idol, America’s Next Top Model, etc.). As they usually use a single judge’s preferences to break ties, it is indeed linear-order tie-breaking.

rium than Nash—a method also utilized by Dhillon and Lockwood [4]). However, one of the main limitations of many of the papers mentioned above is that they assume some knowledge of other players’ preferences.

Attempting to investigate the role of knowing other players’ knowledge, Chopra et al. [2] examined iterative voting with Plurality, and showed the effects of limiting a player’s knowledge of the other players’ preferences. Another interesting model, proposed in Myerson and Weber [12], found a Nash equilibrium for scoring rules, assuming that voters have some knowledge of which candidates have a better chance of winning (based, for example, on pre-election polls), but this does not mean that every election results in an equilibrium.

1.2 Overview of the Paper

In the following section, we give a brief overview of elections, and describe the model of iterative voting that we will be exploring throughout the paper. In Section 3 we show that the characteristics of the tie-breaking rule can affect convergence; we give examples showing that for every scoring rule, as well as for Maximin, there is a (non-linear-ordered) tie-breaking rule for which it will not always converge.

We show in Section 4 that the Borda voting rule is not guaranteed to converge, regardless of the tie-breaking rule used. After that, we limit ourselves to linear-ordered tie-breaking rules, and in Section 5 we show that using the Veto voting rule to create the Iterative Veto procedure results in a voting rule that always converges to a Nash equilibrium. However, we also show that generally, when using k -approval voting rules, they do not always converge, even when using linear-ordered tie-breaking rules. Finally, we discuss various open problems, and the issues that make them difficult (and interesting).

2. DEFINITIONS

2.1 Elections and Voting Systems

Before detailing our iterative game, we first define elections, and how winners are determined.

Definition 1. Let C be a group of m candidates, and let A be the group of all possible preference orders over C . Let V be a group of n voters, and every voter $v_i \in V$ has some element in A which is its true, “real” value (which we shall mark as a_i), and some element of A which it announces as its value, which we shall denote as \tilde{a}_i .

Note that our definition of a voter incorporates the possibility of its announcing a value different than its true value (strategic voting).

Definition 2. A voting rule is a function $f : A^n \rightarrow 2^C \setminus \emptyset$.

There are many known voting systems; one category among them is the family of scoring rules.

Definition 3. A scoring rule is a voting rule that uses a vector $(\alpha_1, \alpha_2, \dots, \alpha_{m-1}, 0) \in \mathbb{N}^m$ such that $\alpha_i \geq \alpha_{i+1}$. Each voter gives α_1 points to its first choice, α_2 points to the second, and so on. The candidates with the highest scores are the winners.

There are several well-known scoring rules.

- **Plurality:** The scoring vector is $(1, 0, 0, \dots, 0)$ —a point is only given to the most preferred candidate.
- **Veto:** The scoring vector is $(1, 1, 1, \dots, 1, 0)$ —a point is given to everyone except the least-preferred candidate.
- **Borda:** The scoring vector is $(m-1, m-2, \dots, 2, 1, 0)$ —a candidate receives points according to its preference rank.
- **k-approval (or k-veto):** The scoring vector is $(1, 1, \dots, 1, 0, 0, \dots, 0)$ —a point is given to the most preferred k candidates (or points are given to all except the least-preferred k candidates).

There are also voting systems that are not scoring rules, such as Maximin.

Definition 4. The Maximin voting rule defines for every two candidates $x, y \in V$ a score $N(x, y)$ which is the number of voters who preferred x over y . Each candidate then receives the score $S_x = \min_{y \in V \setminus x} N(x, y)$. The winners are the candidates with the maximal score.

Our definition of voting rules allows for multiple winners. However, in many cases what is desired is a single winner; in these cases, a tie-breaking rule is required.

Definition 5. A tie-breaking rule is a function $t : 2^C \rightarrow C$ that, given a set of elements in C , chooses one of them as a (unique) winner.

There can be many types of tie-breaking rules, such as random or deterministic, lexical or arbitrary. One family of tie-breaking rules that will be of interest to us is the family of linear-ordered tie-breaking rules.

Definition 6. Linear-ordered tie-breaking rules are tie-breaking rules that decide upon a winner based on some preference order over C (an element of A). Practically, this means that if $a, b \in D \subseteq C$ and $t(D) = a$, then if $a, b \in D' \subseteq C$, then $t(D') \neq b$.

While this paper does not deal with weighted games, expanding the above definitions to games that allow weighted voters is straightforward: a voter v_i with weight w_i is considered as if it were w_i different voters with the same preferences and strategy.

2.2 The Iterative Game

The following definitions and explanations follow the framework established by [10]. We do not assume that every voter knows the preferences of the others; on the contrary, we assume that each player only knows the current results (and scores) of the game, and is not aware of other voters' preferences. Hence, voters are myopic; they only think of changing their vote so as to improve the current situation, as they do not take into account future steps by other players (we also assume they are not trying to learn what their rivals' preferences are, based on their strategies).

Formally, we are viewing the election as a game, in which each player has an internal preference (a_i) , and a strategy (\tilde{a}_i) . The outcome of the game is $t(f(\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}_n))$. We

wish to find a Nash equilibrium, in which no player wishes to change his strategy, i.e., a situation in which, for any voter $v_i \in V$ and any $a'_i \in A$:

$$t(f(\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}_i, \dots, \tilde{a}_n)) \succeq_{a_i} t(f(\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}'_i, \dots, \tilde{a}_n))$$

However, we do not just want to prove that such an equilibrium exists; rather, we wish to show a process that makes that equilibrium reachable from the original starting point (which in many cases might be, due to lack of prior information, truthful on the part of the voting agents—though this is not a necessary requirement for our proofs).

Definition 7. An iterative election game G is made up of an initial election (which we shall mark as G_0), followed by further elections (G_1, G_2, \dots) , with the difference between election G_i and G_{i+1} being that one voter changed his declared preference. A game is stable if there is an n such that for all $i > n$, $G_i = G_n$.

G_0 may be a truthful state (i.e., voters vote according to their real preferences), but it is not necessarily so.

Obviously, since at every step some voter may change something about his reported preferences, no election need be stable. However, analysis becomes more interesting once we limit the voter's possible changes, requiring individual rationality. In that case, a valid step is one in which the winner of the election changes (due to the myopic steps of the players, a move that does not change the winner is pointless), and one which changes the winner to one that the voter who changed his strategy finds more preferred than the previous winner (according to the voter's internal, real preferences). Formally, a step from G_i to G_{i+1} is one in which all voters played in G_i according to the strategies $(\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}_n) \in A^n$, and for some player $v_j \in V$, there is a strategy $\tilde{a}'_j \in A$ such that:

$$t(f(\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}'_j, \dots, \tilde{a}_n)) \succ_{a_j} t(f(\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}_j, \dots, \tilde{a}_n))$$

Under the assumption of individual rationality, a stable game is one that has reached a Nash equilibrium.

Furthermore, we define a specific type of step for our analysis, following [10].

Definition 8. A *best response* step is one in which the voter changing his strategy cannot cause a more preferred candidate to win using a different strategy.

In specific voting rules below, we shall further refine what a best response move means in various circumstances.

3. USING ARBITRARY TIE-BREAKING RULES

3.1 Scoring Rules

THEOREM 1. *An iterative scoring rule game with a deterministic tie-breaking rule, even for unweighted voters that use best-response moves and start from a truthful state, does not always converge.*

PROOF. Our examples will be somewhat complex, as we deal with a large family of voting rules. In some cases using a best response is obvious, as there is only one choice that results in making a specific candidate the winner. In other cases, there may be multiple options to reach the same

winner. In the examples below, we use a “natural” definition, in which players who are taking off points from the current winner will a) give him 0 points, b) award the new winner maximal points, and c) if all things are equal, will prefer to be as close as possible to their truthful preferences. However, even if one does not use such a definition, cycles are still created—just longer cycles, as they may go through several more steps than detailed here.

First, we shall deal with the case where there are at least three candidates that do not receive maximal scores (i.e., $\alpha_{m-2}, \alpha_{m-1}, \alpha_m < \alpha_1$). We have at least four candidates, a, b, c and d . Our tie-breaking rule is that when c is tied with others, except b, c wins. When b is tied with others, except d, b wins. When d is tied with others, except a and c, d wins. We write $a \succ b \succ c$ to express that a is a voter’s favorite candidate, c is his least preferred candidate, and b is ranked in between. We have two voters:

$$\begin{aligned} \text{Voter 1: } & a \succ b \succ c \succ d \\ \text{Voter 2: } & c \succ d \succ b \succ a \end{aligned}$$

We can add several dummy candidates so the score given by voter 1 to b is less than is given to a , and the score voter 2 gives to d is less than given to c (and dummy voters, making these dummy candidates irrelevant as winner possibilities). The winner in this state is c (either he is the sole winner, or through a tie with a). The only option for improving the result for voter 1 is to make b victorious, changing his preference to $b \succ a \succ d \succ c$. Voter 2 can improve the result by changing his preference to $d \succ c \succ a \succ b$, making d the winner (possibly through winning the tie between b and d). The option available to voter 1 is to return to his original preference order, making a , his favorite, the winner. However, now voter 2 will return to his original preference as well, as it ensures the victory of c , his most preferred candidate.

If there are only two candidates that receive less than the maximal score, then we use a different game, one with six candidates. Our tie-breaking rule makes b win when a, b, c, d are tied; a wins when a, c, d are tied; c wins when a and c are tied; and c wins when a, c, d, e are tied (other ties that include d make him the winner; if they do not include d , but do include e or f , then e/f is the winner, with f triumphing over e). Let us look at two voters:

$$\begin{aligned} \text{Voter 1: } & a \succ b \succ c \succ d \succ e \succ f \\ \text{Voter 2: } & b \succ c \succ a \succ d \succ e \succ f \end{aligned}$$

The winner here is candidate b (since a, b, c, d are tied). However, when voter 1 changes his stated preference to $a \succ c \succ d \succ e \succ f \succ b$, then a , his favorite, becomes the winner (since a, c, d are tied). Voter 2 can only improve this situation by changing his stated preference to $a \succ c \succ d \succ e \succ b \succ f$, making c victorious. Voter 1 can now improve his situation by returning to his original preference, making a the winner. In this case, voter 2 will gladly return as well to his original preference, as that will make his favorite candidate, b , win.

If there is only one candidate that receives the less-than-maximal score, this is the Veto voting rule, for which there is a similar, but simpler, example. We shall use two voters, and we can describe the voting rule and tie-breaking rule fully with a table, marking the victor according to whom the voters chose to veto.

	a	b	c	d	e
a	b	c	d	e	d
b	c	d	a	a	a
c	d	a	b	a	a
d	e	a	a	a	a
e	d	a	a	a	a

In our case, the voters’ real preferences are:

$$\begin{aligned} \text{Voter 1: } & c \succ b \succ d \succ e \succ a \\ \text{Voter 2: } & b \succ d \succ c \succ e \succ a \end{aligned}$$

The truthful starting point would result in b being the winner. As voter 1 would rather that c win, he will move to veto b . Following that, voter 2 would move to veto b as well, as that would result in d winning. Voter 1, who would rather that c win, will return to vetoing a , and as voter 2 would rather that b be victorious, would return to vetoing a as well, returning to our original starting point. \square

3.2 Maximin

THEOREM 2. *An iterative Maximin game with deterministic tie-breaking, even for unweighted voters that use best-response moves and start from a truthful state, does not always converge.*

PROOF. We shall again use two voters and four candidates. The voters’ preferences are:

$$\begin{aligned} \text{Voter 1: } & c \succ d \succ b \succ a \\ \text{Voter 2: } & b \succ d \succ c \succ a \end{aligned}$$

We define the tie-breaking rule as follows: $b = c = d \Rightarrow b$; $c = b \Rightarrow c$; $a = b = c \Rightarrow b$; $a = b = c = d \Rightarrow c$; $c = d \Rightarrow c$; $b = d \Rightarrow b$. All the rest include a , and result in a being the winner.

Beginning in a truthful state, the scores of c, b and d are tied at the top, hence b is the winner. Voter 1 has no better manipulation than one that makes c victorious, and changes his preference to $c \succ b \succ d \succ a$, which evens the score of b and c , and c is the winner. Voter 2 now seeks to make b the winner, and succeeds by announcing his preferences to be $a \succ b \succ d \succ c$ (which ties, with the top score, a, b , and c). Voter 1, by returning to his original preference list, makes the score of a, b, c and d equal, resulting in c being the winner, and Voter 2 can retaliate by returning to his original preference list as well, under which b , his favorite, was victorious. \square

4. USING BORDA

Despite the significance of tie-breaking rules, there are voting rules that will not converge, regardless of the tie-breaking rule used.

THEOREM 3. *An iterative Borda game with every type of tie-breaking rule, even for unweighted voters that use best-response moves and start from a truthful state, does not always converge.*

PROOF. The example here is the same as the first example used in the proof of Theorem 1. However, the analysis in the Borda rule is much simpler, as in this case ties never occur, and hence, there is no need to rely on tie-breaking rules to achieve the necessary result: at every stage the winner will




Figure 1: The cycle of Borda non-convergence (top left is the truthful state)

have 4 points, while the other candidates will have 3 points or fewer (see Figure 1).

□

Notice that this proof stands for all scoring rules for which $m \geq 4$ and for which $\alpha_1 > \alpha_2$ and $\alpha_3 > 0$, and Borda is just one example of such a scoring rule.

5. USING VETO AND K-APPROVAL

If we confine our work to tie-breaking rules that enforce a linear order on candidates, most of our counter-examples no longer work, and convergence becomes a possibility.

5.1 Veto

Definition 9. A best response in the case of the Veto voting rule implies that the current (undesired) winner is vetoed.

THEOREM 4. *An iterative Veto game with deterministic linear-order tie-breaking and unweighted voters which use a best-response strategy, converges even when not starting from a truthful state.*

PROOF. Suppose there is an iterative election game G that includes a cycle. We shall confine our game to the cycle only, mark an arbitrary state in the cycle as G_0 , and enumerate the rest of the cycle accordingly. Note that G_0 is not necessarily the opening state of the original game.

Definition 10. $score_i(x)$ is defined as the score of candidate x in game state G_i . $max(G_i)$ is defined as the score of the winning candidate in G_i .

LEMMA 5. *If there is a cycle, then for $j < i$, $max(G_i) \leq max(G_j) + 1$, and if $max(G_i) = max(G_j) + 1$, there is only one candidate with that score.*

PROOF. By induction: for $n = 0$ the lemma is true by definition. Assuming it is true for $n' < n$, we shall prove it for n . If for some j , $max(G_{n-1}) = max(G_j) + 1$, then it is a single candidate, and therefore, the next stage will make that candidate's score go down to $max(G_j)$, and add a point to another candidate. As that candidate's score was less than $max(G_{n-1})$, its new score will be, at most, $max(G_{n-1})$, and if it is exactly $max(G_{n-1})$, it is a single candidate. If $max(G_n) < max(G_{n-1})$, then due to the induction assumption, $max(G_n) < max(G_i) + 1$ for $i < n$ (there cannot be equality, for that means $max(G_{n-1}) > max(G_i) + 1$).

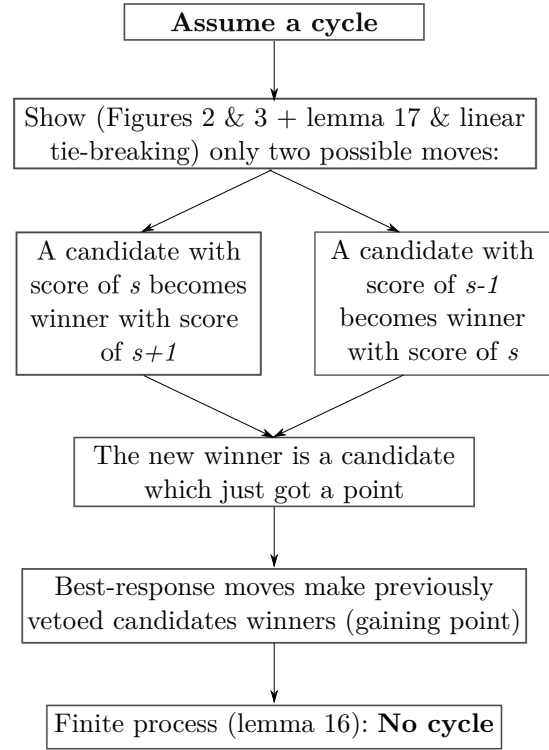


Figure 2: General overview of Veto proof

If, for every $i < n$, $max(G_{n-1}) \leq max(G_i)$, then if the candidate that gets an additional point at stage n has a lower score than $max(G_{n-1})$, this means $max(G_n) \leq max(G_{n-1})$, and the induction requirements still stand. If it has the score $max(G_{n-1})$, it becomes the single candidate with a score of $max(G_{n-1}) + 1$, and $max(G_n) \leq max(G_i) + 1$ for $i < n$. □

Notice that since we can choose G_0 arbitrarily from the cycle, due to the last lemma, $max(G_0) + 1 \geq max(G_i) \geq max(G_0) - 1$, otherwise, there will be no possibility for the cycle to return to its starting point.

LEMMA 6. *There can be at most $n \cdot (m - 2)$ consecutive steps in which the voter changed his veto from candidate a to candidate b , and candidate a became the winner.*

PROOF. Every time a voter changes his veto, he indicates that he prefers the current vetoed candidate to the current winner; that is, the winner is someone he likes less and less as the game progresses. Since there are n voters and, at most, $m - 1$ candidates that are worse than the current one, and as the voter will not vote for the very worst candidate, there are $n \cdot (m - 2)$ steps. □

We shall deal, first of all, with the easiest case, solved by the lemma above, when there is always only one candidate with the winning score (the tie-breaking rule is never used). In this case, at every step, the old winner loses a point, and the new winner gets a point. This is the case dealt with in Lemma 6, and as the number of steps is limited, there can be no cycle.

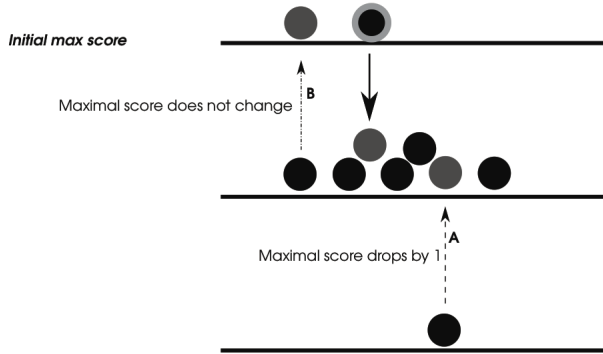


Figure 3: When only one candidate has maximal score

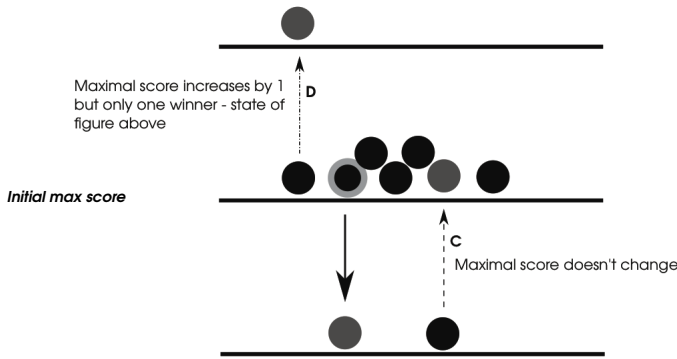


Figure 4: When multiple candidates have the maximal score

Diagrams showing why there is a limit on the increase and decrease of maximal score. When there is a state with only one candidate with maximal score, the maximal score will either remain the same with a single winner (move type A) or decrease (move type B). If it is a state where there are several candidates with the maximal score, the maximal score will either (move type D) increase while creating a situation with a single winner (which cannot increase) or the maximal score will remain the same (move type C). This also illustrates why the score cannot go down much if there is a cycle—it can only increase by one in the whole cycle; at no point can we reach a maximal score 2 points higher than another.

Having dealt with that case, let us take a closer look at G_0 , which we can define as one of the states in which there is more than one candidate with a maximal score. Note that there must be more than one of these states, since if there were a single winner in G_i and more than that in G_{i+1} , a candidate received a point and did not become a unique winner, i.e., his score in G_i was, at most, $\max(G_i) - 2$. Since this is a cycle, there must be a step in which he returns to that score (if it is G_{i+1} , then for a cycle to happen, the same candidate will need to rise again so the voter that increased his score in G_i will veto him again).

LEMMA 7. *For every state G_i in which there is more than one candidate scoring $\max(G_i)$, $\max(G_i) = \max(G_0)$, $|\{x | \text{score}_i(x) = \max(G_i)\}| = |\{x | \text{score}_0(x) = \max(G_0)\}|$ and $|\{x | \text{score}_i(x) = \max(G_i) - 1\}| = |\{x | \text{score}_0(x) = \max(G_0) - 1\}|$. This means there is always the same number of candidates with the maximal score, and with maximal score -1 . Furthermore, these are always the same candidates, switching between the two scores: $\{x | \text{score}_i(x) \in \{\max(G_i), \max(G_i) - 1\}\} = \{x | \text{score}_0(x) \in \{\max(G_0), \max(G_0) - 1\}\}$.*

PROOF. According to Lemma 5, if $\max(G_i) = \max(G_0) + 1$, there is only one candidate with the winning score, and we are not dealing with such a state. Suppose $\max(G_i) = \max(G_0) - 1$; according to the same lemma, this means there is only one candidate with the winning score in G_0 , which we defined as a state having at least two.

At any step in the game, one player loses a point and another gets it. Hence, if the number of those with the maximal score and maximal -1 score is not the same as in G_0 , some candidate lost (or gained) a point, which has a score lower than maximal -1 . However, as the maximal score will never be $\max(G_0) - 1$ (otherwise, according to Lemma 5, there would only be one candidate with winning score in G_0), there is no way in the cycle for the candidate to be vetoed when it has a score of $\max(G_0) - 1$, and get a lower score. As no candidate that has a score of $\max(G_0)$ or $\max(G_0) - 1$ can get a smaller score, the group of candidates with these scores stays fixed throughout the cycle. \square

Definition 11. Let B be the group of candidates who changed places in states in which $\max(G_i) = \max(G_0)$: $\{x | \exists i \text{ such that } \text{score}_i(x) = \max(G_0) \text{ and } \exists j \text{ such that } \text{score}_j(x) = \max(G_0) - 1\}$.

Let $z \in B$ be the lowest ranked candidate, according to the linear tie-breaking rule, in B . However, as at some state i it has a score of $\max(G_0)$, and at state j has a score of $\max(G_0) - 1$, there is a state i' in which it is vetoed and its score drops. At state i' it is the winner, meaning that all $b \in B \setminus z$ have a score of $\max(G_0) - 1$. This further means that there is always only one element of B with the score of $\max(G_0)$ (since the number of candidates with that score is constant, and candidates not in B never have a score of $\max(G_0) - 1$), and that candidate is always victorious over any other candidates with that score (since it is a part of B , it needs to be vetoed).

Thus a step during the game will either give a candidate with a score of $\max(G_0)$ a point that will make him win by giving him a score of $\max(G_0) + 1$, or give a point to a candidate with the score of $\max(G_0) - 1$ (and veto the single candidate with a score of $\max(G_0) + 1$, or, if such does not exist, a candidate in B with a score of $\max(G_0)$),

making him the winner. Hence, the voted-for candidates always become the winners, and according to Lemma 6, this is a finite process.

5.2 k -Approval, $k \geq 3$

For the k -approval voting rule, for $k \geq 3$ we prove, as for Borda, a stronger claim than for general scoring rules—we prove that even when using linear-ordered tie-breaking rules, k -approval is not guaranteed to converge.

THEOREM 8. *An iterative k -approval or k -veto game, when $k \geq 3$, with linear-ordered tie-breaking rule, even for unweighted voters that use best-response moves and start from a truthful state, does not always converge.*

PROOF. We shall provide a proof for 3-approval—it can be expanded for any larger k by adding additional dummy variables. Our tie-breaking rule is linear, with the preference: $b \succ c \succ a \succ d \succ e \succ f$. Let us assume the existence of 50 voters whose preferences are $a \succ b \succ c \succ d \succ e \succ f$. 50 others prefer $a \succ b \succ d \succ c \succ e \succ f$; 50 others prefer $b \succ c \succ d \succ a \succ e \succ f$, and 50 others $a \succ c \succ d \succ b \succ e \succ f$. So a , b , c and d have the same number of points, which is maximal—150, and e and f have 0 points. Another voter votes for $a \succ d \succ e \succ b \succ c \succ f$, and the 3 voters we will deal with have the following preferences:

Voter 1: $b \succ a \succ e \succ f \succ c \succ d$
 Voter 2: $c \succ b \succ e \succ f \succ d \succ a$
 Voter 3: $d \succ c \succ a \succ e \succ f \succ b$

Following all voters, a is the winner with 153 points. b , c and d have 152 points each, e has 3 points, and f has no points. Voter 1 realizes that he can make his favorite, b , win, by changing his vote to $b \succ e \succ f \succ c \succ d \succ a$ (a , b , c and d are now all tied with 152 points). Voter 2 now sees he can make his favorite, c , win, by changing his vote to $c \succ e \succ f \succ d \succ a \succ b$ (a , c and d are tied with 152 points, b has 151 points). At this point, voter 3 realizes he too can make his favorite the victor, by changing his vote to $d \succ e \succ f \succ b \succ c \succ a$ (so d has 152 points, a , b and c have 151 points). Now, voter 1 understands that returning to his original vote would make a the winner, which he prefers over d (now a and d are tied with 152 points). Following that, voter 2 sees that he can make b victorious by returning to his previous preference (since a , b and d will be tied with 152 points). Returning to our starting position, voter 3 sees that returning to his original vote would make a the winner, which is preferable, for him, over b . \square

6. DISCUSSION AND FUTURE WORK

An iterative voting process has a certain natural attractiveness, allowing voters to modify their stated preferences, in light of what they see about the results of an election. Assuming that all voters are equivalently entitled to make modifications, it seems an appealing way to acknowledge the strategic nature of voters, allowing them to change their votes to get results they prefer. If the process converges, we reach some stable expression of the aggregated group preference—but the process may not converge.

We began by shedding some light on the limitations of this mechanism, and on some of the elements that enable it to converge, under specific circumstances. We showed that the makeup of the tie-breaking rule is critical for iterative

voting to become a useful, converging mechanism. This is due to the basic construction of iterative voting; if ties never occur, the analysis is straightforward, either towards guaranteed convergence or not (Iterative Borda does not always converge; Iterative Plurality and Veto always will).

As ties are a significant element of what complicates the iterative convergence problem, the specific mechanism used to resolve them is part of what guarantees convergence (or lack of it): Some tie-breaking rules in certain circumstances ensure convergence; others do not. There is still much to clarify regarding this interaction between tie-breaking rules and equilibria. We have yet to establish the necessary requirements on tie-breaking rules that ensure convergence even when dealing with Iterative Plurality, let alone with other voting rules. We conjecture that requirements may be different for weighted and unweighted voting games.

However, even when eliminating considerations of tie-breaking rules, and even when we limit ourselves to scoring rules, we see that some voting rules—in fact, most of them—will never give us guaranteed convergence (such as Borda, and similar scoring rules with more than three different values). Furthermore, even if we allow ourselves to use only 1s and 0s in our scoring rules, we reach the surprising conclusion that other than the edges of this space (i.e., preference vectors where all but one element is 0, or all but one element is 1), almost no other part of this space can guarantee convergence.

This points to the basic difficulty of the iterative process—in many types of voting rules, a voter’s change of stated preference may have unintended side-effects, so when a player wishes to make a certain candidate victorious, he may be inadvertently setting the stage for another candidate to become a viable contender, to be made the winner by another player. Contrast that with Iterative Plurality or Veto, in which only one candidate benefits (and as a corollary, only one candidate is damaged). However, this is still highly dependent on tie-breaking rules, and hence there might be some tie-breaking criteria that would enable these voting rules to converge as well.

The problem of unexpected candidates becoming viable is further exacerbated with the Maximin voting rule. In that case we encounter a problem of actually defining a “best response”—no longer is there a straightforward definition of a single candidate gaining or losing points. Rather, many candidates may be affected, but which ones will be affected and made viable in the long run cannot be computed or analyzed in a simple or predictable way. While we have begun to analyze this particular non-scoring-rule voting mechanism, the iterative process for these types of voting rules remains mostly unexplored.

The non-convergence of many voting rules may suggest that it would be useful to consider different solution strategies, instead of best-response. While we know [10] that using a better-response strategy does not help, as it will not guarantee convergence even for Iterative Plurality, other solution strategies may enable convergence for a wider range of voting rules. However, we have yet to find a satisfactory voting strategy, which is both natural and ensures convergence.

The new voting rule we explored, **Iterative Veto**, despite having some superficial resemblance to Iterative Plurality, does not have the “self-reinforcing” dynamic that Iterative Plurality has, in which once a candidate has become non-viable he will never return to relevance. In Iterative Veto,

candidates' scores can increase very little from the initial stage, and when their score decreases, it may increase the number of viable candidates, making the process more turbulent than its Plurality equivalent. By making the ultimate winner potentially a candidate which was not viable at the outset, Iterative Veto enables us to reach Nash equilibria that were impossible using Iterative Plurality.

On a final note, we have not dealt with computational complexity issues here, as they were not relevant in the scenarios we considered. However, when expanding the analysis to other voting rules, such issues may arise. For example, [15] showed that finding a manipulation, even for a single manipulator in an unweighted game, is NP-complete for "ranked-pair" games such as STV and second-order Copeland. Therefore, each voter may struggle to find the step to improve his situation (and of course, struggle to find his best response).

7. ACKNOWLEDGMENTS

This work was partially supported by Israel Ministry of Science and Technology grant #3-6797, and the Google Inter-University Center for Electronic Markets and Auctions.

8. REFERENCES

- [1] S. Airiau and U. Endriss. Iterated majority voting. In F. Rossi and A. Tsoukias, editors, *Algorithmic Decision Theory*, volume 5783 of *Lecture Notes in Computer Science*, pages 38–49. Springer Berlin / Heidelberg, 2009.
- [2] S. Chopra, E. Pacuit, and R. Parikh. Knowledge-theoretic properties of strategic voting. In J. J. Alferes and J. Leite, editors, *Logics in Artificial Intelligence*, volume 3229 of *Lecture Notes in Computer Science*, pages 18–30. Springer Berlin / Heidelberg, 2004.
- [3] P. de Trenqualye. An extension of Bowen's dynamic voting rule to many dimensions. *Social Choice and Welfare*, 15(1):141–159, 1998.
- [4] A. Dhillon and B. Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46(1):55–75, January 2004.
- [5] E. Ephrati and J. S. Rosenschein. Deriving consensus in multiagent systems. *Artificial Intelligence*, 87(1-2):21–74, November 1996.
- [6] T. J. Feddersen, I. Sened, and S. G. Wright. Rational voting and candidate entry under plurality rule. *American Journal of Political Science*, 34(4):1005–1016, November 1990.
- [7] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–602, July 1973.
- [8] M. J. Hinich, J. O. Ledyard, and P. C. Ordeshook. Nonvoting and the existence of equilibrium under majority rule. *Journal of Economic Theory*, 4(2):144–153, April 1972.
- [9] J.-J. Laffont. Incentives and the allocation of public goods. In A. J. Auerbach and M. Feldstein, editors, *Handbook of Public Economics*, volume 2 of *Handbook of Public Economics*, chapter 10, pages 537–569. Elsevier, April 1987.
- [10] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria of plurality voting. In *The Twenty-Fourth National Conference on Artificial Intelligence*, pages 823–828, Atlanta, Georgia, July 2010.
- [11] M. Messner and M. K. Polborn. Strong and coalition-proof political equilibria under plurality and runoff rule. *International Journal of Game Theory*, 35(2):287–314, 2007.
- [12] R. B. Myerson and R. J. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, March 1993.
- [13] M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, April 1975.
- [14] K. A. Shepsle. Institutional arrangements and equilibrium in multidimensional voting models. *American Journal of Political Science*, 23(1):27–59, February 1979.
- [15] L. Xia, M. Zuckerman, A. D. Procaccia, V. Conitzer, and J. S. Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence, IJCAI '09*, pages 348–353, Pasadena, California, July 2009.