Convergence and Quality of Iterative Voting Under Non-Scoring Rules

Aaron Koolyk
Hebrew University of Jerusalem
Jerusalem, Israel
aaron.koolyk@mail.huji.ac.il

Omer Lev
University of Toronto
Toronto, Canada
omerl@cs.toronto.edu

Jeffrey S. Rosenschein
Hebrew University of Jerusalem
Jerusalem, Israel
jeff@cs.huji.ac.il

ABSTRACT

Iterative voting is a social choice mechanism whereby voters are allowed to strategically change their stated preferences as the vote progresses, until an equilibrium is reached; at this point, no voter can make a beneficial strategic change. We study iterative voting for several common voting rules and show that, for these rules, such an equilibrium may never be reached. We also consider several variations of iterative voting (restrictions on the allowable changes to votes), and show that with these variations equilibrium may also not be reached. Finally, we do an empirical analysis of the quality of candidates elected through iterative voting.

1. INTRODUCTION

The topic of voting, that is, how to aggregate diverse individual preferences into a collective decision, is of great importance in many automated agent scenarios; it has thus been the topic of much research in multiagent systems. One innovative voting model that was recently proposed is that of *iterative voting* [12]. Whereas classic voting rules usually consist of a single round of ballot submission and announcement of the winner, in iterative voting there can be many such rounds. After each iteration, if any voter wishes to change their vote they may do so, and potentially a new winner replaces the previous one (when multiple such voters exist, an arbitrary voter is chosen). The process terminates when no voter wishes to change their vote. Iterative voting thus embraces the inevitable manipulability of voting shown in the Gibbard-Satterthwaite theorem [20, 6], and considers agents' uniform ability to vote strategically as a collective opportunity.

Besides being an intriguing method for reaching consensus, iterative voting has been proposed as a formal solution concept for voting games. Standard Nash equilibria are of limited usefulness in voting games where the group outcome is generally robust to changes in any single voter's action. The set of iterative voting equilibria, however, is a subset of Nash equilibria, and in particular those iterative voting equilibria reachable from the truthful profile could be considered a more natural (or meaningful) solution concept.

The most salient questions regarding iterative voting thus have two interpretations. Regarding iterative voting as a method for reaching an outcome, we ask whether the process terminates; if so, with what complexity; and does it arrive at "good" outcomes. Regarding iterative voting as a solution concept, we must explore the existence of solutions; the equilibria computation; and notions of price of stability/anarchy.

Most previous research on iterative voting has focused on plurality, with several extensions to other scoring rules. However, in this work we look into previously unexplored voting rules which are not scoring rules—Maximin, Copeland, Bucklin, STV, Second Order Copeland (SOC), and Ranked Pairs. In the process of investigating these voting rules, we design dynamics under which the iterative voting might evolve. While we show that convergence is not guaranteed, we proceed to analyze the outcomes of the iterative dynamic empirically, showing that cycles are not very common, and moreover, the outcomes are generally very good.

2. RELATED LITERATURE

There has been extensive research on solution concepts of voting games, and an overview of the research can be seen in Meir et al. [11]. Due to space constraints, we focus in this section on the iterative model which we extend.

Our model of iterative voting was initiated by Meir et al. [12], who showed that Plurality voting converges under a natural restricted best-response dynamic and linear ordered tie-breaking (a dynamic refined in [10]). Lev and Rosenschein [8] (and in parallel [19]) later showed that Veto, with a similarly natural restricted best-response dynamic, also converges. However, [8] and further work showed that no other scoring rules converge for best-response dynamics, as well as showing that allowing non-linear tie-breaking will result in Maximin not converging (though they showed no result on Maximin with linear tie-breaking).

Reijngoud and Endriss [18] added an epistemic element by varying the amount of information revealed at each stage, and also showed that any scoring rule converges under the k-pragmatism dynamic. Grandi et al. [7] showed that, for two additional restrictive dynamics, scoring rules (as well as Copeland and Maximin), converge, and Loreggia [9] added another very restrictive dynamic, showing that Copeland and Maximin converge under it. Obraztsova et al. [16] abstracted these ideas and put forth two theoretical properties which suffice to guarantee convergence. Not in connection to iterative voting, Obraztsova and Elkind [13] proposed several dynamics, of which we adopt, for example, the Kendall-Tau dynamic.

Research examining the properties of iterative voting oc-

Appears at: 3rd Workshop on Exploring Beyond the Worst Case in Computational Social Choice. Held as part of the 15th International Conference on Autonomous Agents and Multiagent Systems. May 9th-10th, 2016. Singapore.

curred in parallel. Meir et al. [11] showed that a generalization of iterative voting, where voters act under uncertainty, also converges for Plurality. [17, 14] considered questions of computational complexity related to iterative voting, with and without assumptions about "truth bias" and "lazy bias" on the part of voters.

Brânzei et al. [3] examined the *quality* of iterative voting, via the notion of the dynamic price of anarchy. They showed lower bounds for Plurality, Veto, and Borda, and a tight upper bound for Plurality. Additional work on the quality of iterative voting includes that of [11, 18, 7] (all mentioned in other contexts, above) who showed through simulations some improvements in the outcome of elections, in their various versions of iterative voting. However, the closest work in its pattern of simulations and quality measures is Thompson et al. [21], which analyzed truth-biased equilibria, without any assumption regarding their dynamics.

3. PRELIMINARIES

Our setting will be the standard voting model that includes a set of voters $V, |V| = n$, and a set of candidates $C, |C| = m$. Each voter i has a strict preference order $\succ_i$ over C, that is, a complete, reflexive, transitive, and anti-symmetric binary relation over C. Denote the set of all such preference orders as $\mathcal{P}(C)$. A profile

$$\vec{\succ} = (\succ_1, \succ_2, ..., \succ_n) \in \mathcal{P}(C)^n$$

is a vector of n preference orders, one for each voter. We denote by

$$\vec{\succ}_{-i} = (\succ_1, ..., \succ_{i-1}, \succ_{i+1}, \dots, \succ_n) \in \mathcal{P}(C)^{n-1}$$

the profile of the voters excluding i and $(\vec{\succ}_{-i}, \succ_i) = \vec{\succ}$. We shall denote the truthful preferences of voters as $\vec{tr} = (\succ_1^{tr}, \dots, \succ_n^{tr})$.

We model a collective decision through one of two functions. A *social welfare function* is a function $f : \mathcal{P}(C)^n \to \mathcal{P}(C) \backslash \{\emptyset\}$ and a *voting rule* is a function $F : \mathcal{P}(C)^n \to 2^{\mathcal{C}} \backslash \{\emptyset\}$. So, given a (not necessarily truthful) vector of preferences, a social welfare function chooses a preference order and a voting rule chooses a set of candidates. When a voting rule is irresolute, and we would like a unique winner, we use a *tie breaking rule*, a function $t : 2^{\mathcal{C}} \to \mathcal{C}$. A *linear-ordered tie breaking rule* is a rule that breaks ties according to a fixed linear order. It will be without loss of generality throughout this paper that the linear-ordered tie breaking rule we use is the *lexicographic tie breaking rule*, where ties are broken according to the lexicographic order of candidates' names.

3.1 Voting Rules

We shall investigate the following voting rules:

Maximin For each pair of candidates c_1, c_2 let $P(c_1, c_2) = |\{x \in V | c_1 \succ_x c_2\}|$. For each candidate c, let $sc(c) = \min_{c' \neq c \in C} P(c, c')$. The candidates with the maximum score, $\arg\max_{c \in C} sc(c)$, win.

Copeland Let $P(c_1, c_2) = |\{x \in V | c_1 \succ_x c_2\}|$, as above. For $\alpha \in [-1, 1]$, let $sc(c) = |\{c' | P(c, c') > n/2\}| - |\{c' | P(c, c') < n/2\}| + \alpha \cdot |\{c' | P(c, c') = n/2\}|$, and the candidates with the maximum score, $\arg\max_{c \in C} sc(c)$, win. (Generally $\alpha = 0$ is assumed).

Bucklin For each $c \in C$, let $sc(c) = \min_{k<m} |\{x \in V | \exists c_1, \dots c_{m-k} \text{ s.t. } c \succ_x c_i\}| > n/2$. The winner is the candidate with the smallest score, $\arg\min_{c \in C} sc(c)$.

STV Under Single Transferable Voting (STV), the election proceeds in rounds. In each round, the candidate with the lowest plurality score is eliminated and any voter voting for them transfers their vote to their next ranked candidate. The last remaining candidate is the winner.

SOC Second Order Copeland (SOC) chooses winners as in Copeland, except that ties are broken according to the score of defeated candidates. If $sc(c)$ is the Copeland score of c, then Second Order Copeland chooses $c \in \arg\max_{c \in C} sc(c)$ s.t. $\sum_{c' : P(c, c') > n/2} sc(c')$ is maximal.

Ranked Pairs (RP) Let $P(c_1, c_2) = |\{x \in V | c_1 \succ_x c_2\}|$, as above. Let

$$O = (P(c_{i_{1,1}}, c_{i_{1,2}}), P(c_{i_{2,1}}, c_{i_{2,2}}), ..., P(c_{i_{\binom{m}{2},1}}, c_{i_{\binom{m}{2},2}}))$$

be the sorted list of pairs of candidates' P-score such that

$$P(c_{i_{j,1}}, c_{i_{j,2}}) \geq P(c_{i_{j+1,1}}, c_{i_{j+1,2}}).$$

If $P(c_{i_{j,1}}, c_{i_{j,2}}) = P(c_{i_{j+1,1}}, c_{i_{j+1,2}})$, then

$$P(c_{i_{j,1}}, c_{i_{j,2}}) \succ_O P(c_{i_{j+1,1}}, c_{i_{j+1,2}})$$

iff $i_{j,1} < i_{j+1,1}$ or $i_{j,1} = i_{j+1,1}$ and $i_{j,2} < i_{j+1,2}$; i.e., break ties in order lexicographically (first candidate, second candidate). A ranking is constructed by the following algorithm. For j=0 to $\binom{m}{2}$ fix $c_{i_{j,1}} \succ c_{i_{j,2}}$ unless this contradicts a previous step (including by transitivity). The candidate at the top of the constructed ranking is selected as the winner.

An interesting property of which we will make use regards the *Condorcet winner*. A Condorcet winner is a candidate who is preferred to each other candidate by more than half of the voters; however, such a winner does not always exist. A voting rule is *Condorcet consistent* if whenever there is such a Condorcet winner, it is the election's outcome. Among the voting rules we discuss, Maximin, Copeland, SOC and Ranked Pairs are Condorcet consistent, while Bucklin and STV are not.

3.2 Dynamics

We will call a binary relation $\mathcal{D} \subset \mathcal{P}(C)^n \times \mathcal{P}(C)^n$ a *dynamic*. We call a (possibly finite) sequence of profiles $(\vec{\succ}_1, \vec{\succ}_2, ...) \in \mathcal{P}(C)^*$ a *profile sequence* and a (possibly finite) sequence of voters $(v_1, v_2, ...) \in V^*$ a *voter sequence*. A profile sequence $(\vec{\succ}_1, \vec{\succ}_2, ...)$ for which $\vec{\succ}_1$ are the truthful preferences, is called an *initially truthful* profile sequence.

We will say a profile sequence is *valid* for a dynamic $\mathcal{D}$ if $\forall i (\vec{\succ}_i, \vec{\succ}_{i+1}) \in \mathcal{D}$. We will mainly be concerned with dynamics for which all elements differ in a single preference, i.e.,

$$\forall [(\vec{\succ}^{(1)}, \vec{\succ}^{(2)}) \in \mathcal{D}] \;\; \exists i \in V \;\; \text{s.t. } \vec{\succ}^{(1)}_{-i} = \vec{\succ}^{(2)}_{-i}.$$

In such a case, a profile sequence induces a voter sequence $(v_1, v_2, ...)$ where v_i is the voter whose preference changed at stage i. Likewise, a voter sequence defines a set of profile sequences by which it is induced. A voter sequence will be called *valid* if it is induced by some valid profile sequence.

The final element of a (finite) valid profile sequence $(\vec{\succ}_1, \vec{\succ}_2, ..., \vec{\succ}_k)$ will be called an *equilibrium* if there is no $\vec{\succ}_{k'}$ such that $(\vec{\succ}_k, \vec{\succ}_{k'}) \in \mathcal{D}$.

For a dynamic $\mathcal{D}$ and voting rule F with tie breaking rule t, let $\mathcal{I}(\mathcal{D}, F_t) = \{s | s \text{ is a valid profile sequence for } \mathcal{D}(F_t)\}$. We will say that *iterative-F converges* (or *always converges*) under $\mathcal{D}$ if every element of $\mathcal{I}(\mathcal{D}, F_t)$ is finite. Otherwise, we will say that iterative-F under $\mathcal{D}$ *cycles* (or *may cycle*) or does not converge (or may not converge). Notice that, as defined, the semantics of convergence are asynchronous. $\mathcal{I}(\mathcal{D}, F_t)$ converges if *every* element is finite, and is not limited to, say, a "fair schedule of play."

The dynamics we shall consider will be influenced by the truthful preferences, i.e., a dynamic in which a voter's vote changed must have increased the utility of that vote. Two main dynamics have been investigated (e.g., in Meir et al. [12]). An ordered pair of profiles is in the *better response* dynamic if the preferences of all voters but one are identical in the two profiles, and the voter whose preference changes prefers the outcome of the second profile to that of the first profile. In game-theoretic terms, any time a single player can make a better response to a given state, such a move is included in the dynamic. Formally, for two profiles $\vec{\succ}^{(1)}, \vec{\succ}^{(2)}$ and a voting rule F, $(\vec{\succ}^{(1)}, \vec{\succ}^{(2)}) \in$ BetterResponse iff:

$$\exists i \in V \text{ s.t. } \vec{\succ}^{(1)}_{-i} = \vec{\succ}^{(2)}_{-i} \text{ and } F(\vec{\succ}^{(2)}) \succ^{tr}_i F(\vec{\succ}^{(1)}).$$

Such an i is called the *manipulator*, $\succ^{(2)}_i$ is called the *new vote*, and $\succ^{(1)}_i$ is called the *old vote*. Notice that a stable state under this dynamic is a Nash equilibrium.

Similarly, an ordered pair of profiles is in the *best response* (BR) dynamic if the preferences of all voters but one are identical; the voter whose preference changes prefers the outcome of the second profile to that of the first profile (so it is contained in the better response dynamic); and of all possible changes to his preferences, the outcome under the second profile is preferred at least as much as the outcome under any other possible profile. Formally, $(\vec{\succ}^{(1)}, \vec{\succ}^{(2)}) \in$ BR iff:

$$\exists i \in V \text{ s.t. } \vec{\succ}^{(1)}_{-i} = \vec{\succ}^{(2)}_{-i} \text{ and } F(\vec{\succ}^{(2)}) \succ^{tr}_i F(\vec{\succ}^{(1)})$$

and

$$\forall \succ'' \in \mathcal{P}(C) \text{ s.t. } (\vec{\succ}^{(1)}_{-i}, \succ'') \neq \vec{\succ}^{(2)}, F(\vec{\succ}^{(2)}) \succeq^{tr}_i F(\vec{\succ}^{(1)}_{-i}, \succ'').$$

The above description clearly defines a game form. The set of voters is the set of players, the set of preferences is the set of strategies available to each player, and the voting rule determines the outcome of a strategy profile. Ordinal utilities are given by true preference orders. An equilibrium under Best Response (or Better Response) is a Nash equilibrium.

4. DYNAMICS

The study of best response dynamics is prolific, but in the iterative voting context, particular forms of best response have been utilized in the convergence proofs of both plurality [12] and veto [8]. For non-scoring rules, however, there is no immediately clear choice of best response form (indeed, in some cases, like STV, it is NP-complete to calculate what it is). We present here several dynamics that may serve as natural heuristics for a potential voter. There have been dynamics designed with the express purpose of ensuring convergence, as in k-pragmatism, M1, and M2 [18, 7]. However, we propose the following as possibly more natural correspondences to the strategic behavior of self-interested agents.

TOP: This dynamic assigns the candidate which the voter wishes to make a winner the top spot in the new preference order. In many of the voting rules we consider (and any weakly-monotone rule) this dynamic is a subset of the best-response dynamic (i.e., $TOP(\mathcal{P}(C)) \subset BR(\mathcal{P}(C))$), and, indeed, it generalizes the dynamic used in Meir et al. [12].

TB: This dynamic requires the new winner to be at the top of the new ballot, and the previous winner to be at the bottom. While in many scoring rules (e.g., plurality and veto) this is a subset of best response moves (and generalizes those used in Lev and Rosenschein [8]), this is not true in general, and particularly in the voting rules we study in this work.

KT: This dynamic restricts best response to those with minimum Kendall-Tau distance from the previous vote. That is, among all possible moves whose outcome will be the most preferred possible candidate, one with the minimal Kendall-Tau distance[1] from the current vote is chosen.

SWAP: This dynamic, inspired in part by notions from the literature on bribery (see, e.g., [5, 4]), is quite restrictive. It restricts manipulations to a single swap (called a 'shift' in the bribery literature) or even a single adjacent swap (i.e., changing to a vote within Kendall-Tau distance of one from the current vote; a 'swap' in the bribery nomenclature).

5. CONVERGENCE

In this section we consider the convergence of iterative voting for several voting rules. We distinguish between the first three, for which there exists a polynomial time algorithm for a single voter to compute a best response manipulation, and the last three for which such a computation is NP-Complete [2, 1, 22]. In reversal of the common situation in computational social choice, for iterative voting polynomial manipulation is actually quite felicitous.

A note on reading the examples: each column represents a profile of submitted ballots (beginning with the truthful one). The final row in the column indicates the winner of the profile. The i-th entry in a column represents voter i's submitted preferences, where, for example, ABC is to be read $A \succ_i B \succ_i C$. Arrows highlight the changed preference between two profiles at a given stage. The profile sequence formed by continual repetition of the indicated profiles thus forms an infinite element of $\mathcal{I}(\mathcal{D}, F_t)$ and proves non-convergence. Due to space constraints, we omit several proofs and examples.

5.1 Maximin

Similar to plurality and veto, Maximin changes gradually. The difference in score between the previous winner and the new one, when a single voter manipulates, can go up or down by at most one point. One might thus expect there to be an argument for convergence, similar to plurality/veto. But in fact, convergence with Maximin turns out to be elusive even

[1]For $a, b \in \mathcal{P}(C)$, the Kendall-Tau distance between them is defined as $dist(a, b) = |\{i, j\} \in V | (i \succ_a j \text{ and } j \succ_b i) \text{ or } (j \succ_a i \text{ and } i \succ_b j)\}|$.

after major restrictions on the allowable moves.

Theorem 1. *Maximin with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT and SWAP.*

Proof. For BR, the example is:

$$\begin{array}{ccccccc} BAC & & BAC & \longrightarrow & CBA & & CBA \\ CAB & \longrightarrow & ABC & & ABC & \longrightarrow & CAB \\ CAB & & CAB & & CAB & & CAB \\ BCA & & BCA & & BCA & & BCA \\ \mathbf{B} & & \mathbf{A} & & \mathbf{B} & & \mathbf{C} \end{array}$$

For TOP, the example is:

$$\begin{array}{ccccccc} CDAB & & CDAB & \longrightarrow & DABC & & DABC \\ ACBD & \longrightarrow & ADBC & & ADBC & \longrightarrow & ACBD \\ BCDA & & BCDA & & BCDA & & BCDA \\ \mathbf{C} & & \mathbf{A} & & \mathbf{D} & & \mathbf{A} \end{array}$$

For TB, the example is:

$$\begin{array}{ccccccccc} ACDBE & \rightarrow & ACEBD & & ACEBD & \rightarrow & ADEBC & & ADEBC \\ DBCEA & & DBCEA & \rightarrow & CBDEA & & CBDEA & \rightarrow & DBCEA \\ EBDCA & & EBDCA & & EBDCA & & EBDCA & & EBDCA \\ \mathbf{D} & & \mathbf{A} & & \mathbf{C} & & \mathbf{A} & & \mathbf{D} \end{array}$$

For KT, the example is:

$$\begin{array}{ccccccc} DBACE & \longrightarrow & BDCAE & & BDCAE & & BDCAE \\ EDCAB & & EDCAB & \rightarrow & DECAB & & DECAB \\ CEADB & & CEADB & & CEADB & \rightarrow & EACBD \\ \mathbf{A} & & \mathbf{B} & & \mathbf{D} & & \mathbf{A} \end{array}$$

$$\begin{array}{ccccc} DBCAE & & DBCAE & & DBCAE \\ EDCAB & & EDCAB & \leftarrow & DECAB \\ CEADB & \leftarrow & EACBD & & EACBD \\ \mathbf{C} & & \mathbf{E} & & \mathbf{D} \end{array}$$

And, finally, for SWAP, the example is:

$$\begin{array}{ccccccccccc} CABD & & CABD & & CABD & & CABD & \rightarrow & CADB & & CADB \\ BDCA & & BDCA & & BDCA & & BDCA & & BDCA & & BDCA \\ DBAC & & DBAC & & DBAC & & DBAC & & DBAC & & DBAC \\ DCBA & & DCBA & \rightarrow & CDBA & & CDBA & & CDBA & \rightarrow & DCBA \\ ABCD & \rightarrow & ACBD & & ACBD & \rightarrow & ABCD & & ABCD & & ABCD \\ \mathbf{B} & & \mathbf{A} & & \mathbf{C} & & \mathbf{B} & & \mathbf{A} & & \mathbf{D} \end{array}$$

Although the changes to the winner's score are as gradual in Maximin as in plurality and veto, the exponential blowup in strategy space seems to make convergence harder. Whereas in plurality and veto, a voter's ballot reduces to a single candidate, in Maximin a ballot depends on the entire ranking. □

5.2 Copeland

Theorem 2. *Copeland with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT and SWAP. This holds for Copeland$^\alpha$ for any α.*

Proof. Since the number of voters in all our examples is odd, they hold for Copeland$^\alpha$ for any α.

The example for BR:

$$\begin{array}{ccccccc} BDCA & & BCDA & \longrightarrow & DBCA & & DBCA \\ CDAB & \longrightarrow & DABC & & DABC & \longrightarrow & CDAB \\ ABCD & & ABCD & & ABCD & & ABCD \\ \mathbf{B} & & \mathbf{A} & & \mathbf{D} & & \mathbf{C} \end{array}$$

The example for the TOP dynamic:

$$\begin{array}{ccccccc} DABC & & DABC & \longrightarrow & ACBD & & ACBD \\ BDAC & \longrightarrow & BACD & & BACD & \longrightarrow & BDAC \\ CDBA & & CDBA & & CDBA & & CDBA \\ \mathbf{D} & & \mathbf{B} & & \mathbf{A} & & \mathbf{B} \end{array}$$

Using the TB dynamic and moving the desired winner to the top and the current undesired winner to the bottom does not suffice to avoid cycles:

$$\begin{array}{ccccccc} ACB & & ACB & \longrightarrow & ABC & & ABC \\ BCA & \longrightarrow & CBA & & CBA & \longrightarrow & BCA \\ \mathbf{A} & & \mathbf{C} & & \mathbf{A} & & \mathbf{B} \end{array}$$

The exact same example as TB also serves to show restricting best response by minimum Kendall-Tau distance does not suffice to avoid cycles.

Finally, restrictions to a single adjacent swap does not suffice:

$$\begin{array}{ccccccc} BCA & & BCA & \longrightarrow & CBA & & CBA \\ ABC & \longrightarrow & ACB & & ACB & \longrightarrow & ABC \\ \mathbf{B} & & \mathbf{A} & & \mathbf{C} & & \mathbf{A} \end{array}$$

□

5.3 Bucklin

Theorem 3. *Bucklin with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT and SWAP.*

Proof. The example for BR:

$$\begin{array}{ccccccc} ACB & & ACB & \longrightarrow & CBA & & CBA \\ BAC & \longrightarrow & CBA & & CBA & \longrightarrow & BAC \\ ABC & & ABC & & ABC & & ABC \\ CBA & & CBA & & CBA & & CBA \\ A & & \mathbf{B} & & \mathbf{C} & & B \end{array}$$

The example for TOP:

$$\begin{array}{ccccccccccc} BACD & & BACD & \rightarrow & CBAD & & CBAD & \rightarrow & ADBC & & ADBC \\ DBAC & \rightarrow & DCAB & & DCAB & \rightarrow & DABC & & DABC & \rightarrow & DBAC \\ BDAC & & BDAC & & BDAC & & BDAC & & BDAC & & BDAC \\ CDBA & & CDBA & & CDBA & & CDBA & & CDBA & & CDBA \\ ADCB & & ADCB & & ADCB & & ADCB & & ADCB & & ADCB \\ \mathbf{B} & & \mathbf{D} & & \mathbf{C} & & \mathbf{D} & & \mathbf{A} & & \mathbf{D} \end{array}$$

The example for the TB dynamic:

$$\begin{array}{ccccccc}
ADBC & & ADBC & \longrightarrow & ACBD & & ACBD \\
CBDA & \longrightarrow & DBCA & & DBCA & \longrightarrow & CBDA \\
CADB & & CADB & & CADB & & CADB \\
DBAC & & DBAC & & DBAC & & DBAC \\
\mathbf{A} & & \mathbf{D} & & \mathbf{A} & & \mathbf{C}
\end{array}$$

The example for the KT dynamic:

$$\begin{array}{ccccccc}
ABCD & & ABCD & \longrightarrow & ACBD & & ACBD \\
DCBA & \longrightarrow & DBCA & & DBCA & \longrightarrow & DCBA \\
CBAD & & CBAD & & CBAD & & CBAD \\
DACB & & DACB & & DACB & & DACB \\
\mathbf{A} & & \mathbf{B} & & \mathbf{A} & & \mathbf{C}
\end{array}$$

The example for the SWAP dynamic:

$$\begin{array}{ccccccc}
DCBA & & DCBA & \longrightarrow & DBCA & & DBCA \\
ABCD & \longrightarrow & ACBD & & ACBD & \longrightarrow & ABCD \\
CDAB & & CDAB & & CDAB & & CDAB \\
BDAC & & BDAC & & BDAC & & BDAC \\
\mathbf{D} & & \mathbf{C} & & \mathbf{D} & & \mathbf{B}
\end{array}$$ □

5.4 STV

Theorem 4. *STV with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT and SWAP.*

Proof. The example for the BR dynamic:

$$\begin{array}{ccccccccc}
BCDA & \rightarrow & DCAB & & DCAB & \rightarrow & BADC & & BADC \\
ADBC & & ADBC & \rightarrow & DBCA & & DBCA & \rightarrow & ADBC \\
CDAB & & CDAB & & CDAB & & CDAB & & CDAB \\
\mathbf{A} & & \mathbf{C} & & \mathbf{D} & & \mathbf{B} & & \mathbf{A}
\end{array}$$

The example for the TOP dynamic:

$$\begin{array}{ccccccccc}
CDAB & \rightarrow & DABC & & DABC & \rightarrow & CABD & & CABD \\
ABCD & & ABCD & \rightarrow & BACD & & BACD & \rightarrow & ABCD \\
DACB & & DACB & & DACB & & DACB & & DACB \\
CBDA & & CBDA & & CBDA & & CBDA & & CBDA \\
\mathbf{A} & & \mathbf{D} & & \mathbf{B} & & \mathbf{C} & & \mathbf{A}
\end{array}$$

The example for the TB dynamic:

$$\begin{array}{ccccccc}
CABD & & CABD & \longrightarrow & ACDB & & ACDB \\
DBCA & \longrightarrow & BADC & & BADC & \longrightarrow & DBCA \\
DBCA & & DBCA & & DBCA & & DBCA \\
CDAB & & CDAB & & CDAB & & CDAB \\
\mathbf{C} & & \mathbf{B} & & \mathbf{A} & & D
\end{array}$$

The example for the KT dynamic:

$$\begin{array}{ccccccc}
DBAC & & DBAC & \longrightarrow & BDAC & & BDAC \\
ACBD & \longrightarrow & CABD & & CABD & \longrightarrow & ACBD \\
CDBA & & CDBA & & CDBA & & CDBA \\
DABC & & DABC & & DABC & & DABC \\
\mathbf{D} & & \mathbf{C} & & \mathbf{B} & & \mathbf{A}
\end{array}$$

The example for the SWAP dynamic:

$$\begin{array}{ccccccc}
ADCB & & ADCB & \longrightarrow & DACB & & DACB \\
DBAC & \longrightarrow & BDAC & & BDAC & \longrightarrow & DBAC \\
CBDA & & CBDA & & CBDA & & CBDA \\
\mathbf{A} & & \mathbf{B} & & \mathbf{C} & & \mathbf{D}
\end{array}$$ □

5.5 Second Order Copeland

Theorem 5. *SOC with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT and SWAP.*

Proof. The example for BR is:

$$\begin{array}{ccccccc}
BCDA & & BCDA & \longrightarrow & CBDA & & CBDA \\
DABC & \longrightarrow & DACB & & DACB & \longrightarrow & DABC \\
\mathbf{B} & & \mathbf{D} & & \mathbf{C} & & \mathbf{A}
\end{array}$$

The examples for the other dynamics are the same as those for Copeland. □

5.6 Ranked Pairs

In Ranked Pairs, as in other voting rules which output a complete ranking, a stronger convergence property could be defined for the entire ranking, but convergence is elusive even for just the top element of the ranking (the winner of Ranked Pairs).

Theorem 6. *Ranked pairs with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT and SWAP.*

Proof. The example for the BR dynamic:

$$\begin{array}{ccccccccccc}
BDCA & & BDCA & \rightarrow & DACB & & DACB & \rightarrow & CBAD & & CBAD \\
CBAD & \rightarrow & CDAB & & CDAB & \rightarrow & BACD & & BACD & \rightarrow & CDAB \\
\mathbf{B} & & \mathbf{C} & & \mathbf{D} & & \mathbf{A} & & \mathbf{B} & & \mathbf{C}
\end{array}$$

The example for the TOP dynamic:

$$\begin{array}{ccccccc}
ADBC & & ADBC & \longrightarrow & ACBD & & ACBD \\
CDAB & \longrightarrow & DABC & & DABC & \longrightarrow & CDAB \\
BCDA & & BCDA & & BCDA & & BCDA \\
\mathbf{A} & & \mathbf{D} & & \mathbf{A} & & \mathbf{C}
\end{array}$$

The example for the TB dynamic:

$$\begin{array}{ccccccc}
DABEC & & DABEC & \longrightarrow & BADEC & & BADEC \\
CADEB & \longrightarrow & CABED & & CABED & \longrightarrow & CADEB \\
ECADB & & ECADB & & ECADB & & ECADB \\
BDECA & & BDECA & & BDECA & & BDECA \\
\mathbf{D} & & \mathbf{C} & & \mathbf{B} & & \mathbf{C}
\end{array}$$

The example for the KT dynamic:

$$\begin{array}{cccccc} DABC & DABC & DABC & DABC \rightarrow & ADBC & ADBC \\ BDCA & BDCA \rightarrow & BDAC & BDAC & BDAC \rightarrow & BDCA \\ CABD \rightarrow & CADB & CADB \rightarrow & CABD & CABD & CABD \\ CBDA & CBDA & CBDA & CBDA & CBDA & CBDA \\ \textbf{B} & \textbf{C} & \textbf{D} & \textbf{B} & \textbf{A} & \textbf{C} \end{array}$$

The above example, is also for the SWAP dynamic, as all changes are of Kendall-Tau distance of one.

□

6. EMPIRICAL ANALYSIS OF OUTCOMES

We now consider the behavior of iterative voting, and the quality of its outcomes, through the results of empirical simulations. Assessing the quality of a voting method can be subtle. One general methodology is the *a posteriori* approach, to judge a rule by the quality of its outcome. Yet there is no definitive agreed-upon measure of quality of voting rules. Moreover, some voting rules have been designed with a particular measure of quality in mind, such as Maximin, ensuring the core number of supporters a candidate has, against any other one, is maximal.

Another compelling criterion is Condorcet efficiency. If a candidate is preferred to each other candidate by a majority of voters, there is reason to think it should be the winner. Thus it could be interesting to consider how often, for a given distribution of voter preferences, a rule chooses the Condorcet winner.

Furthermore, we wish to estimate social welfare. Social welfare's utility is limited in voting settings as we generally do not have the cardinal utility functions of our voters. However, as has been suggested in previous research, we can use the Borda score on the truthful preferences, in which the utility of each voter of an outcome is $m - i$, if the winner is candidate $c \in C$ which the voter ranks in place i.

As noted before, among the rules we considered in the previous section, Maximin, Copeland, Ranked Pairs, and Second Order Copeland are all Condorcet consistent. For these, therefore, we can only consider how much worse iterative voting can be than truthful voting (though, of course, we know voters do not actually always vote truthfully, so it is not as if iterative plurality is necessarily worse than a non-iterative model). For both STV and Bucklin—which are not Condorcet consistent—there is a possibility that iterative voting could have greater Condorcet efficiency than static voting. And for all rules, we can compare the Borda score of the truthful winner with the (truthful) Borda score of the equilibrium winner.

Our method is influenced by those of [21, 18, 7, 11]. Unlike [18, 7], however, which study iterative voting under restrictive dynamics (M1, M2, and k-pragmatism), we choose to analyze iterative voting under best-response dynamics. Unrestricted best-response is both computationally taxing as well as possibly cyclical. Nevertheless, as the most basic form of iterative voting, it seems to us to be of the greatest interest.

Our findings are the results of simulations of iterative voting for the six rules we have studied. Simulations were run for each rule twice, once with 10 voters and once with 25. Both runs had 4 candidates. For each set of parameters, 10,000 initial (truthful) profiles were sampled uniformly at random. Each profile evolved, for each voting rule, under best response dynamics and was run to completion (or detection of a cycle) 20 times. In keeping with the asynchronous conception of iterative voting, each of the 20 executions were developed at each step by uniformly sampling a voter with a potential move and uniformly sampling a move from all of that voter's possible best-response moves.

We begin with statistics regarding the behavior of best response in the iterative version of the various rules. Table 1 shows the average number of steps in each of the 200,000 paths considered per setting. Many of the 10,000 initial profiles were Nash equilibria, and so non-manipulable. Therefore, a significant fraction of the paths were of length zero. For convenience, we also include the average path length among non-zero-length paths. Finally, we show the maximum length among the 200,000 paths.

Voting rule	Avg. length	Corrected avg. length	Max length	Nash equilibrium share in truthful prefs
Maximin 10	2.29	8.66	143	73.5%
Maximin 25	4.97	25.19	325	80.24%
Copeland 10	6.42	19.42	269	66.55%
Copeland 25	5.51	24.69	504	77.65%
Bucklin 10	3.24	7.32	232	55.55%
Bucklin 25	4.67	11.66	301	59.94%
STV 10	0.94	4.46	97	78.84%
STV 25	1.65	8.00	182	79.32%
SOC 10	5.55	17.38	254	67.76%
SOC 25	5.66	26.00	306	78.18%
RP 10	2.13	8.55	131	75.03%
RP 25	4.18	22.83	376	81.7%

Table 1: Path Lengths

Unsurprisingly, as the number of voters grows, the probability of sampling a Nash equilibrium grows as there is a larger probability of the difference between the winner and the runner up being large enough so a single voter cannot change it (using similar reasoning, increasing the number of candidates would have increased the chances of strategic moves). On the other hand, the path lengths with more voters are longer than those with fewer voters. In the (rarer) case that an elections is close, more voters can participate in the strategic process. Copeland (and Second Order Copeland) tended to have longer paths and STV had especially short paths, but in general the NP-Complete rules did not have shorter paths than polynomial rules.

Before we begin analyzing the quality of the outcomes, we remark on an important point of relevance to the previous section.

For all of the rules, cycles occur quite rarely—the highest share of cycles was 0.57%, though most were well under 0.1%; see Figure 1. So although we have shown that all these rules *can* cycle, the frequency with which they do is very low. Copeland (and SOC), which exhibited the greatest (non-cyclic) path length, also tends to cycle more often than other rules, but it too cycles quite infrequently. STV, which has especially short paths, also cycles less frequently, but

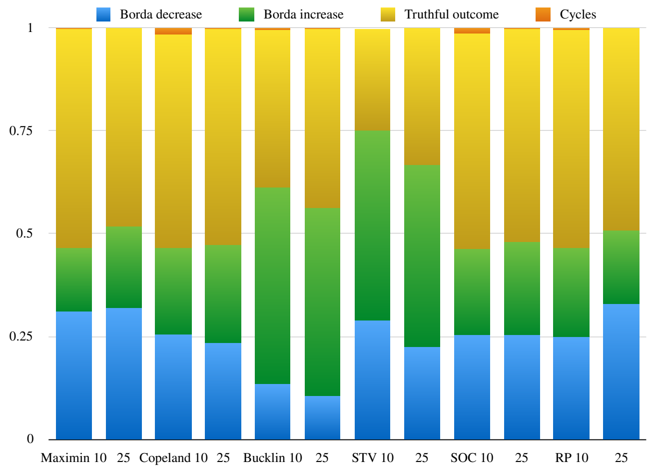

Figure 1: Out of all iterative processes (i.e., not including cases where the truthful preferences were a Nash equilibrium), the share of outcomes which decreased the Borda score, increased it, truthful outcome (of each voting rule), and the share of runs which ended in a cycle.

in general there does not appear to be a distinction in the number of cycles between NP-Complete rules and P rules.

For more voters, there are many fewer cycles, apparently because of the sparsity of cycles and uniform choice of paths.

The rarity of cycles suggest that perhaps iterative voting could be used even with these rules and best response dynamics. In the rare case of a cycle it could be detected and turned over to some cycle-breaking rule, e.g., either running the election again or deciding amongst the different outcomes in the cycle.

Turning to assess the observed outcomes of iterative voting, we first note that quite often iterative voting leads to the original outcome. Many such instances are the result of original profiles which are non-manipulable. But many are also the result of manipulations, whose equilibrium reverted to the original winner. Of the manipulable profiles (see Figure 1), the non-Condorcet consistent rules (Bucklin and STV) behave differently than the others—both of them have fewer than 50% of the outcomes truthful, but their ratio increases as voters grow, unlike the rest of the voting rules. Copeland exhibited the most consistency when increasing the number of voters.

Next we assess the change in Borda score (our proxy for social welfare) and Condorcet efficiency. As can be seen in Figure 1, once again Bucklin and STV behave significantly differently than other voting rules—for them, a significant number of outcomes increase the Borda score as compared to the truthful outcome. This is almost always not the case for the Condorcet consistent rules. However, in all Condorcet consistent rules except Maximin the share of outcomes which decreased the Borda score was close to the share of those that increased it. In all of them but Ranked Pairs this difference was decreased further when the number of voters increased. In all Condorcet consistent voting rules the average Borda score of the outcome was below that of the truthful outcome, but only slightly so—less than 1 point difference. Contrary to that, both Bucklin and STV's average Borda score was above their truthful one, and for Bucklin significantly so (above a 2 point difference).

With regard to Condorcet efficiency, we consider, for each of the 10,000 profiles, whether a Condorcet Winner existed in the original profile, whether it is selected by the voting rule, and whether it is was the outcome in a reached equilibrium. The latter is presented in Table 2 in terms of efficiency (out of 10,000) after aggregating equilibria over non-cycling paths.

Voting rule	# of Profiles with a Condorcet winner	Share of outcomes with Condorcet winner
Maximin 10	4764	0.47
Maximin 25	8413	0.82
Copeland 10	4764	0.47
Copeland 25	8413	0.81
Bucklin 10	3717	0.45
Bucklin 25	4461	0.58
STV 10	4610	0.47
STV 25	7795	0.83
SOC 10	4764	0.47
SOC 25	8413	0.81
RP 10	4764	0.47
RP 25	8413	0.82

Table 2: Comparing Condorcet Efficiency (of 10,000 profiles)

Of the 10,000 profiles sampled with n=10, there were 4764 for which a Condorcet Winner existed; among those with n=25, there were 8413 with a Condorcet Winner. As mentioned, Maximin, Copeland, Ranked Pairs, and Second Order Copeland are Condorcet-consistent, so efficiency has only one direction to move (downward). Yet it does so by very little, although slightly more when there are more voters.

Of the two rules that are not Condorcet consistent, Bucklin and STV improve their efficiency under iterative voting. These two rules also fared well under Borda criteria, suggesting that iterative Bucklin and iterative STV could be considered improvements on their static counterparts. Interestingly, Bucklin is also the most manipulable among the rules (it contained the fewest number of paths of size zero).

7. CONCLUSION AND DISCUSSION

In this work we have continued the exploration of iterative voting. We have done so in two dimensions. In the first, we expanded the set of dynamics to include some which reflect strategic behavior, but restrict best response in a, to a certain extent, natural way—whether by constraining the placement of affected candidates, or by prioritizing minor ballot changes. In the second dimension, we have ventured beyond scoring rules, and have shown that for Maximin, Copeland, Bucklin, STV, Second Order Copeland, and Ranked Pairs, iterative voting under best response dynamics does not always converge. Even after restricting the dynamics to allow voters only limited changes to their ballots they still do not always converge.

On the other hand, we have shown empirically that cycles seem to occur rather infrequently with all of these rules. Furthermore, we have shown that iterative voting, according to certain common criteria, does not perform much worse, and sometimes does better, than non-iterative voting. Notably, in non-Condorcet consistent rules—Bucklin and STV—the winners tend to improve significantly through iterative voting with regard to both their Borda score and Condorcet efficiency.

Continuation of this line of work would include analysis of convergence conditions for more voting rules and additional dynamics, with an aim towards discovering convergence dynamics, or establishing broader impossibility results. The empirical aspect of this work would benefit from expanding the analysis, for example by analyzing more distributions than we had space to include here (e.g., the Mallows model). More generally, the study of iterative voting would be greatly enhanced by incorporating voter learning into the model, and endowing voters with a greater degree of strategic (non-myopic) capabilities (early work in this direction includes Obraztsova et al. [15]).

Acknowledgements

This research has been partly funded by Microsoft Research through its PhD Scholarship Program, and by Israel Science Foundation grant #1227/12. This work was also supported by NSERC grant 482671.

REFERENCES

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

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

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

[4] R. Bredereck, J. Chen, P. Faliszewski, A. Nichterlein, and R. Niedermeier. Prices matter for the parameterized complexity of shift bribery. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, Québec City, Québec, Canada.*, pages 1398–1404, July 2014.

[5] E. Elkind, P. Faliszewski, and A. M. Slinko. Swap bribery. In *Proceedings of the Second International Symposium on Algorithmic Game Theory SAGT*, pages 299–310, Paphos, Cyprus, October 2009.

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

[7] 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 3rd International Conference of Algorithmic Decision Theory (ADT)*, pages 181–192, Brussels, Belguim, November 2013.

[8] O. Lev and J. S. Rosenschein. Convergence of iterative voting. In *The Eleventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pages 611–618, Valencia, Spain, June 2012.

[9] A. Loreggia. Iterative voting and multi-mode control in preference aggregation. Master's thesis, University of Padova, 2012. `http://www.math.unipd.it/~loreggia/pdf/tesi2012.pdf`.

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

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

[12] R. Meir, M. Polukarov, J. S. Rosenschein, and N. Jennings. Convergence to equilibria of plurality voting. In *The Twenty-Fourth National Conference on Artificial Intelligence*, pages 823–828, Atlanta, Georgia, July 2010.

[13] S. Obraztsova and E. Elkind. Optimal manipulation of voting rules. In *Proceedings of the 26th National Conference on Artificial Intelligence (AAAI)*, pages 2141–2147, Toronto, Canada, July 2012.

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

[15] S. Obraztsova, O. Lev, M. Polukarov, Z. Rabinovich, and J. S. Rosenschein. Farsighted voting dynamics. In *The First IJCAI Workshop on Algorithmic Game Theory, at the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)*, Buenos Aires, July 2015.

[16] 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 Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 993–999, 2015.

[17] 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)*, pages 1007–1013, Austin, Texas, January 2015.

[18] A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *Proceedings of the 11th International Con- ference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pages 635–644, Valencia, June 2012.

[19] R. Reyhani and M. C. Wilson. Best reply dynamics for scoring rules. In *The 20th European Conference on Artificial Intelligence (ECAI 2012)*, pages 138–144, Montpellier, France, August 2012.

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

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

Paul, Minnesota, May 2013.

[22] L. Xia, M. Zuckerman, A. D. Procaccia, V. Conitzer, and J. S. Rosenschein. Complexity of unweighted manipulation under some common voting rules. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pages 348–353, Pasadena, California, July 2009.