Trends in Computational Social Choice

4

Cite as: Reshef Meir. Iterative Voting. In Ulle Endriss (editor), *Trends in Computational Social Choice*, chapter 4, pages 69–86. AI Access, 2017.

http://www.illc.uva.nl/COST-IC1205/Book/

CHAPTER 4

Iterative Voting

Reshef Meir

4.1 Introduction

In typical theoretical models of voting, all voters submit their vote at once, without an option to change or revise their decision. While this assumption fits some political voting settings, it fails to hold in most realistic scenarios: committees often follow an informal voting process where members are free to revise their votes or hold noncommittal straw votes; online voting tools such as Facebook and Doodle allow voters to see previous votes and to change their vote by logging in later on; and even in traditional political voting, polls broadcast in the media may prompt voters to change their vote.

Iterative voting games aim to capture such settings. We assume voters have fixed preferences and start from some announcement (e.g., they might sincerely report their preferences). Votes are aggregated via some predefined rule (e.g., Plurality), but voters may change their votes after observing the current announcements and outcome. The game proceeds in turns, where a single voter changes his vote at each turn, until no voter has objections and the final outcome is announced. Crucially, the outcome of iterative voting (like the strategies themselves) may depend on the *order* of voters, which may be affected by external constraints (such as voters' availability to answer an online poll), internal incentives (such as voting early to signal other voters) or other factors.

The lack of a well-defined voting order prevents solutions such as a backward-induction that are common in game theory, even if we are willing to assume that voters are fully rational.

The common assumption in iterative voting is that voters do not know the other voters' preferences or who might change their vote, and thus act in a *myopic way*. That is, vote in every round as if it is the last one, since there is no reliable information for any future prediction. In game-theoretic terms, each voter will play a best reply to the *current action profile* of the other voters. If no voter wants to change his vote, then by definition the current profile is a pure Nash equilibrium (PNE). Voters who are more sophisticated on the one hand, or have less accurate information about the current state on the other hand, may not follow their best reply and instead use other heuristics, in which case equilibria (profiles where no voter wants to move) may not correspond to Nash equilibria.

For any voting rule and type of behavior we are interested in the following questions:

- Are voters guaranteed to converge to an equilibrium? If so, how fast?
- What are all the equilibria reachable from a particular initial state?
- Is the iterative process leading the society to a socially good outcome?

In order to answer these questions, we will first introduce formal game-theoretic definitions for equilibrium and convergence (Section 4.2). Section 4.3 shows how these notions apply in the simple Plurality rule and demonstrates some analysis techniques. Section 4.4 overviews most known results on convergence of iterative voting for myopic rational agents with complete information. In Section 4.5 we relax the model to allow various voting heuristics for iterative voting, focusing on some selected models and convergence results. Section 4.6 concludes and suggests future research directions.

4.2 Preliminaries

For a finite set X, we denote by $\mathcal{L}(X)$ the set of all linear (strict) orders over X. For $L \in \mathcal{L}(X)$, denote by $top(L)$ the first element of L.

A voting instance is defined by a set of *candidates*, or *alternatives*, A, a set of *voters* N, and a *preference profile* $\boldsymbol{L} = (L_1, \dots, L_n)$, where each $L_i \in \mathcal{L}(A)$. For $a, b \in A, i \in N$, candidate a precedes b in L_i (denoted $a \succ_i b$) if voter i prefers candidate a over candidate b. Thus $top(L_i) \in A$ is i's most preferred candidate. In this chapter we assume a voter is never indifferent, i.e., $a \succeq_i b$ means that either $a \succ_i b$ or $a = b$.

A *voting rule* (or, in the game theory literature, *game form*) defines a set of actions A_i for each player, and a function $f : \times_{i \in N} A_i \to A$ from joint actions to alternatives. However most common voting rules assume a certain structure on the action sets. Typically, that the possible actions are preferences over alternatives.

Following Zwicker (2016), a *social choice function* (SCF) is a function that accepts a preference profile $\boldsymbol{L}$ as input, and outputs a nonempty set of winning candidates. Formally: $f : \mathcal{L}(A)^n \to 2^A \setminus \{\emptyset\}$. An SCF f is *resolute* if $|f(\boldsymbol{L})| = 1$ for all $\boldsymbol{L}$. In this chapter we mainly discuss resolute SCFs, which we also refer to as *standard voting rules* or *standard game forms*. However all definitions, except for the notion of truthful vote, naturally extend to non-standard voting rules.[1] See Figure 4.1 for examples of standard and non-standard game forms.

A pair $\langle f, \boldsymbol{L} \rangle$ defines an *ordinal game*, where players' actions are preference orders. The preference of player i over outcomes is given by L_i. In the case of Plurality (or Veto), the set of actions is the set of alternatives A, rather than $\mathcal{L}(A)$. Since A is a coarsening of $\mathcal{L}(A)$, Plurality is still a well-defined SCF.

[1]We emphasize that game forms or voting rules where the actions are not permutations over A (or coarsenings of such permutations) cannot be written as an SCF. An example of a common voting rule that is non-standard is Approval voting.

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

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

f_3	$\emptyset$	a	b	c	ab	ac	bc	abc
a	a	a	a	a	a	a	a	a
b	b	a	b	b	a	a	b	a
c	c	a	b	c	a	a	b	a

f_4	x	y
a	a	b
b	b	c
c	c	a

Figure 4.1: Four examples of game forms with two agents. f_1 is a dictatorial game form with 3 candidates (the row agent is the dictator). f_2 is the Plurality voting rule with 3 candidates and lexicographic tie-breaking. f_3 and f_4 are non-standard game forms. In f_3 player 1 has a single vote, whereas player 2 may approve any subset of candidates. f_4 is a non-standard game form, where the action sets are $A_1 = A = \{a, b, c\}, A_2 = \{x, y\}$.

Convergence and Equilibrium. Any game G induces a directed graph whose vertices are all action profiles (states) $\mathcal{A}$, and edges are all local improvement steps (better replies) (Young, 1993; Andersson et al., 2010). That is, there is an edge from profile $\boldsymbol{a}$ to profile $\boldsymbol{a}'$ if there is some agent i and some action a'_i such that $\boldsymbol{a}' = (\boldsymbol{a}_{-i}, a'_i)$ and i prefers $f(\boldsymbol{a}')$ to $f(\boldsymbol{a})$. A better reply $a_i \stackrel{i}{\to} a'_i$ is called a *best reply* if i has no better reply at profile $(\boldsymbol{a}_{-i}, a'_i)$.

The *sinks* of G are all states with no outgoing edges. Clearly, a state is a sink iff it is a PNE. Since a state may have multiple outgoing edges, we need to specify which one is selected in a given play: in particular, *which player i makes a move* and *which of i's available better replies is selected.*

Much attention has been given in the game theory literature to the question of *convergence*, and several notions of convergence have been defined (Monderer and Shapley, 1996; Milchtaich, 1996; Kukushkin, 2011; Apt and Simon, 2012). For more detailed definitions using schedulers see Meir et al. (2017).

A game G has the *finite individual improvement property* (we say that *G has FIP*), if the corresponding improvement graph has no cycles.

In other words, any sequence of better replies from any initial state $\boldsymbol{a}^0$ reaches a PNE. Games that have FIP are also known as *acyclic games* and as *generalized ordinal potential games* (Monderer and Shapley, 1996). Two weaker notions of acyclicity are as follows.

- A game G has *weakly-FIP* if from any initial state $\boldsymbol{a}^0$ there is *some* path in the improvement graph that reaches a PNE. Such games are known as *weakly acyclic.*
- A game G has *restricted-FIP* (Kukushkin, 2011) if from any initial state $\boldsymbol{a}^0$ and *any order of players* there is *some* path in the improvement graph that reaches a PNE. We refer to such games as *order-free acyclic.*

Intuitively, restricted FIP means that there is some restriction players can adopt such that convergence is guaranteed regardless of the order in which they play. Kukushkin identifies a particular restriction of interest, namely restriction to best reply improvements, and defines the *finite best reply property* (FBRP) and its weak and restricted analogs. The *Finite direct reply property* (FDRP) is only

$$\begin{array}{ccc|c|ccccc}
\text{FBRP f.t.} & \Leftarrow & \text{FBRP} & & \text{restricted-FBRP} & \Rightarrow & \text{weak-FBRP} & & \\
\Uparrow & & \Uparrow & & \Downarrow & & \Downarrow & & \\
\text{FIP f.t.} & \Leftarrow & \text{FIP} & \Rrightarrow & \text{restricted-FIP} & \Rightarrow & \text{weak-FIP} & \Rightarrow & \text{pure Nash} \\
\Downarrow & & \Downarrow & & \Uparrow & & \Uparrow & & \text{exists} \\
\text{FDRP f.t.} & \Leftarrow & \text{FDRP} & & \text{restricted-FDRP} & \Rightarrow & \text{weak-FDRP} & &
\end{array}$$

Figure 4.2: A double arrow X $\Rightarrow$ Y means that any game or game form with the X property also has the Y property. A triple arrow means that any property on the premise side entails all properties on the conclusion side. 'f.t.' means 'from truth.' The third row is only relevant for Plurality/Veto, where direct reply is well defined.

relevant for certain voting rules and is defined later on. Figure 4.2 demonstrates entailment relations among the various acyclicity properties.

We emphasize that the playing agent *must* select an available action, if one exists. For example, we can imagine a dynamics where a voter that only votes for a candidate that is much more preferred than the current winner (say, ranked at least 3 positions above). Such a voter may not move even though he has available better replies. Thus convergence of this dynamic does not imply restricted-FIP.

We say that a game G has *FIP from state* $\boldsymbol{a}$ if all paths from $\boldsymbol{a} \in \mathcal{A}$ reach a PNE. G has *FIP from the truth* if it has FIP from the truthful state $\boldsymbol{a}^* = \boldsymbol{L}$ (for standard rules). We say that a voting rule f has FIP if for *any* preference profile $\boldsymbol{L}$ the induced game $\langle f, \boldsymbol{L} \rangle$ has FIP. The definitions for all other notions of finite improvement properties are analogous.

4.3 Iterative Plurality Voting

Plurality is a particularly simple voting rule, where f returns the candidate ranked at the top position by the largest number of voters.

The *final score* of c for a given profile $\boldsymbol{a} \in A^n$ in the Plurality game form f^{PL} is the total number of voters that vote c. We denote the final score vector by $\boldsymbol{s_a}$ (often just $\boldsymbol{s}$ when the other parameters are clear from the context), where $s(c) = |\{i \in N : a_i = c\}|$. Thus the Plurality rule f^{PL} selects the candidate from $W = \text{argmax}_{c \in A}\, s_{\boldsymbol{a}}(c)$ with the lowest lexicographic index.

Unfortunately, Plurality is not acyclic, and this holds even if voters are restricted to best replies.

Proposition 4.1. *f^{PL} does not have FBRP. In particular it does not have FIP.*

Proof. There are three candidates $A = \{a, b, c\}$ and three voters. We have a single fixed voter voting for a whose preferences are irrelevant. The preference profile of the two other voters is defined as $a \succ_1 b \succ_1 c,\ \ c \succ_2 b \succ_2 a$. The following cycle consists of better replies ((a_1, a_2) are the votes at time t, the winner appears in curly brackets): $(b,c)\{a\} \xrightarrow{2} (b,b)\{b\} \xrightarrow{1} (c,b)\{a\} \xrightarrow{2} (c,c)\{c\} \xrightarrow{1} (b,c)\{a\}$. $\square$

Direct Replies. Meir et al. (2010, 2017) identify a different restriction, namely *direct reply*, which is well defined under the Plurality rule. Formally, a step $\boldsymbol{a} \xrightarrow{i} \boldsymbol{a}'$ is a direct reply if $f(\boldsymbol{a}') = a'_i$, i.e., if i votes for the new winner. An example of an indirect step is when a voter who votes for the winner changes the outcome by moving to a candidate with a low score (e.g., the steps of voter 1 in the example above).

Theorem 4.2 (Meir et al. 2010, 2017). *f^{PL} has FDRP. Moreover, any path of direct replies will converge after at most m^2n^2 steps. In particular, Plurality is order-free acyclic.*

The number of steps until convergence drops to $O(mn)$ if players start from the initial truthful state (Meir et al., 2010) or follow their (unique) direct best reply (Reyhani and Wilson, 2012).

We will demonstrate some of the ideas often used in such proofs, by proving a *weaker* result, namely that any sequence of direct best replies (FDBRP) from the truth converges.

Proof of FDBRP from the truth. Denote by $w^t = f^{PL}(\boldsymbol{a}^t)$ the winner after step t, and by W^t all candidates that can become winners by at most one additional vote. Denote $a = a_i^{t-1}$. We claim that at any step $a \xrightarrow{i} a_i^t$, the following invariants hold:

(1) $a \neq w^{t-1}$ (the manipulator never leaves the current winner);

(2) a_i^t is i's most preferred candidate in $W^{t-1} \setminus \{a\}$;

(3) $a_i^t = w^t$ (vote goes to the new winner);

(4) $a_i^{t-1} \succ_i a_i^t$ (voter always compromises for a less preferred candidate);

(5) $W^t \subseteq W^{t-1}$ (set of possible winners always shrinks); and

(6) For any voter j, either $a_j^t \notin W^t$, or a_j^t is j's most preferred candidate in W^t.

Assume all of (1)-(6) hold until time $t-1$ and consider step t. We prove by induction that all invariants still hold after step t.

Due to (6), we know that either $a \notin W^{t-1}$, or a is i's most preferred in W^{t-1}. Suppose that $a = w^{t-1}$, then we are in the latter case (a is most preferred in W^{t-1}), which means that $w^t = f(\boldsymbol{a}_{-i}^{t-1}, a_i^t) \prec_i a = w^{t-1}$. Thus this cannot be a manipulation step, and $a \neq w^{t-1}$. That is, invariant (1) holds.

Now, since $a \neq w^{t-1}$ the score of the winner after step t does not decrease, only voting candidates in W^{t-1} may change the outcome. Then, invariant (2) follows immediately from our direct best reply assumption.

Invariant (3) follows immediately from the definition of direct replies.

As for (4), either t is the first move of i, in which case $a = top(L_i) \succ_i a_i^t$, or there had been a step $a' \xrightarrow{i} a$ at some time $t' < t$, in which case a is the most preferred in $W^{t'}$. By inductively applying (5), we have that $W^{t-1} \subseteq W^{t'}$, and thus $a \succ_i c$ for all $c \in W^{t-1} \setminus \{a\}$, and in particular $a \succ_i a'$. Thus (4) holds at step t.

We have $s_{\boldsymbol{a}^t}(w^t) \geqslant s_{\boldsymbol{a}^t}(w^{t-1}) = s_{\boldsymbol{a}^{t-1}}(w^{t-1})$ (the equality is due to (1)), which means that the score of the winner weakly increased. Thus the threshold to become a possible winner also weakly increased, whereas the score of all $c \neq a_i^t$ weakly decreased. This means that for any $c \neq a_i^t$ we have $c \notin W^{t-1} \Rightarrow c \notin W^t$, and (5) holds at step t.

Invariant (6) holds at $\boldsymbol{a}^0$ by our assumption of truthful initial vote, and continues to hold as long as (5) does, since a more preferred candidate cannot join W^t. Thus (6) holds at step t. Finally, note that by (3), each voter can move at most $m-1$ times, and thus convergence is achieved in at most $n(m-1)$ steps. □

4.4 Myopic Rational Voters

We overview most of the known results on convergence of iterative voting under various notions of acyclicity, summarized in Table 4.1. Some of the results we cite require some tweaks, for details see Meir et al. (2017). We then study whether equilibria of iterative voting are beneficial for the society.

4.4.1 Strongly Acyclic Voting Rules

It is not hard to see that any *dictatorial rule* (where a single voter determines the outcome) has FIP, due to transitivity of preferences. Interestingly, dictatorships are not the only FIP rules.

In the *direct kingmaker* voting rule (Dutta, 1984) all voters $i \in N \setminus \{1\}$ specify a single candidate $a \in A$, whereas voter 1 selects $i \in N \setminus \{1\}$ to be a "dictator of the day." Note that the direct kingmaker is a non-standard voting rule.

Theorem 4.3. *The direct kingmaker has FIP.*

Proof. Denote $d^t = a_1^t$ as the dictator in $\boldsymbol{a}^t$. In every state $\boldsymbol{a}^t$, only agents 1 and d^t may have a better reply. Further, any better reply of d^t is selecting a more-preferred candidate, i.e., $a_{d^t}^{t+1} \succ_{d^t} a_{d^t}^t$. Thus any agent except agent 1 may move at most $m-1$ times. Since any cycle implies an unlimited number of steps by at least 2 agents, there can be no cycles. □

Characterizing all FIP voting rules is an important and nontrivial problem. For some partial results, see Boros et al. (2010); Kukushkin (2011); Meir et al. (2017).

4.4.2 Order Free and Weak Acyclicity

Veto was shown to converge under direct replies from any initial state (Lev and Rosenschein, 2012; Reyhani and Wilson, 2012). For other common voting rules, results are not as rosy: it is usually possible to construct examples of cycles, even when voters start by voting truthfully. Table 4.1 summarizes known results.

Further, even variations of the Plurality rule, such as adding voters' weights and/or changing the tie-breaking method may result in games with cycles. At least for Plurality with a random tie-breaking rule, it can be shown that it is *weakly acyclic*, thereby providing partial explanation to the fact that simulations almost never hit a cycle (Meir et al., 2017). Whether other common voting rules are also weakly acyclic is an open question, which is particularly of interest for a large number of voters.

Other Notions of Convergence. The above model only considers voters who change their vote one-by-one. Other iterative models exist, that make different assumptions. For example, we can consider voters that make coordinated coalitional moves (Kukushkin, 2011; Gourvès et al., 2016), simultaneous (non-coordinated) moves (Meir, 2015), or a different dynamic where in each step a voter proposes one alternative to replace the current winner using a Majority vote (Airiau and Endriss, 2009).

4.4.3 Reachable Equilibria

Depending on the initial profile and the order of voters, the game may reach one of several equilibria (or none at all, for some voting rules), possibly with different winners. We want to know what these equilibria are. Also, in those cases where an iterative voting game converges, we would like to know "how good" the outcome is to the society. As with other game-theoretic analyses of voting outcomes, there are at least two different approaches to measure outcome quality:

- with respect to the particular voting rule in question,
- with respect to an objective measure, such as social welfare, Condorcet efficiency, etc.

Characterization. The structure of equilibria attained under iterative Plurality voting was studied by Rabinovich et al. (2015), who considered both the model above, and variations where voters are truth-biased (weakly prefer to vote truthfully) or lazy (weakly prefer to abstain). In general, they show that finding whether a particular outcome is reachable via iterative voting from the truthful state, is NP-complete, suggesting that a simple characterization may not exist. In contrast, they provide an efficient algorithm to test which equilibrium states are reachable when the voters a truth-biased or lazy. In fact, when voters are lazy, then in equilibrium at most one voter remains active.

We emphasize that under truth-bias or lazyness, the equilibrium outcomes do not coincide with Nash equilibria.

Dynamic Price of Anarchy. A common way to measure the inefficiency in a game due to strategic behavior is the *Price of Anarchy*: the ratio between the quality of the outcome in the worst Nash equilibrium, and the optimal outcome (Christodoulou and Koutsoupias, 2005). In the context of voting, this translates to the question of how far the equilibrium outcome can be from the truthful voting outcome (seeing the truthful outcome as "optimal" according to the voting rule in use).

As we have seen, Nash equilibria in most voting rules can be arbitrarily far from the truth. Thus, Brânzei et al. (2013) suggested instead to restrict attention to the set of Nash equilibria that are the outcome of some iterative voting procedure, starting from the truthful vote.

Formally, consider a score-based voting rule f, and denote by $s_f(c, \boldsymbol{L})$ the score of candidate c in action profile $\boldsymbol{L}$. The *reachable equilibria* of f, denoted $EQ^T(f, \boldsymbol{L})$, are the set of all profiles $\boldsymbol{L}'$ s.t. :

- There is a path of best-replies from the truthful profile $\boldsymbol{L}$ to $\boldsymbol{L}'$.
- $\boldsymbol{L}'$ is a Nash equilibrium of $\langle f, \boldsymbol{L} \rangle$.

The dynamic Price of Anarchy is defined as

$$\text{DPoA}(f) = \min_{\boldsymbol{L}} \min_{\boldsymbol{L}' \in EQ^T(f,\boldsymbol{L})} \frac{s_f(f(\boldsymbol{L}'), \boldsymbol{L})}{s_f(f(\boldsymbol{L}), \boldsymbol{L})}.$$

Brânzei et al. (2013) show that the DPoA of Plurality is close to 1 (i.e., a winner in equilibrium must have a very close score to the truthful winner); and that the DPoA in Veto depends on the number of candidates m. In particular for $m \leqslant 3$ the DPoA in Veto is constant, regardless of n. The DPoA in Borda, on the other hand, is $\Omega(n)$, meaning that equilibria can be arbitrarily bad.

Objective Quality Metrics. The fact that we chose to use a particular voting rule does not necessarily mean that this rule represents the optimal outcome for every profile. The selection of the rule might be affected by the simplicity of the rule, due to tradition, and so on. We may thus have multiple criteria for a "good outcome," and ask how well a given voting rule satisfies them in equilibrium.

For example, we may be interested in the social welfare of the voters (as measured by Borda score), in the likelihood of finding the Condorcet winner when one exists, or avoiding the Condorcet loser, and so on.

This question was studied using extensive simulations by Meir et al. (2014) for the Plurality rule, where it was shown that equilibria outcomes are *better* than the truthful outcome under most metrics observed. Koolyk et al. (2017) performed similar simulations for several other voting rules, and obtained mixed results w.r.t. the the social welfare. More interestingly, rules that are not Condorcet consistent (Plurality, Bucklin, STV) are more likely to find the Condorcet winner under rational play (i.e., in equilibrium) than under truthful voting, whereas the Condorcet efficiency of Condorcet-consistent rules only slightly declines.

That said, when the number of voters is large, the initial outcome is almost always an equilibrium (whether truthful or not), and thus in most games no voter will move, and the positive effect of best reply dynamics becomes negligible.

4.5 Voting Heuristics

The equilibrium analysis and the best reply dynamics considered in the previous sections makes the following implicit assumptions.

- Voters know exactly how other voters currently vote.
- Voters are myopic: always vote as if the game ends after the current turn.
- Voters are rational: always vote in a way that improves or maximizes their utility.

4.5.1 Ad hoc Heuristics

Most heuristics are similar to best reply in that they only assume the voter knows his own preferences, and has some information about the current voting profile (e.g., the score of each candidate, or their current ranking). However in contrast to best or better reply, a heuristic step may or may not change the outcome. It thus reflects the belief of the voter that she might be pivotal even if this is not apparent from the current state. In the next subsection we will look more closely into such a rational (or bounded-rational) justification, but for now we will be satisfied with just describing some heuristics that have been proposed.

In the following description, we assume f is some score-based rule, unless specified otherwise. Let L_i be the real preferences of voter i, and let $\boldsymbol{a} = (a_1, \ldots, a_n)$,
$\boldsymbol{s} = (s_1, \ldots, s_m)$ be the current action profile and current scores of all candidates. We denote by c_j the candidate with the j'th highest score in $\boldsymbol{s}$.

Crucially, some of these heuristics depend on some internal private parameters, which can be used to explain behavioral differences among voters.

- "k-pragmatist" (Reijngoud and Endriss, 2012): Here, each voter has a parameter k_i, and ranks candidates $W = \{c_1, \ldots, c_{k_i}\}$ at the top according to her real preferences L_i. All other candidates are ranked below W according to their order in L_i.[2]
- "Threshold": Similar to k-pragmatist, except instead of a fixed parameter k_i, the set of "possible winners" W consists of all candidates whose score s_j is above some threshold $T_i(\boldsymbol{a})$ (i.e., the threshold may depend both on i and on the current state).
- "Second Chance" (Grandi et al., 2013): If the current winner is not i's best or second-best choice according to L_i, she moves her second-best alternative to the top position.
- "Best Upgrade" (Grandi et al., 2013): This is a restriction of best reply to candidates that are ranked above the current winner $f(\boldsymbol{a})$ in L_i.
- "Upgrade" (Obraztsova et al., 2015): Similar to Best Upgrade, except the upgraded candidate is not necessarily placed first (only high enough to win).
- "Unit Upgrade" (Obraztsova et al., 2015): Similar to Upgrade, except the upgraded candidate is moved exactly one step up (if this is enough to win).

Simulations show that these heuristics almost always lead to convergence when applied in an iterative voting setting (Grandi et al., 2013).

4.5.2 Strict Uncertainty and Bounded Rationality

A different approach to derive heuristic voting behavior is to consider a formal way to model voter's *uncertainty* regarding the outcome. Then, based on his beliefs, the voter selects the action (ballot) that is best for him.

[2]There is another variation where only the most preferred candidate in W is moved to the top, without other changes.

To see why this may differ from a purely rational behavior, note that:

1. The beliefs of the voter may not be correct or justified.
2. The response of the voter may not maximize his expected utility, which may not even be well-defined.

Thus we can think of such approaches as models of voters with *bounded rationality*. In contrast to Bayesian game-theoretic models of strategic voting (Myerson and Weber, 1993; Messner and Polborn, 2005), voters may not assign exact probabilities to outcomes, and in particular cannot compute expected utilities.[3]

Several papers studied beliefs based on strict uncertainties. One that we already mentioned is (Reijngoud and Endriss, 2012), which assumes the voter may only be aware of the current winner, the order of candidates according to scores, etc. Similar models are considered in other papers without studying iterative moves or convergence (Conitzer et al., 2011; van Ditmarsch et al., 2013). We can also think of voters whose beliefs depend on some internal parameter, which we can think of as their *uncertainty level*. Intuitively, as the voter is more uncertain, she considers more outcomes as possible.

Two such models based on voter's optimism were suggested by Reyhani et al. (2012) and Obraztsova et al. (2016), where in the first paper the optimism is regarding the actual unknown scores, and in the second it is about the voter's ability to prompt other voters into supporting the same candidate. In both models, more optimistic voters will consider a larger set of candidates as possible winners, and will vote strategically to one of them.

Local Dominance. A third model of strict uncertainty is based on local dominance (Meir et al., 2014; Meir, 2015), and can be applied to scoring-based rules. This model explicitly separates the beliefs of the voter on candidates' scores and his strategic actions:

- All voters share some prospective score vector $\boldsymbol{s} = (s_1, \dots, s_m)$.
- Each voter i has an uncertainty parameter r_i.
- Voter i considers as *possible* all outcomes $\boldsymbol{s}'$ such that $|s_c - s'_c| \leqslant r_i$ for all $c \in A$.[4] Denote all possible states in voting profile $\boldsymbol{a}$ by $S_i(\boldsymbol{a}) = \{\boldsymbol{s}' : \|\boldsymbol{s}' - \boldsymbol{s}_{\boldsymbol{a}}\|_{\ell_\infty} \leqslant r_i\}$.
- Given this belief, voter i will change his action from a_i to a'_i if action a'_i *dominates* action a_i. Formally, if $f(\boldsymbol{s}', a'_i) \succeq_i f(\boldsymbol{s}', a_i)$ for all states $\boldsymbol{s}' \in S_i(\boldsymbol{a})$, and $f(\boldsymbol{s}'', a'_i) \succ_i f(\boldsymbol{s}'', a_i)$ for at least one state $\boldsymbol{s}'' \in S_i(\boldsymbol{a})$.

This behavior encodes bounded rationality under *loss aversion*: the voter will make a strategic move only if certain (according to his beliefs) that this move will not hurt him, and might be beneficial.

---

[3]Also note that expected utility is undefined for a voter with ordinal preferences, even if we have such a distribution.

[4]The paper also considers other distance metrics, but the ℓ_∞ metric is the simplest one.

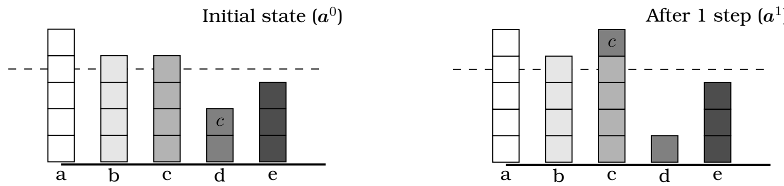

Figure 4.3: A single Local-Dominance step. The letter inside a voter is his *second* preference, thus the highlighted voter has preferences $d \succ_i c \succ_i \{a, b, e\}$. The dashed line marks the threshold of possible winners W_i for voters of type $r_i = 2$. Although $d \xrightarrow{i} c$ is not a better reply (since it does not change the outcome), it is a valid LD move according to Lemma 4.4: $d = a_i^0 \notin W_i$, $c = a_i^1 \in W_i$, and there is a candidate $a \in W_i$ such that $c \succ_i a$.

Denote by $W_i \subseteq A$ the set of candidates whose Plurality score (without voter i) is at least $max_c s(c) - 2r_i$. Note that these are exactly the candidates considered as possible winners by voter i, since there is a possible state $\boldsymbol{s}'$ where $j \in W_i$ gets r_i more votes, and the current winner gets r_i votes less. See Figure 4.3 for an example of a Local-Dominance move.

Local dominance is in fact a special case of π-manipulation (Reijngoud and Endriss, 2012), where in the general case $S_i(\boldsymbol{a})$ may be an arbitrary set of "possible profiles." One other special case of interest is W-manipulation, where the manipulator is assumed to know only the identity of the winner, i.e., $S_i(\boldsymbol{a}) = \{\boldsymbol{a}' \text{ s.t. } f(\boldsymbol{a}) = f(\boldsymbol{a}')\}$.

Worst-Case Regret minimization (WCR) (Meir, 2015) and Non-Myopic voting (NM) (Obraztsova et al., 2016) are similar to local dominance in the way they derive the set of possible winners W_i, but then make some different behavioral assumptions on action selection, which we will not specify explicitly here.

It turns out that the local dominance provides a (bounded) rational justification to the threshold heuristic we described above, at least for Plurality.

Lemma 4.4 (Meir et al. 2014; Meir 2015). *a'_i locally dominates a_i only if:*

1. *either $a_i \notin W_i$, or a_i is the least preferred candidate in W_i;*
2. *$a'_i \in W_i$;*
3. *there is some $c \in W_i$ that is less preferred than a'_i.*

In addition, if the above conditions apply, then the most preferred candidate $a'_i \in W_i$ always locally dominates a_i.

Therefore, a voter that simply follows the threshold heuristics is essentially strategizing according to local dominance.[5] If the voter selects a step *minimizing his worst-case regret* rather than following local dominance moves, then this coincides with the threshold heuristics exactly (Meir, 2015).

Local dominance is strongly related to models based on modal logic and epistemology (Chopra et al., 2004; van Ditmarsch et al., 2013).

[5]Note that while any LD move is consistent with the threshold heuristics, the converse does not always hold. E.g., if there are 5 possible winners above the threshold, then a move from the third-preferred to the most preferred is not a local dominance move.

4.5.3 Equilibrium and Convergence

Given a voting rule f and a population of voters with well defined heuristics, a *voting equilibrium* is simply a profile of valid votes $\boldsymbol{a}$, such that for each $i \in N$, the heuristic action of i in profile $\boldsymbol{a}$ is his current action a_i. That is, a state where no voter wants to change his vote.

Observation 4.5. *Consider an arbitrary voting rule f, and any restricted better reply dynamics (including best reply and better reply). Then for any preference profile $\boldsymbol{L}$, a state $\boldsymbol{a}$ is a voting equilibrium if and only if $\boldsymbol{a}$ is a pure Nash equilibrium of the game $\langle f, \boldsymbol{L} \rangle$.*

Given a voting rule and a heuristic, two important questions are: (a) does an equilibrium exist? and (b) will voters converge to equilibrium? The latter question can be further split to whether convergence is guaranteed from arbitrary initial states or from the truthful state.

Recall that the FIP property means that convergence is guaranteed regardless of the initial state, the order of the voters, and which available reply they choose. As any heuristic simply replaces the (possibly empty) set of better replies with some other set, we can modify the definition of FIP, or FIP from the truth, accordingly.

These questions were studied in several recent papers. Some heuristics are very easy to analyze. For example, when voters start from the truthful vote, then voters using the k-pragmatist or the Second Chance heuristics will move at most once (Reijngoud and Endriss, 2012; Grandi et al., 2013). Therefore, FIP from the truth is immediate. Obraztsova et al. (2015) identified some common structure for heuristic dynamics, which can be used to prove convergence for various combinations of voting rules and heuristics. This was further developed by Endriss et al. (2016) to prove convergence under W-manipulation.

However, all these studies are restricted to voters that start by reporting the truth, and use exactly the same heuristics. Results are summarized in Table 4.2.

Uncertainty-based Heuristics. Uncertainty-based heuristics are more involved, especially when the society is composed of voters with different uncertainty levels. Therefore, they have been studied mostly for the Plurality rule. On the other hand, it turns out that the Local Dominance heuristics has very strong convergence properties. An example of a Plurality game where voters use the Local Dominance heuristics is given in Figure 4.4.

Theorem 4.6 (Meir et al. 2014; Meir 2015). *Plurality with the Local Dominance heuristics has FIP. This holds for any population of voters with either homogeneous or diverse uncertainty levels.*

Limited convergence properties were also shown for other uncertainty-based heuristics. We summarize them in Table 4.2.

4.5.4 Equilibrium Properties

Equilibrium properties are typically studied using simulations, so that uncommon or unlikely equilibria can be ignored. Simulations are carried out by gen-

Voting rule	FIP	FBRP	restricted-FIP	Weak-FIP
Dictator	✓	✓	✓	✓
Direct Kingmaker	✓[M16]	✓	✓	✓
Plurality	×	×[MP+10]	✓[MP+10,MP+17]	✓
Veto	×	×[M16]	✓[RW12,LR12]	✓
k-approval ($k \geqslant 2$)	×	×[LR12,L15]	×	×[M16]
Borda	×	×[RW12,LR12]	×	×[RW12]
PSRs (except k-approval)	×	×[LR12,L15]	?	?
Approval	×	×[M16]	✓[M16]	✓
Other common rules	×	×[KT+17]	?	?

Table 4.1: Positive results carry over to the righthand side, negative to the lefthand side. All rules in the table use lexicographic tie-breaking. Reference codes: MP+10 (Meir et al., 2010), RW12 (Reyhani and Wilson, 2012), LR12 (Lev and Rosenschein, 2012) (see (Lev and Rosenschein, 2016) for the full version), M15 (Meir, 2015), L15 (Lev, 2015), M16 (Meir, 2016), KS+17 (Koolyk et al., 2017), MP+17 (Meir et al., 2017).

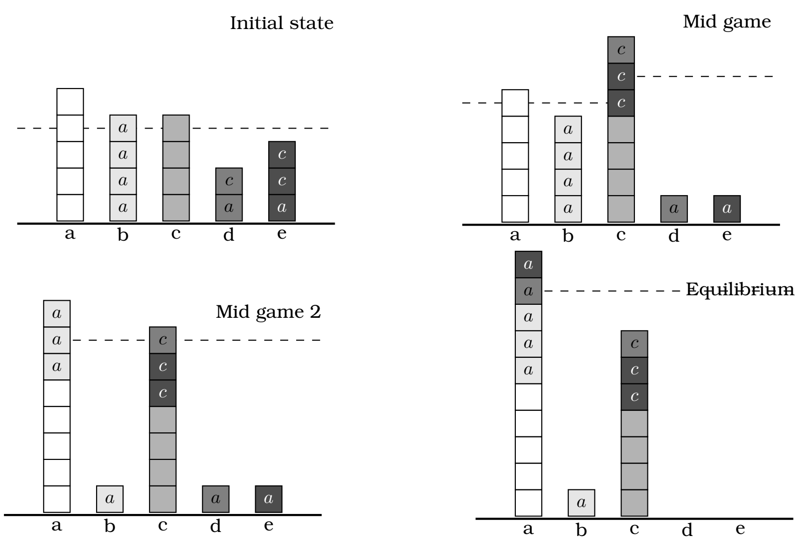

Figure 4.4: Convergence under Local-Dominance. The top left figure shows the initial (truthful) state of the game. The letter inside a voter is his second preference. The dashed line marks the threshold of possible winners W_i for voters of type $r_i = 2$. Note that due to tie breaking it is not the same for all candidates. For example, since a beats b in tie-breaking, b needs 2 more votes to win in the initial state. In the next two figures we can see voters leaving their candidates (who are not possible winners for them) to join one of the leaders. The last figure shows an equilibrium that was reached.

Voting rule	k-prag. [RE12]	Second Chance [GL+13]	Best Upgrade [GL+13]	Upgrade [OM+15]	Unit Upgrade [OM+15]	W-manip. [EO+16]
PSRs	✓	✓	✓	?	✓*	✓#
Maximin	✓	✓	✓	✓	✓	✓
Copland	✓	✓	✓	?	?	✓
Bucklin	-	✓	?	?	✓	?
all rules	-	✓	?	?	?	?

Table 4.2: Positive results mean FIP from the truth for uniform population. * - common PSRs including Borda and Plurality. # - means restricted-FIP. Reference codes: RE12 (Reijngoud and Endriss, 2012), GL+13 (Grandi et al., 2013), OM+15 (Obraztsova et al., 2015), EO+16 (Endriss et al., 2016).

Heuristic	Population	FIP	FIP from truth	equilibrium exists
LD	uniform	✓[M15]	✓[MLR14]	✓
LD	uniform + truth bias	?	✓[MLR14]	✓
LD	diverse	✓[M15]	✓	✓
WCR	uniform	?	✓[M15]	✓
WCR	diverse	×	×	×[M15]
NM	uniform	?	✓[OL+16]	✓
NM	diverse	×	×[OL+16]	?

Table 4.3: Convergence results for Local-Dominance and Worst-Case Regret minimization. All results are for Plurality. Reference codes: MLR14 (Meir et al., 2014), M15 (Meir, 2015), OL+16 (Obraztsova et al., 2016).

erating preference profiles from some distribution (e.g., Impartial culture, Urn, Plackett-Luce, etc.), setting the initial profile (truthful or other), and then sampling voters randomly to make a heuristic move, until an equilibrium is reached. For some heuristics voters' parameters should also be decided up front.

It should first be noted that convergence to equilibrium is achieved in practice (i.e., in simulations) almost always, whether or not this is guaranteed by theorems. Moreover, this convergence is typically very quick.

The *ad hoc* heuristics we mentioned typically lead to a better winner in terms of Condorcet efficiency and Borda score (Grandi et al., 2013). As with the best reply simulations mentioned in Section 4.4.3, this improvement is mild, possibly since the aforementioned heuristics give rise to a small amount of strategic behavior, and thus in many profiles the equilibrium is simply the initial state (note that some of these heuristics are just restrictions of best reply).

Extensive simulations of Plurality voting with Local Dominance heuristics show the following (Meir et al., 2014):

- As uncertainty level r increases, there is more strategic interaction among voters (more moves) until a certain point, from which strategic interaction declines.

- With more strategic interaction, the welfare measures (including Borda score, Condorcet consistency and others) tend to improve, reaching a significant improvement around the peak of strategic activity.
- With more strategic interaction, votes are more concentrated around two candidates (Duverger Law).

These findings are consistent among a broad class of preference distributions, and for different numbers of voters and candidates. Therefore, at least for Plurality, it seems that equilibria reached under Local Dominance resemble outcomes we observe in reality, and avoids unreasonable or highly inefficient Nash equilibria.

4.6 Conclusion

Iterative voting provides a natural tool to analyze strategic voting, which avoids the need to introduce cardinal utilities and probabilities, yet allows for flexible models of bounded rationality and restricted information. It turns out that iterative voting can also be a useful tool for automated systems that aggregate information from many sources (Hassanzadeh et al., 2013).

Theoretical analysis can be used to characterize the conditions under which convergence is expected, as well as properties of the attained equilibria. Future work could focus on general properties of voting rules that lead to positive results, and on designing simple iterative mechanisms that improve efficiency and welfare.

Theory alone, however, cannot settle the question of *which models and assumptions are more plausible*, especially when heuristics are involved. To answer these questions, theoretical models should be combined with empirical data and behavioral experiments (Forsythe et al., 1996; Palfrey, 2009; Kearns et al., 2009; Mattei et al., 2012; Tal et al., 2015; Bassi, 2015). Such interdisciplinary study would help social choice researchers to better understand iterative voting and design better voting mechanisms.

Acknowledgments

The author thanks Nick Jennings, Omer Lev, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S. Rosenschein for many helpful discussions and useful comments.

Bibliography

S. Airiau and U. Endriss. Iterated majority voting. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT)*, 2009.

D. Andersson, V. Gurvich, and T. D. Hansen. On acyclicity of games with cycles. *Discrete Applied Mathematics*, 158(10):1049–1063, 2010.

K. R. Apt and S. Simon. A classification of weakly acyclic games. In *Proceedings of the 5th International Symposium on Algorithmic Game Theory (SAGT)*, 2012.

A. Bassi. Voting systems and strategic manipulation: An experimental study. *Journal of Theoretical Politics*, 27(1):58–85, 2015.

E. Boros, V. Gurvich, K. Makino, and D. Papp. Acyclic, or totally tight, two-person game forms: Characterization and main properties. *Discrete Mathematics*, 310 (6):1135–1151, 2010.

S. Brânzei, I. Caragiannis, J. Morgenstern, and A. D. Procaccia. How bad is selfish voting? In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 2013.

S. Chopra, E. Pacuit, and R. Parikh. Knowledge-theoretic properties of strategic voting. In *Proceedings of the 9th European Conference On Logics In Artificial Intelligence (JELIA)*, 2004.

G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on the Theory of Computing (STOC)*, 2005.

V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*, 2011.

B. Dutta. Effectivity functions and acceptable game forms. *Econometrica: Journal of the Econometric Society*, 52(5):1151–1166, 1984.

U. Endriss, S. Obraztsova, M. Polukarov, and J. S. Rosenschein. Strategic voting with incomplete information. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

R. Forsythe, T. Rietz, R. Myerson, and R. Weber. An experimental study of voting rules and polls in three candidate elections. *International Journal of Game Theory*, 25(3):355–383, 1996.

L. Gourvès, J. Lesca, and A. Wilczynski. Strategic voting in a social context: Considerate equilibria. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*, 2016.

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

F. F. Hassanzadeh, E. Yaakobi, B. Touri, O. Milenkovic, and J. Bruck. Building consensus via iterative voting. In *Proceedings of the IEEE International Symposium on Information Theory Proceedings (ISIT)*, 2013.

M. Kearns, S. Judd, J. Tan, and J. Wortman. Behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences*, 106(5): 1347–1352, 2009.

A. Koolyk, T. Strangway, O. Lev, and J. S. Rosenschein. Convergence and quality of iterative voting under non-scoring rules. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.

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

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

O. Lev and J. S. Rosenschein. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.

O. Lev and J. S. Rosenschein. Convergence of iterative scoring rules. *Journal of Artificial Intelligence Research*, 57:573–591, 2016.

N. Mattei, J. Forshee, and J. Goldsmith. An empirical study of voting rules and manipulation with large datasets. *Proceedings of the 4th International Workshop on Computational Social Choice (COMSOC)*, 2012.

R. Meir. Plurality voting under uncertainty. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

R. Meir. Strong and weak acyclicity in iterative voting. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory (SAGT)*, 2016.

R. Meir, M. Polukarov, J. S. Rosenschein, and N. Jennings. Convergence to equilibria of plurality voting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 2010.

R. Meir, O. Lev, and J. S. Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of the 15th ACM Conference on Electronic Commerce (ACM-EC)*, 2014.

R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Acyclic games and iterative voting. *Artificial Intelligence*, 2017. Forthcoming.

M. Messner and M. Polborn. Robust political equilibria under plurality and runoff rule. *IGIER Working Paper*, 2005.

I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.

D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

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

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

S. Obraztsova, O. Lev, M. Polukarov, Z. Rabinovich, and J. S. Rosenschein. Non-myopic voting dynamics: An optimistic approach. In *Proceedings of the 10th Multidisciplinary Workshop on Advances in Preference Handling (M-PREF)*, 2016.

T. Palfrey. Laboratory experiments in political economy. *Annual Review of Political Science*, 12:379–388, 2009.

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

A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.

R. Reyhani and M. C. Wilson. Best-reply dynamics for scoring rules. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI)*, 2012.

R. Reyhani, M. C. Wilson, and J. Khazaei. Coordination via polling in plurality voting games under inertia. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI)*, 2012.

M. Tal, R. Meir, and Y. Gal. A study of human behavior in voting systems. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2015.

H. van Ditmarsch, J. Lang, and A. Saffidine. Strategic voting and the logic of knowledge. In *Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2013.

H. P. Young. The evolution of conventions. *Econometrica: Journal of the Econometric Society*, 61(1):57–84, 1993.

W. S. Zwicker. Introduction to the theory of voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.