Journal of Artificial Intelligence Research 57 (2016) 573–591 Submitted 04/16; published 12/16

Convergence of Iterative Scoring Rules

Omer Lev OMERL@CS.TORONTO.EDU
University of Toronto, 10 King's College Road
Toronto, Ontario M5S 3G4 Canada

Jeffrey S. Rosenschein JEFF@CS.HUJI.AC.IL
Hebrew University of Jerusalem, Givat Ram
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 plausible outcomes, which are *stable* states, 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 no iterative scoring rule converges for all tie-breaking. However, using a restricted tie-breaking rule (such as the linear order rule used in previous work) does not by itself *ensure* convergence. We prove that in addition to plurality, the veto voting rule converges as well using a linear order tie-breaking rule. However, we show that these two voting rules are the only scoring rules that converge, regardless of tie-breaking mechanism.

1. Introduction

Any system—mechanical or human—that is composed of different elements, employs some mechanism of aggregating the opinions of its components in order to reach a joint decision or action. These mechanisms can be thought of as elections, in which each participant ranks several options; after receiving each agent's preference order over these options, the mechanism picks a specific option as its result.

People have long been aware that sometimes, by misreporting their preference order, the outcome may become more favorable to them ("strategic/tactical voting"). Such phenomena have bothered social choice theorists, such as Condorcet and Borda, as democracy and voting became more common in the late 18th century. However, efforts to develop voting methods that would encourage voters to be truthful were revealed to be futile, once the Gibbard-Satterthwaite theorem (Gibbard, 1973; Satterthwaite, 1975) showed that no non-dictatorial "reasonable" voting rule can be strategyproof. This result led to a line of research commenced by Bartholdi, Tovey, and Trick (1989), Bartholdi and Orlin (1991), and expanded further by many researchers (Conitzer, Sandholm, & Lang, 2007; Procaccia, Rosenschein, & Zohar, 2007; Procaccia & Rosenschein, 2007; Xia, Zuckerman, Procaccia, Conitzer, & Rosenschein, 2009; Zuckerman, Procaccia, & Rosenschein, 2008; Zuckerman, Lev, & Rosenschein, 2011), which assess the complexity of finding a beneficial manipulation for voters in various settings. See also overviews in the work of Rothe (2015) and Brandt, Conitzer, Endriss, Lang, and Procaccia (2016).

Despite significant results from this line of research, it has not helped us in analyzing election results. When *many* voters are strategic and manipulating, we are still at a loss to understand which election results are possible and which are not, not to mention trying to understand underlying truthful preferences when presented with election results.[1] Furthermore, while game theory is often a powerful tool when analyzing situations where participants manipulate, voting scenarios have some properties that make them more difficult to analyze: participants only supply a ranking of their preferences, rather than the utility value of each option (as is typical in monetary settings), and the assumption that participants have a good perception of others' preferences is not true. Looking at equilibria in voting games results in an enormous number of equilibria (hundreds of thousands even in small games, such as one with 10 voters and 5 candidates; see Thompson, Lev, Leyton-Brown, & Rosenschein, 2013). Much worse from the viewpoint of analysis, many Nash equilibria are practically useless for any analytic purpose, not representing a realistic possible end-state of an election. For example, even if all voters rank the same candidate in last place, there are Nash equilibria where it wins (e.g., in plurality, if all voters vote for this candidate, this is a Nash equilibrium). This renders standard Nash equilibria virtually useless as an analytic tool, despite their fundamental relevance in showing stable states that incorporate voters' manipulations.

One approach to trying to model the election process and outcome is *iterative voting*, suggested by Meir, Polukarov, Rosenschein, and Jennings (2010). This work views decision-making as a process by which voters change their vote according to what they know of the result; for example, if by changing their vote they can alter the outcome to one that is more preferable for them, they will do so. This process not only happens in smaller groups (e.g., a group of friends choosing a restaurant, or using event coordination sites, such as Doodle), but can also be used to model larger-scale elections, where people's knowledge of the state of the election comes from polls, according to which they adapt their vote (some simulation-based support for such a view, in a slightly more complex model, was shown in Meir, Lev, & Rosenschein, 2014). A voter in a plurality election (e.g., for U.S. Congress or U.K. Parliament) that supports a party that has no chance of winning may switch their vote to the major party that they support the most, in order to influence the election outcome. Looking at the set of states that are reachable from true preferences, using this dynamic, gives us (in some sense) a set of plausible outcomes. These states are far fewer than the number of Nash equilibria, and may give us a better understanding of realistic outcomes.

However, the results in the work of Meir et al. (2010) were limited: the process was shown to converge if voters are myopic and pursue at each stage the most beneficial strategy that they can, while the voting system is plurality with a linear tie-breaking rule. While the assumptions on voters were justified—the necessity of the best-response strategy was shown, and it is not clear how to remove the myopic assumption (though some recent work has

1. This problem is only exacerbated when considering independent computer agents. They are motivated to manipulate, and may have the computational resources to calculate the particular strategy to do so optimally.

tried to address this problem)—there is no obvious reason to assume plurality with a linear tie-breaking rule as the underlying voting rule.

In this paper, we examine the robustness of the *iterative voting* 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. We prove that when using a linear-order tie-breaking rule, the iterative process with Veto does converge when voters use the best-response strategy.[2] However, we end with a negative result, showing that regardless of the tie-breaking rule used, iterative voting cannot be generalized to all scoring rules: no scoring rule other than plurality and veto will converge, including the family of voting rules that lies "between" them (k-approval). This negative result was expanded to several other non-scoring rules in the work of Koolyk, Lev, and Rosenschein (2016).

2. Related Work

The common game-theoretic approach to iterative processes is subgame-perfect Nash equilibrium, which entails a backward induction from potential results. This has mostly been used to analyze open serial elections (e.g., roll-calls) and sequential elections with two candidates (Sloth, 1993; McKelvey & Niemi, 1978; Dekel & Piccione, 2000). However this is not well-attuned to reality, as people are notoriously bad at doing backward induction (Johnson, Camerer, Sen, & Rymon, 2002), and in our iterative setting, with its limited information about other voters' preferences and its non-deterministic voting order, it is even less applicable and realistic.

While we use the framework established by Meir et al. (2010), the notions of an iterative approach to voting, as well as of seeking election equilibria, exist in prior research. Dealing with general game structure, and therefore including mixed strategies (which are less relevant in a voting setting), fictitious play was suggested by Brown (1951). It includes agents reacting with a best response strategy to one another, and has been shown to converge for particular types of games (e.g., zero-sum games Robinson, 1951). In decision making settings, an iterative process for reaching decisions was offered for agents by Ephrati and Rosenschein (1996), but it used a mechanism to transfer money-like value among agents, and hence is not relevant to 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 the work of Laffont (1987), 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. More recently, Reijngoud and Endriss (2012) explored an iterated poll setting, but as theirs had only one round where voters change their vote according to a fixed order, the issues discussed here do not arise.

2. Iterative 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.

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 (1979), where an equilibrium was forced by using a combination of preference limitation and organizational limitations. A different approach was used by de Trenqualye (1998), who chose to achieve an equilibrium by using a specific voting rule and Euclidean preferences. More recently, Airiau and Endriss (2009) 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).

Attempting to investigate the role of knowing other players' knowledge, Chopra, Pacuit, and Parikh (2004) examined iterative voting with plurality, and showed the effects of limiting a player's knowledge of the other players' preferences. Another interesting model, proposed by Myerson and Weber (1993), 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); this does not, however, mean that every election results in an equilibrium.

In searching for equilibria (albeit not iteratively), Feddersen, Sened, and Wright (1990) chose (like Laffont, 1987) to limit preferences to single-peaked preferences. Others, like Hinich, Ledyard, and Ordeshook (1972), 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 (2007), analyzed equilibria by coalitional manipulation (hence, using a stronger equilibrium than Nash, a method also utilized in Dhillon & Lockwood, 2004). However, one of the main limitations of many of the papers mentioned above is that they assume some player knowledge of other players' preferences.

2.1 Iterative Voting

The iterative model presented by Meir et al. (2010) has attracted much interest, aiming to expand the model in different directions. Some work has attempted to understand the need for the "best response" strategy requirement, and try to generalize it, which was analyzed by Obraztsova, Markakis, Polukarov, Rabinovich, and Jennings (2015c) (and, in a different way, yet still applicable to the iterative model, in Meir, 2015). Looking into strategies that might ensure convergence in an iterative voting process (and not just best-response) was also explored by Grandi, Loreggia, Rossi, Venable, and Walsh (2013), and by Loreggia (2012). A recent attempt to address the myopic limitation has been done by Obraztsova, Lev, Polukarov, Rabinovich, and Rosenschein (2015b), but it is in its early stages, and is the least-explored direction.

A different approach has been to change the voters' behavior in the iterative process, and make them biased: truth-biased (i.e., all things being equal, prefer truthful vote) or lazy-biased (all things being equal, prefer not participating). This was explored by Rabinovich, Obraztsova, Lev, Markakis, and Rosenschein (2015), which also showed that finding if a state is reachable iteratively with regular voters is NP-complete. In a different approach, Brânzei, Caragiannis, Morgenstern, and Procaccia (2013) addressed some notions of the price of anarchy in the context of iterative plurality, showing that winners will have a very high truthful score.

A different model using an iterative concept is by Meir et al. (2014) (and, continuing it, Meir, 2015), but it assumes that players are not sure of what the state of the world is (though their simulations, even for our stricter setting, seem to vindicate the general iterative approach). While that model encompasses, in some respects, the iterative voting model as a particular case, it has so far only explored plurality.

Moving beyond plurality, Obraztsova, Lev, Markakis, Rabinovich, and Rosenschein (2015a) examine the complexity of knowing if a particular profile is reachable from the truthful state via an iterative process using plurality and veto, and show this is NP-complete, while also showing a particular family of problems where this is polynomial for veto. Grandi et al. (2013) (and Loreggia, 2012) approach non-scoring rules, but only show convergence there with a limited set of strategies, but not best-response. Recently, Koolyk et al. (2016) showed non-convergence in various strategies, including best-response, in various non-scoring rules.

We note that Meir et al. (2010) also showed various non-convergence results for plurality when the iterative stage is simultaneous, and when tie-breaking is non-deterministic. Also, after the publication of the proof below on iterative veto, another proof, very different in structure, for the same result was published by Reyhani and Wilson (2012). They also showed that convergence for veto occurs with complexity $\mathcal{O}(mn)$.

3. Preliminaries

We first set out to define elections. An election $\mathcal{E}$ is composed of several elements:

- *Candidates*: A set C of m elements.
- *Voters*: A set V of size n. Each voter $i \in V$ is associated with an element $a_i \in \pi(C)$, where $\pi(C)$ is the set of linear orders of the elements of C. This represents the truthful preferences of a particular voter. To make this similar to common game-theoretic settings, we can associate this preference order with some utility function $u_i : C \to \mathbb{R}$, in which there is a higher utility from higher-ranked candidates, but this utility function is hidden from anyone but the voter itself. We mark this preference order by $\succ_i$.
- *Voting function*: A function $f : \pi(C)^n \to C$ which is given a set of preferences from each voter and outputs a winner.

A voting function $f : \pi(C)^n \to C$ consists of two parts: a voting rule $\hat{f} : (\pi(C))^n \to 2^C$ itself, and a tie-breaking rule $t : 2^C \to C$. Thus, $f = t \circ \hat{f}$.

Definition 1. *A* tie-breaking rule *is a function $t : 2^C \to C$, which takes a subset of C and returns a single candidate. A particular family of tie-breaking rules is of lexicographic or* linear tie-breaking rules *which are defined by an order $a \in \pi(C)$, and for $\hat{C} \subseteq C$, $t(\hat{C})$ is $c \in \hat{C}$ which is the highest ranked member in $\hat{C}$ according to the order a (i.e., for all $c' \in \hat{C}$, $c' \neq c$, $c \succ_a c'$).*

As a voter may be strategic, it may report to the voting rule a different preference than its truthful $\succ_i$ preference order. That is, each voter i is reporting to the voting rule

$b_i \in \pi(C)$, and b_i might not equal a_i. A set of all such n declared votes is a *voting profile*, which we shall generally denote by $\mathbf{b}$.

Definition 2. *A* strategic deviation *by voter i in a profile $\mathbf{b}$ in which each voter $j \neq i$ is reporting a vote $b_j \in \pi(C)$, is when there is a preference order $\hat{b}_i \in \pi(C)$, $\hat{b}_i \neq b_i$ such that*

$$f(b_1, \ldots, \hat{b}_i, \ldots, b_n) \succ_i f(b_1, \ldots, b_i, \ldots, b_n)$$

We denote the profile $\mathbf{b}$ without the vote of voter $i \in V$ as $\mathbf{b_{-i}}$.

The Gibbard-Satterthwaite theorem tells us that all reasonable voting rules have a profile for which there are voters that have a strategic deviation from the truthful profile.

Definition 3. *A profile $\mathbf{b}$ is a* Nash equilibrium *when no voter has a strategic deviation from the profile. So for any voter i,*

$$f(b_1, \ldots, b_i, \ldots, b_n) \succeq_i f(b_1, \ldots, \hat{b}_i, \ldots, b_n)$$

for any $\hat{b}_i \in \pi(C)$.

We use the terminology of Nash equilibria in voting scenarios to refer only to pure Nash equilibria, as mixed Nash equilibria have little practical meaning in such scenarios.

3.1 Voting Rules

Many voting rules have been suggested in the literature, though we shall focus only on a few of them. In particular, we will pay attention to a particular family of voting rules:

Definition 4. *A* scoring rule *is a voting rule defined by a vector*

$$(\alpha_1, \alpha_2, \ldots, \alpha_{m-1}, 0)$$

in which $\alpha_1 \geq \alpha_2 \geq \alpha_{m-1} \geq 0$. Each voter's vote contributes α_1 points to the highest ranked candidate in its reported preference, α_2 points to the second ranked candidate, and so on. Finally, scores are tallied after all votes have been cast, and the candidates that scored the maximal number of points are the winners.

Many commonly used voting rules are scoring rules. In particular, we shall focus on:

Plurality The most widely-used voting rule. Equivalent to the scoring rule $(1, 0, \ldots, 0)$.

Veto A voting rule complementary to plurality. Equivalent to the scoring rule $(1, 1, \ldots, 1, 0)$.

k-approval/k-veto Lying between plurality and veto, in which we fix the number of candidate getting a score of 1, or the number of those getting a score of 0. Equivalent to the scoring rule $(1, \ldots, 1, 0, \ldots, 0)$.

We shall discuss these voting rules as families, including their variants for any number of candidates. So, for example, when discussing the veto voting rule, we refer to using the veto scoring rule for any number of candidates m. However, our proofs apply to each of these separately, and therefore, for example, any scoring rule that incorporates the plurality scoring rule for a particular number of candidates m will converge for m candidates (though not necessarily for a different number of candidates, if they use a different scoring rule for that number).

3.2 The Iterative Process

The following definitions and explanations follow the framework established by Meir et al. (2010). 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 candidate 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 wish to show a process that examines equilibria 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 5. *An iterative election game G is made up of an initial election (which we shall mark as G_0), followed by further elections ($G_1, G_2, \ldots$), with the difference between election G_i and G_{i+1} being that one voter changed their declared preference. A game converges if it becomes stable, i.e., 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 true preferences), but it is not necessarily so.

Obviously, since at every step some voter may change something about their reported preferences, no election need be stable. However, the 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,[3] and the step must be one that changes the winner to a candidate that the voter who changed their vote prefers over 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 $(b_1, b_2, \ldots, b_n) \in \pi(C)^n$, and for some player $j \in V$, there is a strategy $b'_j \in \pi(C)$ such that:

$$t(f(b_1, b_2, \ldots, b'_j, \ldots, b_n)) \succ_{a_j} t(f(b_1, b_2, \ldots, b_j, \ldots, b_n))$$

Under the assumption of individual rationality, and thanks to voters not changing votes simultaneously,[4] a stable game is one that has reached a Nash equilibrium.

Furthermore, we define a specific type of step for our analysis (following Meir et al., 2010).

Definition 6. *A* best response *step is one in which the voter changing their strategy cannot cause a more preferred candidate to win using a different strategy, i.e., given a strategy profile $(b_1, b_2, \ldots, b_n) \in \pi(C)^n$, strategy $\bar{b}_i \in \pi(C)$ is the best response for player i if for every strategy $b'_i \in \pi(C)$*

$$t(f(b_1, b_2, \ldots, \bar{b}_i, \ldots, b_n)) \succ_{a_i} t(f(b_1, b_2, \ldots, b'_i, \ldots, b_n))$$

In specific voting rules below, we shall further refine what a best response move means in various circumstances. In particular, Meir et al. (2010) defined a best response in plurality to mean the voter is voting for the candidate they are making the winner, and proved:

3. Due to the myopic steps of the players, a move that does not change the winner is pointless.
4. Meir et al. (2010) showed that allowing simultaneous updates prevents convergence even for plurality.

Theorem 1 (Meir et al., 2010 Thm. 3)**.** *An election conducted using iterative plurality with a linear order deterministic tie-breaking rule, when voters are myopic and using best-response strategies, will always converge.*

4. Tie-Breaking Rules

We start our analysis of the iterative voting model by examining if the requirement for linear-ordered tie-breaking rules is necessary. Indeed, we show that using arbitrary tie-breaking rules is not just harmful to the convergence of plurality, but that it prevents convergence for any scoring rule.

We shall be dealing in our proof with some number of candidates, so we first wish to ensure any proof on non-convergence in more cases.

Lemma 1. *If a scoring rule $\alpha = (\alpha_1, \ldots, \alpha_{m-1}, 0)$ does not converge, then for every $m' > m$, and scoring rule $\beta = (\beta_1, \ldots, \beta_{m'-1}, 0)$ such that $\beta_1 = \alpha_1$ and for every α_i there is a j such that $\alpha_i = \beta_j$, β does not converge as well.*

Proof. Suppose the example showing that α does not converge requires a voters. To construct a similar example for β, we first select m candidates. Now, let us take $m!$ voters. Each voter has a different permutation of m candidates in the first m rankings (so, after all $m!$ voters are tallied, all m candidates will have the same score), and then some permutation of the remaining $m' - m$ candidates, but with only a specific one of them in the last place. Now, we create another set of $m!$ voters, and this time, a different candidate of the $m' - m$ is put in the last place for all of them. Overall, $(m' - m)m!$ voters are created, with each of the m voters having an identical score, and each of the $m' - m$ candidates has a score that is lower by $m!\alpha_{m-1}$ than the m candidates (if this is not more than α_1 we repeat the process several times until it is). Finally, we repeat this process a times, so that there is no way for a voters to make any of those $m' - m$ voters a viable winner.

Now, we add the a voters from the α example, placing our m candidates in an order so that they get the scores they get in the α example, and the rest of the $m' - m$ candidates fill in the rest of the ranking. The rest of the iterative process happens by those a voters iteratively doing what they did in the α example. □

We now turn to the main theorem. Our examples will be somewhat complex, as we deal with a large family of voting rules. In some cases the best response strategy is obvious, as there is only one choice that results in making a particular candidate the winner. In other cases there may be multiple options to reach the same outcome, hence we used a "natural" definition for scoring rules, in which players taking off points from the current winner will give it zero points and award the new winner the maximal score possible. Also, all things being equal, we assume that voters will prefer to be as close as possible to their truthful preferences. However, even if one does not use such a definition for best response, cycles are still created—but they might be longer ones, as they may go through several more steps than detailed here.

Theorem 2. *An election conducted using an iterative scoring rule with an arbitrary deterministic tie-breaking rule, when voters are myopic and using best-response strategies, will not converge for some preferences even when starting from the truthful state.*

Proof. First, we deal with scoring rules in which at least three candidates do not receive maximal scores (i.e., $\alpha_{m-2} < \alpha_1$). We have at least four candidates, a, b, c, and d, and 2 voters:

$$\text{Voter 1: } a \succ b \succ c \succ d$$
$$\text{Voter 2: } c \succ d \succ b \succ a$$

Our tie-breaking rule needs to include the following, for ties of two candidates:

- c wins when tied with others, except b.
- b wins when tied with others, except d.
- d wins when tied with others, except a or c.

We can add several dummy candidates so that 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 potential winners). The winner in this truthful state is c (either the sole winner, or through a tie with a). Voter 1's best response is to make b victorious, changing its preference to $b \succ a \succ d \succ c$. Voter 2 can improve the result by changing its preference to $d \succ c \succ a \succ b$, making d the winner (possibly through winning the tie between b and d). The best option available to voter 1 is to return to its original preference order, making a, its favorite, the winner. However, now voter 2 will return to its original preference as well, as it ensures the victory of c, its own most-preferred candidate.

If there are only two candidates that receive less than the maximal score, then we use a different setting, one with six candidates. Our tie-breaking rule follows:

- b wins when a, b, c, d are tied.
- a wins when a, c, d are tied.
- c wins when a and c are tied and when a, c, d, e are tied.
- d wins in other ties that include it.
- f wins in other ties that do not include d.
- e wins in other ties that do not include d or f.

Let us look at two voters:

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

The winner here is candidate b (since a, b, c, d are tied). However, when voter 1 changes its stated preference to $a \succ c \succ d \succ e \succ f \succ b$, then a, its favorite, becomes the winner (since a, c, d are tied). Voter 2 can only improve this situation by changing its stated preference to $a \succ c \succ d \succ e \succ b \succ f$, making c victorious. Voter 1 can now improve the

	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

Table 1: Non-linear tie-breaking rule for veto

situation by returning to its original preference, making a the winner. In this case, voter 2 will also return to its original preference, as that will make its 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 using Table 1, marking the victor according to whom the voters chose to veto.

In our case, the voters' real preferences are:

$$\text{Voter 1: } c \succ b \succ d \succ e \succ a$$
$$\text{Voter 2: } b \succ d \succ c \succ e \succ a$$

The truthful starting point would result in b being the winner. As voter 1 would rather that c win, it 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, which would rather that c win, will return to vetoing a, and as voter 2 would rather that b be victorious, it would return to vetoing a as well, returning to our original starting point. □

5. Veto

The veto voting rule is the only scoring rule, besides plurality, that converges under the iterative mechanism (as will be shown below). In the work on iterative plurality, it was assumed that in a best-response move, the voter votes for the candidate that becomes the winner (because it is always a best-response). Similarly, for veto, we assume that a best-response move is one in which a candidate vetoes the current winner; no move can get a better result.

Observation 1. *A* best response *in the case of the Veto voting rule is equivalent to the current (undesired) winner being vetoed.*

Now we can prove convergence, using a proof by contradiction.

Theorem 3. *Iterative Veto elections with deterministic linear-order tie-breaking and voters that use a best-response strategy, converge even when not starting from a truthful state.*

Proof. Suppose there is an iterative election G that includes a cycle. We shall 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 election.

Definition 7. *$score_i(x)$ 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 2. *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. Proving by induction, the base case is trivial. Assuming it is true after $h - 1$ steps, proving it for step h: examining G_{h-1}, if there was a $j < h - 1$ for which $max(G_{h-1}) = max(G_j) + 1$, there is a single winner in G_{h-1}, which loses a point, and therefore the winner in G_h will have, at most, $max(G_{h-1})$ points. Thus, $max(G_h) \leq max(G_{h-1})$, and the claim stems from its truth for G_{h-1}.

If for every $j < h - 1$ $max(G_{h-1}) \leq max(G_j)$, the maximal score in G_h will rise by at most one point, i.e., $max(G_h) \leq max(G_j) + 1$ for all $j < h$. Furthermore, if it indeed grows, there is only a single candidate with that number of points (as only the candidate that got an extra point has this score). If the maximal score in G_h did not grow, $max(G_h) \leq max(G_{h-1})$, and the claim is true from its correctness for G_{h-1}. □

Notice that since we can choose G_0 arbitrarily from the cycle, due to the previous 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 3. *There can be at most $n \cdot (m - 2)$ consecutive steps in which the voter changed their veto from candidate a to candidate b, and candidate a became the winner.*

Proof. Every time a voter changes its vetoed candidate, it indicates that the previously vetoed candidate is preferable to the current winner; that is, the winner is a candidate less and less liked as the game progresses. So, if a voter's currently vetoed candidate is $c \in C$, and the current winner is $c' \in C$, the voter indicates $c \succ c'$ by their move. Later, in their next move, they prefer making $c' \in C$ the winner over the current one, $c'' \in C$, indicating $c' \succ c''$. When this continues, we pass more and more candidates as winners, ranked lower and lower for each voter. Since there are n voters and, at most, $m - 1$ candidates that are worse than the current one, and as the voter will not choose to make the very worst candidate the winner, 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 gains a point. This is the case dealt with in Lemma 3, and as the number of steps is limited, there can be no cycle.

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 was 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., its score in G_i was, at most, $max(G_i) - 2$. Since this is a cycle, there must be a step in which it returns to that score, i.e., where the maximal score is $max(G_i) - 1$ (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 this candidates's score in G_i will veto it again).

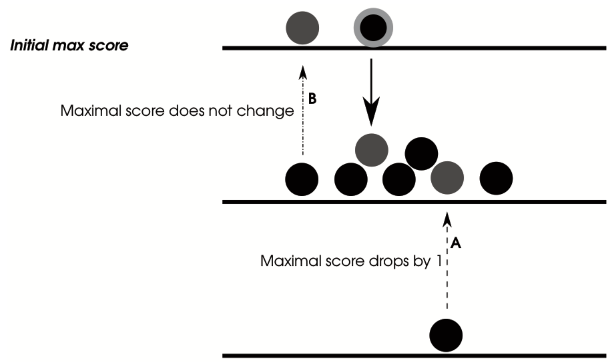

Figure 1: When only one candidate has maximal score

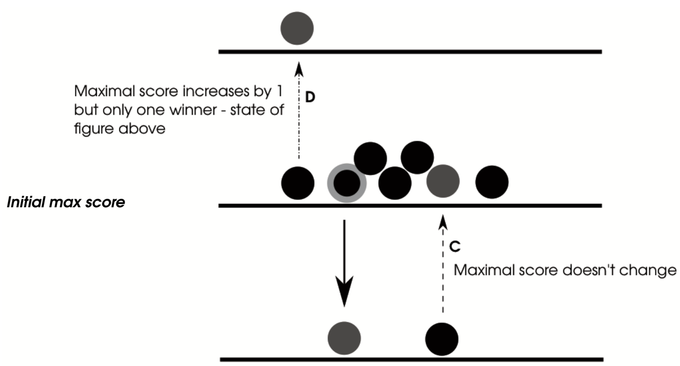

Figure 2: When multiple candidates have the maximal score

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 B) or decrease (move type A). If it is a state where there are several candidates with the maximal score, the maximal score will either increase (move type D) while creating a situation with a single winner 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.

Lemma 4. *For every state G_i in which there is more than one candidate scoring $max(G_i)$, $max(G_i) = max(G_0)$, $|\{x \mid score_i(x) = max(G_i)\}| = |\{x \mid score_0(x) = max(G_0)\}|$ and $|\{x \mid score_i(x) = max(G_i)-1\}| = |\{x \mid score_0(x) = max(G_0)-1\}|$. This means the number of candidates with the maximal score remains fixed, as does the number of candidates with maximal score -1. Furthermore, these are the same candidates, switching between the two scores: $\{x \mid score_i(x) \in \{max(G_i), max(G_i)-1\}\} = \{x \mid score_0(x) \in \{max(G_0), max(G_0)-1\}\}$.*

Proof. According to Lemma 2, if $max(G_i) = max(G_0) + 1$, there is a single candidate with the winning score, and this lemma does not handle this case. 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 candidate loses a point and another gains 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 in G_0. However, as the maximal score will never be $max(G_0) - 1$ (otherwise, according to Lemma 2, 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 set of candidates with these scores stays fixed throughout the cycle. □

Let B be the set of candidates who changed places in states in which $max(G_i) = max(G_0)$ (i.e., $B = \{x \mid \exists i$ such that $score_i(x) = max(G_0)$ and $\exists j$ such that $score_j(x) = max(G_0) - 1\}$). Let $z \in B$ be the lowest-ranked candidate according to the linear tie-breaking rule in B. Since z changes its score, there is a state G_i where z has the score $max(G_0)$ and is vetoed, i.e., z is the winner if G_i. This means there is no other candidate from B with the score $max(G_0)$. As the number of candidates with $max(G_0)$ does not change (according to Lemma 4), this means that at every state G_j in which $max(G_j) = max(G_0)$, there is only a single candidate from B with $max(G_0)$ points, and it always wins (due to the tie-breaking rule, since z wins). This means the candidate getting the point at every stage is the one that becomes the winner—which, as noted in Lemma 3, is a finite process, contradicting the endless cycle. □

6. Other Scoring Rules

Finally, having shown that iterative veto converges (and knowing that iterative plurality converges from Meir et al., 2010), we are ready to show that the iterative process will not converge for any other scoring rule, regardless of the tie-breaking rule it uses.

Theorem 4. *Under the iterative procedure, using a best response strategy and when voters are myopic, no scoring rule apart from plurality and veto converges.*

Proof. Different rules require different proofs, and we consequently divide our proof:

Part 1: Scoring Rules with 2 Values (k-approval)

Scoring rules with only 2 values which are not plurality or veto are equivalent to k-approval for $k > 1$ and $k < m - 1$. We shall show it for $k = 2$ and $m = 4$. Using dummy candidates, this can be extended to any size of k.

Let a be the candidate which a tie-breaking rule selects when all 4 candidates are tied. The voters are:

$$\text{Voter 1: } a \succ b \succ c \succ d$$
$$\text{Voter 2: } d \succ b \succ c \succ a$$

The winner in this case is b, with a score of 2. But voter 1 can change its vote to $a \succ c \succ d \succ b$, making a the winner (thanks to the tie-breaking). But now voter 2 can change to $c \succ d \succ b \succ a$, making c the winner. Now, voter 1, by reverting to its truthful preference makes a the winner again (thanks to tie-breaking rules), and, finally, by reverting to its true preference, voter 2 returns us to the original state, making b the winner and creating the cycle.

Part 2: Scoring Rules with 3 Values

Oddly, this is the most complicated case. Our scoring rule has the values α_1, α_2 and 0. Let us first examine the case when the scoring rule for 4 candidates is $(\alpha_1, \alpha_2, 0, 0)$ and $\alpha_1 < 2\alpha_2$. Any tie-breaking rule has a candidate a for which there are at least 2 other candidates, b and c, which it dominates when tied with them. The voters are:

$$\text{Voter 1: } b \succ d \succ c \succ a$$
$$\text{Voter 2: } a \succ d \succ c \succ b$$

Hence the winner is d, with a score of $2\alpha_2$. Voter 2 can now change its vote to $a \succ c \succ b \succ d$, making a the winner with α_1 points (as d changes to α_2). Now voter 1 can change to $c \succ b \succ d \succ a$, making c the winner, with $\alpha_1 + \alpha_2$ points. By reverting to its truthful vote, voter 2 can make its favorite candidate, a, the winner again, to which voter 1 retaliates by returning to its truthful vote as well, creating a cycle, with d as the winner.

We now create a cycle for the case $\alpha_1 \geq 2\alpha_2$. Our tie-breaking rule is such that candidate a dominates candidate c. The voters are:

$$\text{Voter 1: } a \succ c \succ d \succ b$$
$$\text{Voter 2: } d \succ c \succ b \succ a$$

This means a is the winner with α_1 points, thanks to the tie-breaking rule. Voter 2 now changes to $c \succ d \succ b \succ a$, making c the winner with $\alpha_1 + \alpha_2$ points. Voter 1 retaliates by changing to $a \succ d \succ b \succ c$, making a the winner again. By returning to its truthful preference, voter 2 makes the winner d, its favorite candidate, with $\alpha_1 + \alpha_2$ points. However, by completing the cycle and returning to its truthful preference, voter 1 makes its favorite, a, the winner again.

We now look at the case where the scoring rule for 4 candidates is $(\alpha_1, \alpha_2, \alpha_2, 0)$. Any tie-breaking rule has a candidate (which we name a) which when tied with d, wins, and a candidate b that when tied with candidates c or d it wins as well. The voters are:

$$\text{Voter 1: } c \succ d \succ b \succ a$$
$$\text{Voter 2: } a \succ b \succ d \succ c$$

The winner is either a or b (depending on if $\alpha_1 < 2\alpha_2$ or not). Voter 1 reacts by voting $d \succ c \succ b \succ a$, making d the winner with $\alpha_1 + \alpha_2$ points, to which voter 2 reacts by voting $a \succ b \succ c \succ d$, making a or b the winner. However, by reverting to its truthful preference, voter 1 makes c the winner with a score of $\alpha_1 + \alpha_2$. Voter 2 completes the cycle by reverting to its truthful preference.

The last case is when the scoring rule for 4 candidates is $(\alpha_1, \alpha_1, \alpha_2, 0)$. Let a be a candidate that loses in ties against 2 other candidates, b and c (every tie-breaking rule has such a candidate). The voters are:

Voter 1: $a \succ b \succ c \succ d$
Voter 2: $a \succ c \succ b \succ d$
Voter 3: $a \succ b \succ c \succ d$
Voter 4: $a \succ c \succ b \succ d$
Voter 5: $d \succ b \succ a \succ c$
Voter 6: $d \succ c \succ a \succ b$
Voter 7: $a \succ c \succ b \succ d$
Voter 8: $d \succ b \succ c \succ a$

The winner is a with $5\alpha_1 + 2\alpha_2$ points, which is changed to c (with $5\alpha_1 + 2\alpha_2$ points, using the tie-breaking rule) when voter 8 changes to $c \succ d \succ b \succ a$. Now voter 7 makes a the winner again when it votes $a \succ b \succ d \succ c$. Voter 8 now returns to being truthful, making b the winner, with $5\alpha_1 + 2\alpha_2$ points (and using the tie-breaking rule), causing voter 7 to revert to being truthful, making a the winner again.

Part 3: Scoring Rules with More Values

We shall show this on voting rules with 4 values and 4 candidates. This is easily extendable to more candidates and values by adding dummy candidates (and voters who do not get to manipulate) as necessary.

The scoring rule is $(\alpha_1, \alpha_2, \alpha_3, 0)$ and is strictly monotonic. The voters are:

Voter 1: $a \succ b \succ c \succ d$
Voter 2: $c \succ d \succ b \succ a$

Candidate c is the winner here with $\alpha_1 + \alpha_3$ points (a has only α_1 points; b has $\alpha_2 + \alpha_3$; and d has α_2). Voter 1 can change the outcome by changing its vote to $b \succ a \succ d \succ c$ reducing c to only α_1 points, and raising b to $\alpha_1 + \alpha_3$. Voter 2 retaliates by changing to $d \succ c \succ a \succ b$, making d the winner with $\alpha_1 + \alpha_3$ points, and reducing b to α_1. Voter 1 improves its situation by reverting to its truthful preferences, which make a the winner (with $\alpha_1 + \alpha_3$ points). Voter 2 completes the cycle by also reverting to its truthful preferences, making its favorite candidate, c, the winner. □

7. Conclusion and Discussion

Since the iterative model was introduced for plurality by Meir et al. (2010), many extensions, additions, and further models have been based on it. Almost all of this work has focused on the plurality voting rule, which was there shown to converge. However, relatively little work has gone into examining the mechanism itself. In this work, we examined some varied voting systems that can be used with the iterative dynamic. While for the tie-breaking rule we show that the restriction on linear-ordered rules (for deterministic tie-breaking) has merit (as no larger class ensures convergence even in plurality), things are quite different regarding the voting mechanism itself.

We then show that the iterative process does converge in veto, adding it to plurality as a converging scoring rule. However, we prove that except for these two, all other scoring rules have a setting where they will not converge, regardless of the type of tie-breaking rule used, including, surprisingly, the iterative k-approval voting rules that lie, in a sense, between plurality and veto. By exploring this point we see a limitation of the iterative process—it is mainly relevant, in its current form, to limited options with regard to voting rules. However, these voting rules are among the most popular, so the analysis maintains its applicability. Moreover, the common usage of iterative processes in real life may help explain the prevalence of these voting rules.

In analyzing non-iterative elections, the iterative dynamic helps us consider likely moves by agents (or people), as they try to assess what to do (akin to the effects of polls in real-world elections). Therefore, an important direction of research is to try and find some natural family of strategies that is more limited than just best response, in which other scoring rules may converge. Furthermore, for non-scoring rules, our understanding of what is a reasonable iterative strategy is still limited, and we therefore still seek a better understanding of how an iterative dynamic will behave there; then we can determine whether analysis can yield stable states as an outcome.

Beyond that, we hope to see an undertaking similar to the one taken for plurality by Meir et al. (2014), in which a simulator outputs properties of mechanisms based on various voter distributions. Using such a simulation system to compare different voting rules (as a starting point, plurality and veto) may be an interesting next step in the exploration of iterative models and their representation of common human decision-making processes.

8. Acknowledgements

This work was partially supported by Israel Ministry of Science and Technology grant #3-6797, the Google Inter-University Center for Electronic Markets and Auctions, NSERC grant 482671, Microsoft Research through its PhD Scholarship Program, and by Israel Science Foundation grant #1227/12.

References

Airiau, S., & Endriss, U. (2009). Iterated majority voting. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT)*, pp. 38–49, Venice, Italy.

Bartholdi, III, J. J., & Orlin, J. B. (1991). Single transferable vote resists strategic voting. *Social Choice and Welfare*, *8*(4), 341–354.

Bartholdi, III, J. J., Tovey, C. A., & Trick, M. A. (1989). The computational difficulty of manipulating an election. *Social Choice and Welfare*, *6*(3), 227–241.

Brandt, F., Conitzer, V., Endriss, U., Lang, J., & Procaccia, A. D. (Eds.). (2016). *Handbook of Computational Social Choice*. Cambridge University Press.

Brânzei, S., Caragiannis, I., Morgenstern, J., & Procaccia, A. D. (2013). How bad is selfish voting?. In *Proceedings of the 27th National Conference on Artificial Intelligence (AAAI)*, pp. 138–144, Bellevue, Washington.

Brown, G. W. (1951). Iterative solution of games by fictitious play. *Activity analysis of production and allocation*, *13*(1), 374–376.

Chopra, S., Pacuit, E., & Parikh, R. (2004). Knowledge-theoretic properties of strategic voting. In *Proceedings of 9th European Conference on Logics in Artificial Intelligence (JELIA)*, pp. 18–30, Lisbon, Portugal.

Conitzer, V., Sandholm, T., & Lang, J. (2007). When are elections with few candidates hard to manipulate?. *Journal of the ACM*, *54*(3), 1–33.

de Trenqualye, P. (1998). An extension of Bowen's dynamic voting rule to many dimensions. *Social Choice and Welfare*, *15*(1), 141–159.

Dekel, E., & Piccione, M. (2000). Sequential voting procedures in symmetric binary elections. *Journal of Political Economy*, *108*(1), 34–55.

Dhillon, A., & Lockwood, B. (2004). When are plurality rule voting games dominance-solvable?. *Games and Economic Behavior*, *46*(1), 55–75.

Ephrati, E., & Rosenschein, J. S. (1996). Deriving consensus in multiagent systems. *Artificial Intelligence*, *87*(1-2), 21–74.

Feddersen, T. J., Sened, I., & Wright, S. G. (1990). Rational voting and candidate entry under plurality rule. *American Journal of Political Science*, *34*(4), 1005–1016.

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

Grandi, U., Loreggia, A., Rossi, F., Venable, K. B., & Walsh, T. (2013). Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *Proceedings of 3rd International Conference of Algorithmic Decision Theory (ADT)*, pp. 181–192, Brussels, Belguim.

Hinich, M. J., Ledyard, J. O., & Ordeshook, P. C. (1972). Nonvoting and the existence of equilibrium under majority rule. *Journal of Economic Theory*, *4*(2), 144–153.

Johnson, E. J., Camerer, C., Sen, S., & Rymon, T. (2002). Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory*, *104*(1), 16–47.

Koolyk, A., Lev, O., & Rosenschein, J. S. (2016). Convergence of iterative voting under non-scoring rules (extended abstract). In *Proceedings of the 15th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1329–1330, Singapore.

Laffont, J.-J. (1987). Incentives and the allocation of public goods. In Auerbach, A. J., & Feldstein, M. (Eds.), *Handbook of Public Economics*, Vol. 2 of *Handbook of Public Economics*, chap. 10, pp. 537–569. Elsevier.

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

McKelvey, R. D., & Niemi, R. G. (1978). A multistage game representation of sophisticated voting for binary procedures. *Journal of Economic Theory*, *18*(1), 1–22.

Meir, R. (2015). Plurality voting under uncertainty. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*, pp. 2103–2109, Austin, Texas.

Meir, R., Lev, O., & Rosenschein, J. S. (2014). A local-dominance theory of voting equilibria. In *Proceedings of the 15th ACM conference on Economics and Computation (EC)*, pp. 313–330, Palo Alto, Califronia.

Meir, R., Polukarov, M., Rosenschein, J. S., & Jennings, N. R. (2010). Convergence to equilibria of plurality voting. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pp. 823–828, Atlanta.

Messner, M., & Polborn, M. K. (2007). Strong and coalition-proof political equilibria under plurality and runoff rule. *International Journal of Game Theory*, *35*(2), 287–314.

Myerson, R. B., & Weber, R. J. (1993). A theory of voting equilibria. *The American Political Science Review*, *87*(1), 102–114.

Obraztsova, S., Lev, O., Markakis, V., Rabinovich, Z., & Rosenschein, J. S. (2015a). Beyond plurality: Truth-bias in binary scoring rules. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT)*, pp. 451–468, Lexington, Kentucky.

Obraztsova, S., Lev, O., Polukarov, M., Rabinovich, Z., & Rosenschein, J. S. (2015b). Farsighted voting dynamics. In *AGT@IJCAI*, Buenos Aires, Argentina.

Obraztsova, S., Markakis, E., Polukarov, M., Rabinovich, Z., & Jennings, N. R. (2015c). On the convergence of iterative voting: How restrictive should restricted dynamics be?. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*, pp. 993–999, Austin, Texas.

Procaccia, A. D., & Rosenschein, J. S. (2007). Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, *28*, 157–181.

Procaccia, A. D., Rosenschein, J. S., & Zohar, A. (2007). Multi-winner elections: Complexity of manipulation, control and winner-determination. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1476–1481, Hyderabad, India.

Rabinovich, Z., Obraztsova, S., Lev, O., Markakis, E., & Rosenschein, J. S. (2015). Analysis of equilibria in iterative voting schemes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1007–1013, Austin, Texas.

Reijngoud, A., & Endriss, U. (2012). Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Vol. 2, pp. 635–644, Valencia, Spain.

Reyhani, R., & Wilson, M. C. (2012). Best-reply dynamics for scoring rules. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI)*, Vol. 2, pp. 672–677, Montpellier, France.

Robinson, J. (1951). An iterative method of solving a game. *Annals of Mathematics*, *54*(2), 296–301.

Rothe, J. (Ed.). (2015). *Economics and Computation*. Springer.

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

Shepsle, K. A. (1979). Institutional arrangements and equilibrium in multidimensional voting models. *American Journal of Political Science*, *23*(1), 27–59.

Sloth, B. (1993). The theory of voting and equilibria in noncooperative games. *Games and Economic Behavior*, *5*(1), 152–169.

Thompson, D. R. M., Lev, O., Leyton-Brown, K., & Rosenschein, J. S. (2013). Empirical aspects of plurality election equilibria. In *Proceedings of the 12th International Coference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 391–398, St. Paul, Minnesota.

Xia, L., Zuckerman, M., Procaccia, A. D., Conitzer, V., & Rosenschein, J. S. (2009). Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 348–353, Pasadena, California.

Zuckerman, M., Lev, O., & Rosenschein, J. S. (2011). An algorithm for the coalitional manipulation problem under maximin. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 845–852, Taipei, Taiwan.

Zuckerman, M., Procaccia, A. D., & Rosenschein, J. S. (2008). Algorithms for the coalitional manipulation problem. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 277–286, San Francisco, California.