Iterative Voting in District-Based Elections

Omer Lev[1] , **Tyrone Strangway**[2]
[1]Ben-Gurion University
[2]University of Toronto
omerlev@bgu.ac.il, tyrone@cs.toronto.edu

Abstract

This paper examines the combinations of two decision-making models. The first, district-voting, is a commonly used method both for political elections (in the US, UK, Canada etc,), as well as in any setting with multiple groups: each district or group have an election, and the election results are aggregated in a second stage. The second model is that of iterative voting. In iterative voting, participants are assumed to adapt their vote to the current situation as they see it (e.g., via polls). Voters attempt to increase the chance the winner will be a candidate they prefer, and in order to do so, they vote differently than they believe.

We show that while some voting rules are known to converge when using iterative voting, the process no longer converges when using iterative voting in district-based elections. However, we use simulations to compare the election outcomes using multiple voting rules and multiple distributions, showing which voting rules result in higher quality winners, and which voting rules seem to be less suited for district-base iterative voting.

1 Introduction

Voting, as a method to aggregate the various views of agents – human or artificial ones – is a crucial part in any system which relies on multiple diverse and independent agents, as their viewpoints need to be combined to create a coherent outcome. A particularly interesting subset of aggregation methods are district-based elections. In such elections, the agents are divided into several parts (commonly called "districts"), and their votes are combined first in each district, which then submits its election outcome to a second stage, that chooses the most common outcome as the final result. Much district-based election research, even its computational parts [Bachrach *et al.*, 2016; Lewenberg *et al.*, 2017; Borodin *et al.*, 2018] take its inspiration from political elections – district-based elections choose country leaders (as in Westminster system countries, where the districts are parliamentary constituencies) or can create new laws (e.g., the US Congress). However, district-based elections are applicable in non-political settings as well – organizational decisions, for example, may be decided by a vote between the heads of the organization's sub-units, each with its own decision process; a sensor array with different types of sensors (each type reaching its own combined reading); and other settings in which agents have a natural division between them.

Understanding how individual preferences interact with a voting system – in particular one as relatively little explored as district-based one – is of paramount importance to all participants, both voters and potential winners. However, the Gibbard-Satterthwaite theorem [Gibbard, 1973; Satterthwaite, 1975] showed that under any reasonable voting system, voters might be incentivized to misreport their preferences. In order to understand the voting behavior of agents we wish to find the stable states (or *Nash equilibria*), for each election – states in which voters, after strategizing and knowing the outcome, do not wish to further change their votes. However, the standard Nash equilibrium can be an unsatisfactory solution concept when it comes to voting. For example, if one candidate was universally loathed and ranked last by all voters' private beliefs, the state where it is all voters' top choice is a Nash equilibrium, despite being quite an unlikely voting outcome.

Instead, the iterative voting model was suggested by [Meir *et al.*, 2010], and it considers a subset of Nash equilibria: those reachable by an iterative process, in which voters may change their vote (one at a time) to make the outcome more preferable to them. This model does not shy away from the inevitable strategic voting by agents, and attempts to stay grounded in what voters might plausibly do. The solution concept suggested by this model (and the papers that furthered investigated the model) is that when starting from truthful preferences, the equilibrium reached by the process of iterative voting might represent a more plausible idea of a potential end state of the election.

This paper wishes to combine these two research strains – district-based voting as a common setting for decision making, and iterative voting as a useful analytical tool of possible voter behavior. Such a combination has some relevance in political settings (e.g., US presidential elections, in which the districts are the states, and the iterative dynamic is due to polling), but it might be more interesting exploring it in organizational settings: small organizational units trying to reach a decision, conducting internal debate (which often incorporates an implicit iterative dynamic) on a set of options, with each unit reporting its decision, and they are ultimately aggregated. Such a dynamic can be seen in various organizations, for example, the decision processes in some youth-groups (similar to boy and girl scouts) in some countries, aggregating decisions from local branches.

The combination of district and iterative voting is not always smooth – known results from iterative voting literature do not necessarily apply in district-based settings, which can introduce unintuitive behavior among voters. Consider a voter whose preferred candidate is uncompetitive in their district, but a strong candidate in other districts. They may prefer to vote for a more locally competitive candidate which is not competitive in other districts, and thus deprive their preferred candidate's competition of a district representative, even if it means their own district is represented by a candidate they dislike. This non-monotone logic does not hold in previously studied iterative voting systems, increasing voters' array of possible strategies. Other voters, on the other hand, might be less concerned with the identity of the winning candidate in the district elections, but rather focus on electing a representative they feel will be a strong advocate for them. In this work we examine both types of voters:

- **Globally minded voters**: These voters are concerned with the overall winner outcome. Hence, they will vote strategically to ensure a preferred candidate holds a plurality of the districts, even if this means they are represented by an inferior candidate.
- **Locally minded voters**: These voters are concerned with their local representative, they will vote strategically to ensure their preferred candidate represents them, even if this means the overall winner will not be to their liking.

We begin by showing that previously known results on the convergence of iterative voting no longer hold in district-based settings. We then empirically study these equilibria by employing simulations of iterative voting, comparing various voter distributions and voting methods[1].

2 Related Work

This paper joins two different research threads. The first, and more well established, is the one of iterative model. While work on different models involving iterative updating of preferences existed before it (e.g., [Airiau and Endriss, 2009]), the model we are exploring and expanding was introduced in [Meir *et al.*, 2010], which both introduced the model and showed convergence results for plurality. Convergence for veto (and non-convergence for other scoring rules), as well as the importance of linear ordered tie-breaking was shown in [Lev and Rosenschein, 2016]. Lack of convergence in the main non-scoring rule voting mechanisms was shown in [Koolyk *et al.*, 2017]. [Brânzei *et al.*, 2013] showed the relative strength of the candidates which win in iterative voting, while [Rabinovich *et al.*, 2015] analyzed the complexity of finding Nash equilibria reachable through an iterative process.

The basic iterative voting model was expanded in various ways – [Obraztsova *et al.*, 2015b] discusses different dynamics that can be pursued by voters (a topic also discussed in [Grandi *et al.*, 2013; Koolyk *et al.*, 2017]), while [Rabinovich *et al.*, 2015] discussed lazy and truth biased voters, and [Tsang and Larson, 2016] explored graph-related social network concerns. [Meir *et al.*, 2014] (followed up by [Meir, 2015; Lev *et al.*, 2019]) suggested a more elaborate model, encompassing uncertainty by voters as to what state they are in. [Obraztsova *et al.*, 2015a] used a similar idea to deal with the myopia of voters.

The second research thread this paper connects to is one on district-voting, which has to do with the effect districts have on voting outcomes. While this has been explored with regards to political contexts (usually, country specific) in political science, history and law (e.g., [Erikson, 1972; Schuck, 1987; Wang, 2016]), it has only recently garnered interest in computational fields. [Bachrach *et al.*, 2016] deals with the fundamental representability issues raised by using districts, while most other papers deal more specifically with gerrymandering – the manipulation of electoral districts to the benefit of some candidate [Bervoets and Merlin, 2012; van Bevern *et al.*, 2015; Lewenberg *et al.*, 2017; Pegden *et al.*, 2017; Borodin *et al.*, 2018]. [Lev and Lewenberg, 2019] examined an iterative process in the context of districts, but it only dealt with agents moving between districts, but that is not a possibility in many settings, where the available strategy has to do with who the ultimate winner is.

3 Preliminaries

An election consists of a set of voters V of size n and a set of candidates (or options) C of size m. In our case, our voters are divided between a set of districts $D = \{D_1, D_2, \ldots, D_k\}$. We shall assume the district sizes are equal, so n is divisible by k. Let $\pi(C)$ be the set of all linear orders of the elements of C, and each vote $v \in V$ is associated with such an order $\succ_v \in \pi(C)$, which determines the internal (and unchanging) preference order of each agent over the candidates.

The voting system in each district uses a function $\tilde{f} : (\pi(C))^{\frac{n}{k}} \to 2^C$, which takes as an input all of the district's voters' votes (which do not necessarily have to match their internal preferences), and outputs a set of candidates. Then, a function $t : 2^C \to C$ is used to break ties and select a single winner. Hence, the winner in each district is determined by $f : \tilde{f} \circ t$. Finally, the ultimate election winner is determined in a second stage, which choses the candidate that won most districts (i.e., plurality)[2].

We shall explore the following voting rules (the $\tilde{f}$ above):

Plurality Each voter gives a single point to a particular candidate. The candidate with the most points is the winner (or winners).

Veto Each voter gives a point to all candidates except one (which the voter is "vetoing"). The candidate with the

[1] We significantly enhance and expand the code base from http://www.preflib.org/tools/ivs.php, initially used in [Meir *et al.*, 2014].

[2] Naturally, other voting methods may be considered here, but we chose the system analyzed in previous district-voting literature, and which seems to be the one commonly used.

most points is the winner (or winners).

Borda Each voter gives a set of points to candidates, such that if there are m candidates, the candidate ranked at the ith location (1 being most favored) gets $m-i$ points. So the top candidate gets $m-1$ points, the second ranked gets $m-2$, decreasing point by point until the penultimate candidate gets 1 point, and the least favored candidate gets 0 points. The candidate with the most points is the winner (or winners).

Copeland Each candidate is given a score based on how many candidates it beat in pairwise comparisons. So for a candidate $c \in C$, it gets a point for every candidate $c' \in C$ such that $|\{v \in V | c \succ_v c'\}| > |\{v \in V | c' \succ_v c\}|$, and loses a point for every candidate $c' \in C$ such that $|\{v \in V | c \succ_v c'\}| < |\{v \in V | c' \succ_v c\}|$. The candidate with the maximal score is the election winner (or winners).

STV (Single Transferable Vote) Each voter submits their full ranking, but only their first choices are examined at first, and added up for each candidate (as in plurality). Until there is a candidate which receives 50% or more of the vote, the candidate which received the fewest votes is eliminated, and those voting for them will transfer their vote to the highest candidate in their ranking that has not yet been eliminated.

The voting rules above include examples of scoring rules (in which voters allocate points – as in plurality, veto and Borda) and of tournament-based rules (based on pairwise comparisons – like Copeland). The tie-breaking rule (t) we shall use throughout this paper is the one used in almost all previous iterative voting and district-voting papers – lexicographic (i.e., there is a linear order – an element of $\pi(C)$ – that determines the tie-breaking).

One further definition we shall note is that of the Condorcet winner. This candidate is supported by a majority of voters compared to any other candidate. That is, candidate $c \in C$ is a Condorcet winner if for any other candidate $c' \in C, c \neq c'$, $|\{v \in V | c \succ_v c'\}| > |\{v \in V | c' \succ_v c\}|$. A Condorcet winner does not necessarily exist, and is not guaranteed to be a winner by many voting rules even when it does, but it can be viewed as a "good", or desirable, winner, which can be used, in some instances as a proxy to how good of a candidate is being selected.

3.1 Iterative Voting without Districts

The iterative voting model is one which tries to describe the dynamic by which voters change their vote to maximize their utility (determined by how high in their own internal rankings the ultimate winner is), according to what others are voting. Once knowing what is the election outcome, the voters know which candidates are viable, and if they can change the outcome, they change their vote so they can make someone they prefer over the current winner victorious. That is, if the current vote by all voters is $\vec{p}$, with voter $v \in V$ voting p_v (and the vector of other voters except v being $\vec{p}_{-v}$), then v can make a move if there is a vote $a \in \pi(C)$, $a \neq p_v$ such that $f(\vec{p}_{-v}, a) \succ_v f(\vec{p})$. If there are multiple such moves then the *best response* move is v's most favorable outcome, that is, a vote $a \in \pi(C)$ is a best response if $f(\vec{p}_{-v}, a) \succ_v f(\vec{p})$, and if there is no vote $b \in \pi(C)$ such that $f(\vec{p}_{-v}, b) \succ_v f(\vec{p}_{-v}, a)$.

If this process ends in a state such that no voter wants to change their vote, it is a stable state, also called a *Nash equilibrium* (though note that not all Nash equilibria may be reachable by such a process). If for a particular voting rule the iterative process always converges to a stable state, we say the voting rule converges. Otherwise, for non-converging voting rules, there are settings where the iterative process may never end.

So far, on regular, non-district iterative voting, the following theorems are known

Theorem 1 ([Meir *et al.*, 2010] Theorem 3)**.** *Iterative plurality (with deterministic tie-breaking), when voters are myopic (i.e., only examine the current state, and not predict ahead) and pursue a best-response strategy, will always converge to a Nash equilibrium.*

Theorem 2 ([Lev and Rosenschein, 2016] Theorem 3)**.** *Iterative veto (with deterministic tie-breaking), when voters are myopic (i.e., only examine the current state, and not predict ahead) and pursue a best-response strategy, will always converge to a Nash equilibrium.*

Theorem 3 ([Lev and Rosenschein, 2016] Theorem 4)**.** *For any scoring rule except plurality and veto, iterative voting, even with deterministic tie-breaking, and even when voters are myopic (i.e., only examine the current state, and not predict ahead) and pursue a best-response strategy, have cases in which they will not converge to a Nash equilibrium.*

Theorem 4 ([Koolyk *et al.*, 2017] Theorems 1-6)**.** *Even with deterministic tie-breaking, and when voters are myopic (i.e., only examine the current state, and not predict ahead) and pursue a best-response strategy, iterative maximin, iterative Copeland, iterative Bucklin, iterative STV, iterative second-order-Copeland, and iterative ranked pairs*[3] *will all have cases in which they will not converge to a Nash equilibrium.*

4 Iterative Voting in District-Based Settings Model

When considering what are voters aiming for in district-based settings we examine two options:

Global These voters, as in the "regular", non-district, version of iterative voting, wish to make the overall winner as favorable to them as possible. Hence, they will choose a vote that, regardless of what happens in their own districts, makes the global winner someone they prefer.

Local These voters are only concerned with which candidate wins their own district, and their strategic moves does not take into considerations which candidate is the overall winner.

If all voters are local, then district-based elections are basically run as k separate elections, as no voter changes their vote depending on who the overall winner is. All results regarding iterative elections hold in this case, as the districts are not inter-connected, as far as voters are concerned.

[3]We do not define voting rules which are not examined in this work.

However, when voters are global, their strategic considerations change. In particular, moves which would not have been rational in non-district iterative voting become plausible. For example, voters may chose to vote for a candidate they dislike, and prefer the current winner over, as it would effect the overall winner's identity in a way they prefer (see Example 5).

Example 5. *We have 3 districts (each using plurality) and 3 candidates, a, b, and c (ties broken lexicographically). In one of them, the winner is candidate a, and in another the winner is candidate b. In the third district $\lfloor \frac{n}{2} \rfloor + 1$ voters vote for c, and the rest vote for b. Hence, the winner of this district is candidate c, and the overall winner is candidate a.*

Suppose voter v, in the third district, has the preference order $c \succ b \succ a$. If this voter is a local voter, they do not change their vote, as their favorite candidate is the district winner, and they only care for that. However, if v is a global voter, they will change their vote to b, making the district winner b (which, in their preferences is a worse *candidate than c which the previously voted for), making the overall winner candidate b, which they prefer over the previous overall winner, candidate a.*

5 Convergence

For local voters, all previous convergence results (Theorems 1, 2, 3, 4) still hold. Hence, iterative plurality or iterative veto in district-based elections will converge to Nash equilibria with local voters, and other scoring rules (and well known non-scoring rules) will not converge.

When voters are global, any non-convergence result still holds (Theorems 3, 4) since all candidates may have the same number of districts supporting them, and one single district's winner will determine the outcome. In this case, this district's global voters care about the district's outcome, as it will determine the overall winner. So global voters behave like local voters, and if they might end up in a cycle, so will global voters. We formalize this understanding in a theorem:

Theorem 6. *When voters are global or local, and are pursuing a myopic best-response strategy, for any scoring rule that is not plurality or veto, and for maximin, Copeland, Bucklin, STV, second-order-Copeland, and ranked pairs, iterative voting in district-based elections has cases where they will not converge to a Nash equilibrium, and they will end up in a cycle.*

For the convergence results, however, for iterative plurality and veto, we need to show if these still hold. Sadly, they do not:

Theorem 7. *When voters are global and are pursuing a myopic best-response strategy, both iterative plurality and iterative veto in district-based elections have cases where they do not converge to Nash equilibria and may end up in a cycle.*

Proof. We begin with the proof for iterative plurality. The construction here is somewhat similar to one given in [Meir *et al.*, 2010] for non-convergence of *non-best-response* strategies. However, in our construction, the voters use best-response strategies, which under district-voting end up as being non converging.

Consider a setting with 9 districts, each with 11 voters, and 4 candidates – a, b, c, d (using lexicographic tie-breaking. Each of a, b, d have the support from 2 districts, while candidate c has the support from three districts, making c the global winner. An overview of the cycle can be seen in Figure 1.

We shall look at agents in three districts. Districts I + II, currently supporting candidate d, and district III, currently supporting candidate c. District I is composed of one a voter, 3 voters supporting b, 3 voters supporting c, and 4 voters supporting candidate d. District I has a voter x, whose preference order is $a \succ b \succ c \succ d$.

District II is composed of one b voter, 3 voters supporting a, 3 voters supporting c, and 4 voters supporting candidate d. District II has a voter y whose preference order is $d \succ a \succ b \succ c$.

District III is composed of one d voter, 3 voters supporting a, 3 voters supporting b, and 4 voters supporting c. District III has a voter z whose preference order is $c \succ b \succ a \succ d$.

Let us now examine these voters moves:

1. Voter x, who changes their vote to b, making the global winner candidate b (since they cannot change their district to have a as the winner, making b the winner is their best option).
2. Voter y is now able to change their vote to a (as it prefers a winning to b, and it cannot make its favorite, candidate d the winner), making that candidate the winner of district II, and also the global winner.
3. Voter z can now change their vote to candidate b, making it the district winner as well as the global winner.
4. Voter x now changes their vote to c, making it district I's winner, but it allows the global winner to become a (agent x cannot make a the winner of its district, and a is preferable to b as a global winner for this voter).
5. Voter z reverts to its truthful vote, for c, making in the winner of district III, and the global winner as well.
6. Voter x now reverts to voting for candidate b (as it did in step 1), making a the global winner.
7. We are now back in the state at the beginning of step 3, and steps 3-6 can repeat ad-infinitum.

The example for iterative veto is quite similar (and follows the same overview, in Figure 1), only containing slightly different voters to allow for similar strategic changes (which district votes for which candidates). Therefore, district I is composed of 1 voter vetoing d, 2 voters vetoing b, 2 voters vetoing c, and 6 voters vetoing a. We consider two voters in this district – x' with the preference order $b \succ d \succ c \succ a$ and x with the preference $a \succ b \succ d \succ c$.

District II is composed of 1 voter vetoing d, 2 voters vetoing a, 3 voters vetoing c, and 5 voters vetoing b. We consider one voter in this district – y with the preference $d \succ a \succ b \succ c$.

District III is composed of 1 voter vetoing c, 2 voters vetoing b, 3 voters vetoing a, and 5 voters vetoing d. We consider one voter in this district – z whose preference order is $c \succ b \succ a \succ d$.

Let us now examine these voters moves:

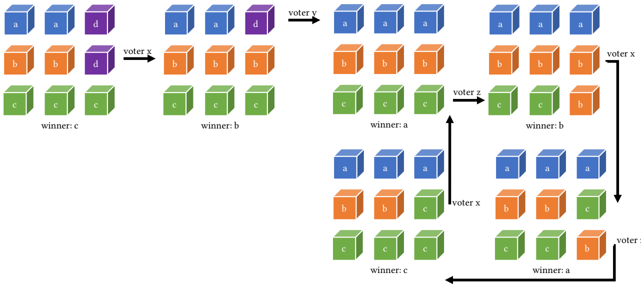

Figure 1: An overview of the voters causing the cycle for both plurality and veto in iterative voting when using district-based elections. Each box is a district, with the letter on it representing the district winner. Voter x in District I makes a change, leading voter y, in district II to make another. Then, a cycle begins, with voters x and voter z (from district III) with alternating moves.

1. Voter x', who changes their veto to d, makes b district I winner and also the global winner.
2. Voter y is now able to change their veto to d, making a the district and global winner (as it prefers a winning to b, and it cannot make its favorite, candidate d the winner).
3. Voter z can now change their vote to veto candidate c, making candidate b the district winner as well as global winner.
4. Voter x now changes their veto to b, making c the district I's winner, but allowing the global winner to become a (agent x cannot make a the winner of its district, and a is preferable to b as a global winner for this voter).
5. Voter z now vetoes b, making c the winner of district III, and the global winner as well.
6. Voter x now vetoes candidate c, making candidate b the district winner and candidate a the global winner.
7. We are now back in the situation at the beginning of step 3, in which agent z had an incentive to veto c, and steps 3-6 can repeat ad-infinitum.

□

We note that theorem 7 does not really require all voters to be global, but only 3 global voters in 3 different districts. If we allow for starting at a non-truthful state, we need only 2 global voters in 2 different districts.

6 Simulation Setup

In an effort to better understand the properties of equilibria arising from iterative voting, we turn to simulations. In order to better perceive the impact of voter type and voting rule on iterative voting we focus on a particular setting of the parameters, 25 voters divided into 5 equal sized districts with the same 5 candidates competing in each district.

Voter profiles are generated by sampling from either a uniform distribution over all preferences or from a uniform distribution over all single-peaked preferences. Moreover, we consider both games containing only globally minded voters and games containing only locally minded voters. For each combination of voting rule, voter preference type (uniform or single-peaked) and voter aims (global or local) we sample 1000 sets of preferences. Since iterative voting is a non-deterministic process, we run the game until convergence 200 times (terminating a run after 1000 steps). We use a deterministic tie-breaking rule.

If a locally minded voter has more than one response profile which leads to the same winner in their district they will opt to take the profile with the closest Kendall-tau distance[4] to their current profile. If a globally minded voter has more than one response profile which can lead to the same global winner they will break ties by opting for the profile which leads to their most preferred winner within their district (subsequent ties are also broken using the Kendall-tau distance to their current profile).

7 Results

We examine the simulation results in two regards: how involved was the iterative process, and what is the quality of the outcome. Using this we can better understand systems which already work in this manner, better comprehend their outcomes (and how they are reached). We can also view the problem with a mechanism design approach, seeing which voting systems are most suitable for a district-based system.

7.1 The Iterative Process

The first issue to note is that despite the theoretical results, *in no case did a simulation end up in a cycle*. That is, despite Theorems 6 and 7 that convergence is not guaranteed, in almost all cases, convergence was reached, and in the few where it was not, it was because number of states exceeded our threshold limits, and even that number was not large (it was less than 1% in almost all cases – mostly significantly less – apart from Borda global voters, uniformly distributed, where 5% of runs timed out after 1000 moves). This is consistent with what was observed in [Koolyk *et al.*, 2017], in which voting rules which were shown to not converge, did not actually end up in a cycle in any randomly-created preference set and dynamic. In general, Borda based systems struggled most with convergence, with 20%-40% of election settings having at least one run (each setting had 200 runs) which failed to converge (other voting rules had far less issues, in many cases no run timed-out at all).

In almost every setting the iterative process settled into equilibrium within twenty steps on average. Once again, a noticeable exception were variants of Borda which averaged a few hundred steps. The iterative process was relatively focused, and on average, for each election setting, about half of the candidates won no equilibrium at all. Overall, voters with single-peaked preferences tended to have fewer winners in each setting, and we conjecture that the existence of a Condorcet winner in each district tended to limit the opportunities for iterative voting in the single-peaked preferences case.

As described earlier, the district-based settings introduce a non-monotone nature to the iterative process. For local voters, this is inadvertent – they do not care for the overall win-

[4]The Kendall-tau distance between two profiles $a, b \in \pi(C)$ is the number of pairwise disagreements between a and b.

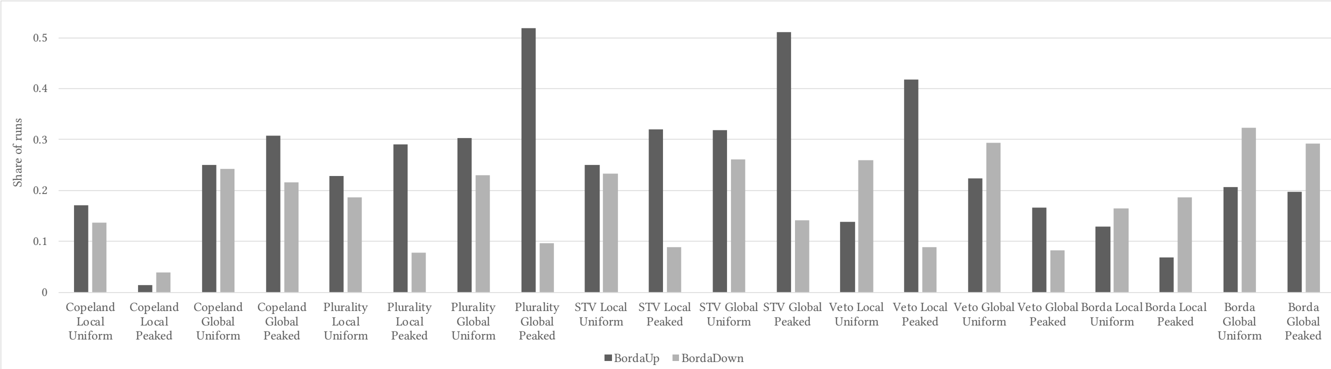

Figure 2: Examining the change in the Borda score of the overall winner, and whether it increased or decreased (from the truthful winner to the iterated winner).

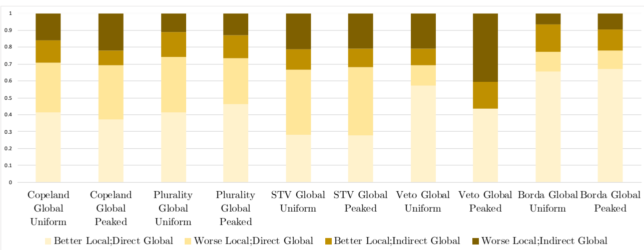

Figure 3: The fraction of move types of globally minded voters.

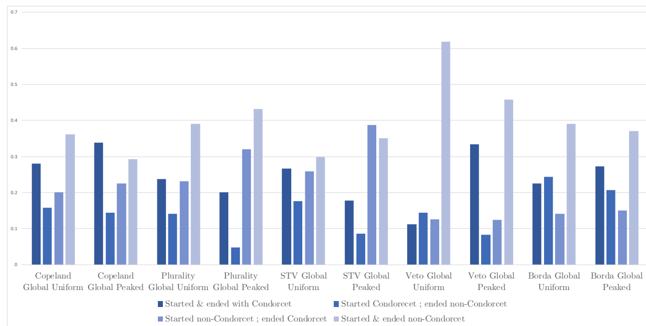

Figure 4: How the Condorcet winner emerges or disappears in runs where iterative voting took place and a Condorcet winner existed.

ner, but their moves may change the winner's identity. However, global voter moves (see Figure 3) may change the outcome of their own district's winner in a deliberate way to hurt a different candidate, rather than to strengthen their desired candidate (a *destructive* move, rather than a *constructive* one). In all cases (except veto with single-peaked preferences) constructive moves were the majority, although in many cases these constructive moves came at the cost of a worse local winner. Destructive moves were not entirely uncommon, constituting at least 20% of the moves in any setting. The equivalent data of how locally minded voters impacted the global winner indicate that only rarely was the share of moves that hurt the overall winner above 30%. However, this also has to do with cases where local voters' manipulation had no effect on the overall winner, when the district winner was not a winner or a runner-up in the overall election.

7.2 Winner Quality

In this section we will examine the quality of the winner found by the iterative voting process. How one measures quality can be tricky since there is no agreed upon definition. This is further complicated by the different voting rules, each with a different goal. For example, Copeland's method rewards candidates who win in pairwise matchups, and cares little for the margin of victory.

One way of measuring quality, used in [Thompson *et al.*, 2013; Koolyk *et al.*, 2017], is how often the **truthful winners** emerges after iterative voting. The truthful winner is the candidate the mechanism intended to pick, if they reemerge as the winner it indicated the voting rule achieved its goal despite the voters' strategizing. For globally minded voters, the truthful winner reemerged 40% to 50% of the time. The variance seems to be a function of the voting rule and not the distribution (veto was something of an exception). This seems consistent with the results of [Koolyk *et al.*, 2017], which found the voting rule tended to determine how often the truthful winner was the iterated winner[5]. Interestingly, locally minded voters, who do not optimize for the global winner, exhibited a similar pattern where the voting rule often determined how the truthful winner emerged as the iterated winner. It was slightly more common for the truthful winner to be the final winner, but this difference seems to be caused by local moves having no impact on the global winner.

Another metric for winner quality is the **Borda score** of the winner ([Koolyk *et al.*, 2017; Bachrach *et al.*, 2016; Meir *et al.*, 2014] all used this as a measure of quality). We are interested in the change in Borda score between the initial (truthful) winner and final winner. Unsurprisingly most settings with globally minded voters saw a high fraction of increasing Borda scores (see Figure 2). Surprisingly, in most settings with locally minded voters there was a large fraction of increasing Borda scores. The main exception were all variants of Borda, which saw a higher fraction of decreasing

[5]Their work adjusted voter's response dynamics, also comparing them between single-peaked or uniform distributions.

scores[6] (as did veto with uniform voters).

A possible partial explanation for the high fraction of increasing Borda scores with globally minded voters, especially in single-peaked cases, has to do with Condorcet winners[7] (their existence is guaranteed in the single-peaked case). There appears a positive correlation between a high fraction of increasing Borda scores and a high ratio of starting with a non Condorcet winner and moving to the Condorcet winner vs starting with the Condorcet winner and moving to a non Condorcet winner (see Figure 4).

As a final metric we examine the **Condorcet winner**, this candidate, preferred by a majority in all pairwise matchups, is a generally desirable candidate (when they exist). We consider the case where an overall Condorcet winner existed (i.e., when all voters, from all districts were examined together). As can be seen in Figure 4, in almost all cases it was more likely to start with a non Condorcet winner and end with the Condorcet winner than the reverse (starting with the Condorcet winner and finishing with a non Condorcet winner). This is despite having at least one alternative strong candidate (the truthful winner) and that the district-based elections we study are not Condorcet-consistent. This was fairly prominent with globally minded voters with single-peaked preferences using plurality and STV. Borda was again the exception, as with globally minded voters it was more likely to leave the Condorcet winner than to find them.

8 Discussion

We explored how district-based elections, a commonly used election mechanism, work when voters employ iterative voting, updating their vote according to what they learn of the vote outcomes. This solution concept raises potential end-states for elections which try to avoid obviously impossible or unreachable states (as Nash equilibria sometimes are). We explore how the district-based structure changes the convergence results, and show both an answer for this (there is a change – convergence which was assured in the non-district case is no longer guaranteed), as well as showing that in many cases it does not matter, as simulations do not end up in a cycle, indicating these are edge cases.

The simulations themselves help us focus on which voting systems are desirable or appropriate for such settings. Somewhat surprisingly, the Borda system with globally minded voters, despite having starting off with relatively high quality winner ends up as being rather mediocre in our setting – it is common for overall winners to be worse than in the starting position (using all of our criteria – is the outcome the truthful winner; the Borda score of the winner; and how likely it is to end up with a Condorcet winner). On the upside, both plurality and STV[8] seem to have very desirable properties for these elections, with STV doing especially well with an electorate sampled from a single-peaked distribution.

This work is the first step in better understanding district-based elections. While most research focused on manipulations of the districts themselves, here we try to better grasp the dynamics which shape these elections outcomes. There is plenty more to do – from a wider variety of voter distributions, through heterogenous voter population, to an examinations of more voting dynamics: we have examined (as most iterative-voting simulations have) the behavior and outcomes of a best response strategy, but other dynamics might yield further insights into the effects of districts on voters behaving iteratively. Finally, merging these simulations with other concepts suggested for district-based elections (such as the "price of districting", suggested by [Bachrach *et al.*, 2016]) might lead to further interesting results.

[6] This is not as obvious as it sounds – the winner is not necessarily the candidate with the highest Borda score, due to the district-based system. [Bachrach *et al.*, 2016] showed the Borda score of the global winner can be about $\frac{1}{m^2}$ (in our case – $\frac{1}{25}$) of the candidate with the highest Borda score.

[7] The Condorcet winner need not have the highest Borda score, but they tend to have a relatively high one.

[8] Incidentally, both systems used in district-based voting in the political world – STV in Australia and Ireland, and plurality almost anywhere else with district-based elections.

References

[Airiau and Endriss, 2009] Stéphane Airiau and Ulle Endriss. Iterated majority voting. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT)*, pages 38–49, Venice, Italy, October 2009.

[Bachrach *et al.*, 2016] Yoram Bachrach, Omer Lev, Yoad Lewenberg, and Yair Zick. Misrepresentation in district voting. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 81–87, New York City, New York, July 2016.

[Bervoets and Merlin, 2012] Sebastian Bervoets and Vincent Merlin. Gerrymander-proof representative democracies. *International Journal of Game Theory*, 41:473–488, 2012.

[Borodin *et al.*, 2018] Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. Big city vs. the great outdoors: Voter distribution and how it affects gerrymandering. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 98–104, Stockholm, Sweden, July 2018.

[Brânzei *et al.*, 2013] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel 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.

[Erikson, 1972] Robert S. Erikson. Malapportionment, gerrymandering, and party fortunes in congressional elections. *American Political Science Review*, 66(4):1234–1245, 1972.

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

[Grandi *et al.*, 2013] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby 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, Belgium, November 2013.

[Koolyk *et al.*, 2017] Aaron Koolyk, Tyrone Strangway, Omer Lev, and Jeffrey S. Rosenschein. Convergence and quality of iterative voting under non-scoring rules. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 273–279, Melbourne, Australia, August 2017.

[Lev and Lewenberg, 2019] Omer Lev and Yoad Lewenberg. ”reverse gerrymandering”: a decentralized model for multi-group decision making. In *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI)*, Honolulu, Hawaii, January-February 2019.

[Lev and Rosenschein, 2016] Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative scoring rules. *Journal of Artificial Intelligence Research (JAIR)*, 57:573–591, December 2016.

[Lev *et al.*, 2019] Omer Lev, Reshef Meir, Svetlana Obraztsova, and Maria Polukarov. Heuristic voting as ordinal dominance strategies. In *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI)*, Honolulu, Hawaii, January-February 2019.

[Lewenberg *et al.*, 2017] Yoad Lewenberg, Omer Lev, and Jeffrey S. Rosenschein. Divide and conquer: Using geographic manipulation to win district-based elections. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 624–632, São-Paulo, Brazil, May 2017.

[Meir *et al.*, 2010] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Convergence to equilibria of plurality voting. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 823–828, Atlanta, July 2010.

[Meir *et al.*, 2014] Reshef Meir, Omer Lev, and Jeffrey 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, California, June 2014.

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

[Obraztsova *et al.*, 2015a] Svetlana Obraztsova, Omer Lev, Maria Polukarov, Zinovi Rabinovich, and Jeffrey S. Rosenschein. Farsighted voting dynamics. In *AGT@IJCAI*, Buenos Aires, Argentina, July 2015.

[Obraztsova *et al.*, 2015b] Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R. Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*, pages 993–999, Austin, Texas, January 2015.

[Pegden *et al.*, 2017] Wesley Pegden, Ariel D. Procaccia, and Dingli Yu. A partisan districting protocol with provably nonpartisan outcomes. ArXiv:1710.08781, October 2017.

[Rabinovich *et al.*, 2015] Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey 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.

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

[Schuck, 1987] Peter H. Schuck. The thickest thicket: Partisan gerrymandering and judicial regulation of politics. *Columbia Law Review*, 87(7):1325–1384, 1987.

[Thompson *et al.*, 2013] David Robert Martin Thompson, Omer Lev, Kevin Leyton-Brown, and Jeffrey S. Rosenschein. Empirical aspects of plurality election equilibria. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 391–398, St. Paul, Minnesota, May 2013.

[Tsang and Larson, 2016] Alan Tsang and Kate Larson. The echo chamber: Strategic voting and homophily in social networks. In *Proceedings of the 15th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 368–375, Singapore, May 2016.

[van Bevern *et al.*, 2015] René van Bevern, Robert Bredereck, Jiehua Chen, Vincent Froese, Rolf Niedermeier, and Gerhard J. Woeginger. Network-based vertex dissolution. *SIAM Journal on Discrete Mathematics*, 29(2):888–914, 2015.

[Wang, 2016] Samuel S.-H. Wang. Three tests for practical evaluation of partisan gerrymandering. *Stanford Law Review*, 68:1263–1321, June 2016.