

“Reverse Gerrymandering”: a Decentralized Model for Multi-Group Decision Making

Omer Lev and Yoad Lewenberg

Abstract

District-based manipulation, or gerrymandering, is usually taken to refer to agents who are in fixed location, and an external division is imposed upon them. However, in many real-world setting, there is an external, fixed division – an organizational chart of a company, or markets for a particular product. In those cases, the agents might desire a “reverse gerrymandering”, as each of them, on their own, wish to maximize their influence in more parts of the company, or resources seeking to be allocated to areas where they will be most needed.

In this paper we explore such an iterative dynamic, finding that allowing this decentralized system results, in some particular cases, in a stable equilibrium, though in general, the setting may end up in a cycle. We further examine how this decentralized process affects the social welfare of the system.

1 Introduction

In October 2016, just before the US presidential elections, the New York Times published an article titled “Go Midwest, Young Hipster” [9]. In it, the author emphasized how crucial elections could be won if supporters of a party would move from where they are concentrated to areas where they are more sparse. While this may be a ridiculous suggestion for national elections and the large number of people involved (and the author did not claim otherwise), in other settings such a proposition is not as preposterous.

An obvious, more realistic, setting of a similar idea in organizations is that of change agents [15]. When trying to change a corporate behavior, change agents move around the organization, trying to form coalitions to bring about some form of a change they are advocating for, so that it encompasses the whole company.¹ Many people in large bureaucratic organizations are familiar with the sets of committees in which decision are made. In such organizations, people seeking to push an idea (or a person) try to gain influence in multiple committees.

A wider angle reveals this to be, more generally, a model of decentralized resource allocation. Understanding what would be the outcome if many of our resources were autonomous and could attempt to allocate themselves to their optimal destination, while their ultimate goal is succeeding in as many destinations as possible. This is a desirable feature in many industrial IoT applications, but more concretely, one can think of a medical resource (or robot) deciding which of multiple field operations it should go to without central direction. It can be a batch of advertising resources directed to the appropriate product/market without the company HQ deciding on it; or it can be grassroots political campaigners going to volunteer in the place they think they will be able to overwhelm rival campaigns, with no need for their candidate telling them where to go.

All these settings have to do with agents realizing whether they are needed in their current location, and if not – looking for a better place they can contribute in, possibly

¹For example, in pushing to change Apple’s corporate culture in its development efforts after Steve Jobs’ NeXT was acquired, NeXT people were placed in various units across the company, in order for them to aid similarly minded Apple personnel in bringing about better practices.

only seeing the places that are reachable by them from their area. In this work, we are less interested in how they might reach their destination, but in the fundamental question of what are the effects of using a decentralized system, instead of a hierarchal, centralized one.

Since this setting is, fundamentally, equivalent to an attempt by supporters of some candidate – which can be a person, a policy, or an idea – to spread themselves so they can encourage other supporters and together form a majority in many sub-units or districts, we will use the terminology of elections and voting. Voting and elections have been widely studied in artificial intelligence as natural tools for aggregating preferences of independent, self-interested agents in multiagent systems. Agents have a preference for some candidate and can see if their involvement in their district is affecting some change. If it does not, they may examine other possible districts, and they can move to and be more helpful in. Each agent operates independently of others, raising the question of whether this process can converge; what are its outcomes when agents move about; and what structural issues influence this outcome.

2 Related Work

As a motivation, resource allocation lies at the core of much economic/Algorithmic Game Theory (AGT) research, and in particular, a desire to avoid a central, all-knowing, entity that divides the resources. Indeed, both auctions and pricing, the main components of the AGT toolkit, are ways in which resources can be allocated more efficiently, without requiring a centralized decision-making process [6, 13]. However, in cases like ours, where the “resources” may be autonomous, these approaches are not practical.

Discussion of gerrymandering and district division and the effects of population distribution on them have been growing significantly in the past several years in the computer-science community [2, 8]. But our particular model is more similar to hedonic games [3]: agents are choosing groups to participate in to increase their utility. However, while in hedonic games agents care for the success of their coalition, in our model they are concerned with the overall outcomes of every coalition in the game.

Somewhat related to our model, van Bevern et al. [18] modeled a setting motivated by gerrymandering, in which agents are moved according to a fixed rule in a graph in order to “reorganize” them. In a certain manner, the work of van Bevern et al. [18] is a less general case of our model.

We employ an iterative framework, which bears similarity to an extensive line of research, such as [12, 7, 4, 10, 14]. While the existence of districts causes some complications, we are less interested in the change of voters’ actual vote, but on their manipulation by changing their district which is a type of iterative manipulation not previously considered.

Calculating the optimal allocation is complicated by the “resources” in this case being the agents with the utility. Finding an optimal allocation using an iterative process (which we use) has been shown to be in most cases NP-hard [1].

3 Model

We have a set C of m candidates $\{c_1, \dots, c_m\}$, and a set V of n voters $\{v_1, \dots, v_n\}$. Each $v_i \in V$ is associated with a preference order of candidates (i.e., if $\pi(C)$ is the set of linear orders of the elements of C , each v_i is associated with its preference order, an item in $\pi(C)$). For any number of voters $1 \leq s \leq n$ we have a voting rule $f : (\pi(C))^s \rightarrow 2^C \setminus \emptyset$, and from here we have two options for tie-breaking:

deterministic Using tie-breaking rule $t : 2^C \rightarrow C$, so that the outcome of a vote by s voters is $t \circ f$.

fractional Each candidate in $f(\vec{p})$ receives $\frac{1}{|f(\vec{p})|}$ of the win (the meaning of this will become apparent shortly).

In this paper, we will focus on plurality, in which each vote gives a point to their top-ranked candidate, and the candidate with the most points is the winner. The model can easily be extended to any other voting rule.

In addition to this basic voting setting, we also have a set \mathcal{D} of k districts $\{D_1, \dots, D_k\}$. There is an initial division of voters to the different districts, which can be considered as $g^0 : V \rightarrow \mathcal{D}$. We will denote the initial voters of district D_i as $V_i^0 = \{v \in V | g^0(v) = D_i\}$. From the voters of (V_1, \dots, V_k) one can construct a district scoring vector $s \in \mathbb{R}_+^m$, giving each candidate a point (or its fraction, in the fractional case) if they have won a district. An inter-district voting rule takes this vector s and outputs a set of overall winners – $\tilde{f} : \mathbb{R}_+^m \rightarrow 2^C$. Once again, a tie-breaking rule can be used, to make the overall election winner $\tilde{t} \circ \tilde{f}$ (here we do not have a fractional variant, as we seek to end with an overall winner). In this paper (as in Bachrach et al. [2] and in Lewenberg et al. [8]), we shall use the plurality rule as our inter-district voting rule, making the winner the candidates which won the most districts.

Each voter v_i has a utility function which gets as an input the district scoring vector $u_i : (\mathbb{R}_+^m) \rightarrow \mathbb{R}$. We are going to consider two types of utility functions for our agents:

global These agents are only concerned with the ultimate outcome. That is, the final outcome of $\tilde{t} \circ \tilde{f}$ and where the winner is located in their own preference order.

lexicographic These agents consider mainly the input district scoring vector. They re-order it according to their own preference order, constructing $(score(c_1^i), \dots, score(c_m^i)) \in \mathbb{R}_+^m$ such that c_1^i is voter v_i 's top ranked candidate, c_2^i the 2nd favorite and so on. The utility function works lexicographically, such that $(x_1, x_2, \dots, x_m) \succeq_i (x'_1, x'_2, \dots, x'_m)$ if there is $1 \leq j \leq m$ such that for all $j' < j$, $x_{j'} = x'_{j'}$, and $x_j > x'_j$.

3.1 Iterative Process

At time h an agent can decide that they wish to move from their district to another. An agent v will move from D_i to D_j if $u_i(s) > u_i(s^{h-1})$ for the district score vector s^{h-1} created from $f(V_1^{h-1}), \dots, f(V_i^{h-1}), \dots, f(V_j^{h-1}), \dots, f(V_k^{h-1})$ and s created from $f(V_1^{h-1}), \dots, f(V_i^{h-1} \setminus \{v\}), \dots, f(V_j^{h-1} \cup \{v\}), \dots, f(V_k^{h-1})$.

If there are multiple changes available to v , the voter will choose the one that maximizes u_i (i.e., a best-response).

That is, an agent will move if they can improve the overall result of the election, regardless of their effect in their previous district. As in other iterative voting papers (from Meir et al. [12] onwards), we will allow only one agent to change their district at each time point.

In various settings (e.g., workplaces), there may be limitations on agents leaving an organizational sub-unit understaffed, or arriving at an already crowded sub-unit. Hence, we will explore the implications of limitations on the size of districts. We will do so by imposing upper and lower thresholds on district sizes, so that for any i , $b_- \leq |D_i| \leq b^+$. That is, if some district is at its lower threshold, no agent will be allowed to leave. If a district is at its upper threshold, no agent will be allowed to join.

Beyond the manipulation of agents moving between districts, they could combine such a move with a strategic changing of their announced vote. This applies less to resource allocation settings (a resource can not change its essence), but more to settings that involve agents pushing for some change (such as in an organizational setting). We do not allow agents to change their vote while in the same district (that would collapse to the iterative

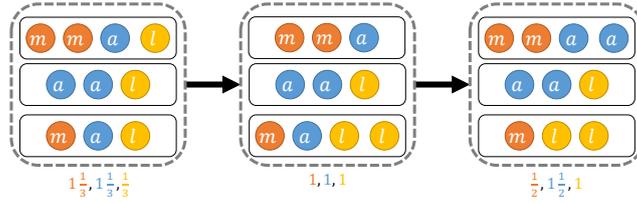


Figure 1: Example 1: The l agent in the top unit moves to the bottom unit, winning it over. This makes it possible for the a agent in the bottom unit to move to the top.

voting setting in any case), but they can change their vote when changing district. Thus, an agent $v \in V$ currently in D_i voting with a preference order $v' \in \pi(C)$ will move to district D_j with a vote v'' in time h if $u_i(s) > u_i(s^{h-1})$ for the district score vector s^{h-1} created from $f(V_1^{h-1}) \dots, f(V_i^{h-1}), \dots, f(V_j^{h-1}), \dots, f(V_k^{h-1})$ and s created from $f(V_1^{h-1}) \dots, f(V_i^{h-1} \setminus \{v'\}), \dots, f(V_j^{h-1} \cup \{v''\}), \dots, f(V_k^{h-1})$.

To recap, we have several variables in flux:

tie-breaking Either deterministic (one winner per district) or fractional. This serves to model different district settings. When, suppose, each district selects a delegate, deterministic tie-breaking makes sense. In cases of, for example, money allocation, a fractional tie-breaking makes more sense.

utility Either globalists (caring only for the global result) or lexicographic. Globalists makes sense, again, for cases of representation when, ultimately, the district process will select an overall winner. When the districts are making a monetary allocation, for example, someone fighting for a particular cause, first wants their cause to get the most it can, and only secondary to that, to allocate money to more minor concerns.

strategic actions Either just moving between districts, or also vote-strategic. This depends on the information structure – if all know a particular agent as a supporter of a particular candidate, they may not be able to switch their allegiance easily.

Naturally, in addition to that we have as parameters the environment’s variables: number of voters (n), candidates (m), districts (k), and bounds on district sizes (b_-, b^+).

Example 1. Consider a multi-unit technology company. Each unit is in charge of one product, and each product can be developed using one of the three available technologies: machine learning (ML), algorithmic game theory (AGT), and logic. There are 10 developers in the company and 3 units, and in every unit there must be either 3 or 4 developers. The developers want that as many as possible products will be developed using their favorite technology. In case of a tie in a unit, the product will be developed using hybrid technology (i.e., fractional districts).

3 developers in the company are machine learning enthusiasts (m), 3 believe in algorithmic game theory (a) and the other 4 support logic (l) (the full preferences order is not needed for this example). The initial company partition is $V_1^0 = \{m, m, a, l\}$, $V_2^0 = \{a, a, l\}$ and $V_3^0 = \{m, a, l\}$. In the initial state, the first product is developed using ML techniques, the second product is developed using AGT and third combines all three approaches. As can be seen in Figure 1, in the initial state, only a developer from the first unit can move to other units, due to the size constraints, and indeed the developer that supports “ l ” wants to move to the third unit, as this move will result in the third product be developed using logic approaches. That is, $V_1^1 = \{m, m, a\}$, $V_2^1 = \{a, a, l\}$ and $V_3^1 = \{m, a, l, l\}$. Now, the developer that supports “ a ” wants to move to the first unit and cause the first product to partly be

developed using algorithmic game theory approaches and $V_1^2 = \{m, m, a, a\}$, $V_2^2 = \{a, a, l\}$ and $V_3^2 = \{m, l, l\}$. Assuming that the machine learning developers from V_1^2 prefer $l \succ a$, then no developer has an incentive to switch units and therefore the third state is stable.

4 Convergence

The first step in our analysis is to try and see if there are cases in which the iterative process will converge.

Theorem 1. *If $b^+ - b_- = 1$ and agents are globalists and not vote-strategic and districts are not fractional, the iterative process will always converge to a stable Nash equilibrium, in which no agent wishes to change.*

Proof. First, we will show that under the theorem's conditions, the score of each district's winner is weakly monotonically increasing with the iterative process. If a district's size is b_- , no agent can leave, and if a new agent arrives, it will do so to change the winner outcome, and hence the new winner's score will be the same as the previous one (since all of its voters are still there), or strictly more. If the district size is b_+ , the last transfer of a voter to the district (which made its size b_+) was needed to make this candidate win (otherwise the voter would not have moved). This means no supporters of the current winner can leave the district, as then the score of the winner of the district will never decrease.

Assume for contradiction that the process does not converge and results in a cycle. Since each district's maximal score is weakly monotonically increasing, districts' maximal scores during the cycle always stay the same. Of all the voters that move let C' be the set of candidates they vote for, and let $c' \in C'$ be the candidate of that set that is highest ranked in the tie-breaking rule. Let v be a voter that votes for c' . Once v moves to a new district and c' wins there, there is no way for any other candidate to dislodge c' . So v has no reason to ever continue moving in the cycle, contradicting its existence. \square

Theorem 2. *If $b^+ - b_- = 1$ with lexicographic, not vote-strategic voters and districts are not fractional the iterative process will always converge to a stable Nash equilibrium, in which no agent wishes to change.*

Proof. Assume for contradiction that there are cycles and let $\mathcal{S} = S^1 \rightarrow S^2 \rightarrow \dots \rightarrow S^l \rightarrow S^1$ be a cycle, where each $S \in \mathcal{S}$ is a state that holds all the relevant information.

First, observe that in every district one voter leaves, one voter joins and so on, because $b^+ - b_- = 1$. Further, voters would join districts only if they could influence the result in their new districts. For every state $S \in \mathcal{S}$, district $D \in \mathcal{D}$ and candidate $c \in C$ let $s(c, D, S)$ be the number of votes in district D at state S for candidate c . Let $w(D, S) \in C$ be the winner in district D at state S .

Let $C' \subset C$ be the set of candidates that become victorious during the cycle \mathcal{S} , and let $z \in C'$ be the lowest candidate in C' according to the tie breaking rule. Let $D \in \mathcal{D}$ be a district in which z becomes victorious during \mathcal{S} and let $(v_1, \dots, v_{k_D}, v_1)$ be the ordered set of voters that joined district D during \mathcal{S} and let $(u_1, \dots, u_{k_D}, u_1)$ be the ordered set voters that left district D during \mathcal{S} . Assume without loss of generality, that first voter v_1 joined district D and then voter u_1 left district D . Further, assume without loss of generality that voter v_1 supports candidate z . Consider the move $S^1 \xrightarrow{v_1 \rightarrow D} S^2 \in \mathcal{S}$ when voter v_1 joined district D .

Now, there are two possible cases:

Case I: Voter u_1 does not support candidate z Since candidate z is the lowest ranked candidate, according to the tie breaking rule, who became victorious in district D

during the cycle \mathcal{S} , it must hold that $s(z, D, S^2) \geq s(c, D, S^2) + 1^2$ for every other candidate $c \in C$, $c \neq z$. As voter u_1 does not support candidate z , before v_2 joined district D , the winner was candidate z , and therefore after v_2 joined district D in state S^h there were exactly two candidates (including z) with $s(\cdot, D, S^h) \geq s(z, D, S^2)$. Because voters alternately join and leave district D and voters would join a district only if they could influence the result, it holds that in every $S \in \mathcal{S}$ there is a candidate with $s(\cdot, D, S) \geq s(z, D, S^2)$. However, if in state S^1 there was a candidate with $s(c, D, S^1) \geq s(z, D, S^2)$ then voter v_1 could not make candidate z victorious.

Case II: Voter u_1 does support candidate z When u_1 left district D they joined district D' . If the next voter to leave district D' does not support candidate z then we are back at Case I. Therefore, there is a cycle of z 's supporters that join and leave districts. Consider voter v that left district D and join district D' in the move $S \xrightarrow{D \rightarrow v \rightarrow D'} S'$. As the last voter to join district D also supports candidate z it must hold that $w(D, S) = z$, $w(D', S) \neq z$, $w(D, S') \neq z$ and $w(D', S') = z^3$. Let $w(D', S) = c$ and $w(D, S') = c'$, because v moved from D to D' it must hold $c \not\prec_v c'$, i.e., voter v prefers candidate c' over candidate c . Consider the directed graph $G_v = (C, E_v)$ where the nodes are the candidates and there is a directed edge $(c, c') \in E_v$ if and only if there is a move $S \xrightarrow{D \rightarrow v \rightarrow D'} S' \in \mathcal{S}$ such that $w(D', S) = c$ and $w(D, S') = c'$. Because \mathcal{S} is a cycle, after v joins a district they leave the district and therefore the in-degree of a node $c \in G_v$ equals to its out-degree, and thus there is a cycle in G_v . However, for every directed edge $(c, c') \in E_v$ it holds that $c \not\prec_v c'$, reaching a contradiction. □

Theorem 3. *If $b^+ - b_- = 1$ with fractional districts and lexicographic, not vote-strategic voters, the iterative process will always converge to a stable Nash equilibrium, in which no agent wishes to change.*

Proof. Assume for contradiction that there are cycles and let $\mathcal{S} = S^1 \rightarrow S^2 \rightarrow \dots \rightarrow S^l \rightarrow S^1$ be a cycle.

As before, note that in every district one voter leaves, one voter joins and so on; and voters would join districts only if they could influence the result in their new districts.

Now there are three possible cases:

Case I: This is the case if during the cycle at some state $S \in \mathcal{S}$ for some $D \in \mathcal{D}$: $|\arg \max_{c \in C} s(c, D, S)| = 1$. Let $c^* \in C$, $D \in \mathcal{D}$, $S^1 \xrightarrow{v \rightarrow D} S^2 \in \mathcal{S}$ such that $\arg \max_{c \in C} s(c, D, S^2) = \{c^*\}$ and in this case I, the next voter to leave district D does not support candidate c^* . Let voter $v \in V$ be the voter that joined D in the move $S^1 \xrightarrow{v \rightarrow D} S^2$, let voter $v' \in V$ be the next voter to join D in the move $S^h \xrightarrow{v' \rightarrow D} S^{h+1} \in \mathcal{S}$, and let voter u be the voter that left D after v has joined and before v' has joined. As voters only join districts if they can influence the results, and voter u did not support candidate c^* it holds that there are two candidates in state S^{h+1} with $s(\cdot, D, S^{h+1}) = s(c^*, D, S^2)$, and therefore for every $S \in \mathcal{S}$ there exists a candidate with $s(\cdot, D, S) = s(c^*, D, S^2)$. It is a contradiction to the fact that before v joined district D , $\max_{c \in C} s(c, D, S^1) = s(c^*, D, S^2) - 1$.

Case II: This is the case if during the cycle at some state $S \in \mathcal{S}$ for some $D \in \mathcal{D}$: $|\arg \max_{c \in C} s(c, D, S)| = 1$, however (unlike case I), this is the case if at any point

² $s(x, y, z)$ denotes the score of candidate x in district y for state z .

³ $w(x, y)$ denotes the winner in district x in state y .

that $\arg \max_{c \in C} s(c, D, S) = \{c^*\}$ for $c^* \in C$, $D \in \mathcal{D}$, $S \in \mathcal{S}$ the next voter to leave district D supports candidate c^* . In that case, let $c^* \in C$, $D \in \mathcal{D}$, $S^1 \xrightarrow{v \rightarrow D} S^2 \in \mathcal{S}$ such that $\arg \max_{c \in C} s(c, D, S^2) = \{c^*\}$ and let $u \in V$ be the next voter to leave district D and joined district D' in the move $S^h \xrightarrow{D \rightarrow u \rightarrow D'} S^{h+1} \in \mathcal{S}$. As voter u supports candidate c^* , and $\arg \max_{c \in C} s(c, D, S^h) = \{c^*\}$ it must hold that $\arg \max_{c \in C} s(c, D', S^{h+1}) = \{c^*\}$ and the next voter to leave district D' also supports candidate c^* . Hence, there is a cycle of c^* 's supporters that leave join and districts. Similarly to the proof of the previous theorem (case II), because voters leave their district only if they strictly prefer the new outcome, this dynamic cannot cycle.

Case III: This is the case if during the cycle \mathcal{S} , for every $S \in \mathcal{S}$, $D \in \mathcal{D}$: $|\arg \max_{c \in C} s(c, D, S)| > 1$. In that case, when a voter $v \in V$ that supports candidate c^* leaves district D and joins district D' in the move $S^1 \xrightarrow{D \rightarrow v \rightarrow D'} S^2 \in \mathcal{S}$ it must hold that $c^* \in \arg \max_{c \in C} s(c, D, S^1)$ and $c^* \in \arg \max_{c \in C} s(c, D', S^2)$ and in addition, due to the fractional character of districts,

$$\left| \arg \max_{c \in C} s(c, D', S^2) \right| < \left| \arg \max_{c \in C} s(c, D, S^1) \right| \quad (1)$$

also note that

$$\left| \arg \max_{c \in C} s(c, D, S^1) \right| = \left| \arg \max_{c \in C} s(c, D, S^2) \right| + 1 \quad (2)$$

and

$$\left| \arg \max_{c \in C} s(c, D', S^1) \right| = \left| \arg \max_{c \in C} s(c, D', S^2) \right| - 1. \quad (3)$$

Consider the following potential function

$$\phi(S) = \sum_{D \in \mathcal{D}} \left| \arg \max_{c \in C} s(c, D, S) \right|^2 \quad (4)$$

due to Equations 1, 2, 3 and 4 it holds that $\phi(S^1) > \phi(S^2)$, which cannot happen in a cycle, as games with potential functions are guaranteed to converge to a Nash equilibrium. \square

These convergence proofs are tight, in the sense that almost any deviation from them, results in a setting that does not converge.

Theorem 4. *The following cases are not guaranteed to converge, and contain cases with cycles:*

1. If $b^+ - b_- \geq 2$ when districts are deterministic (not fractional) and agents are not vote-strategic. [**bounds of Theorem 1 and Theorem 2**]
2. If $b^+ - b_- \geq 2$ when districts are fractional, and agents are lexicographic and not vote-strategic. [**bound of Theorem 3**]
3. If $b^+ - b_- = 1$ when districts are fractional and voters are globalist. [**bound of Theorem 1**]

Proof. For item 1, when districts are deterministic, Figure 2 shows a cycle when tie-breaking t in each district is $a \succ b$, and overall, \tilde{t} is $b \succ a$.

For item 2, when districts are fractional, Figure 3 shows a cycle, when agents are lexicographic.

For item 3, when districts are fractional and voters globalist, Figure 4 shows a cycle. \square

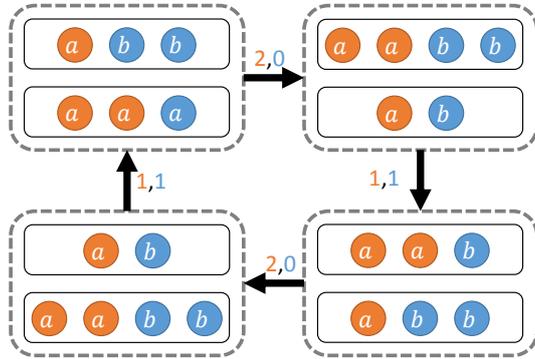


Figure 2: A cycle occurs, assuming districts are deterministic tie-breaking in each district is in a 's favor, and either agents are lexicographic, or the global tie-breaking is for b . On each edge is the score vector for the candidates induced by this move.

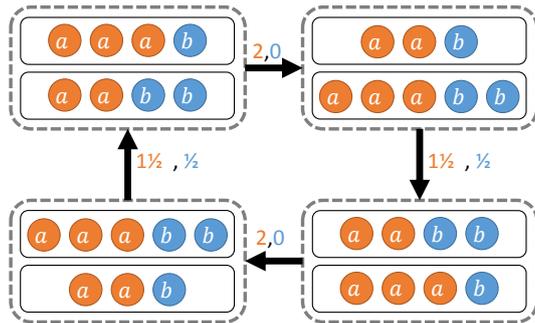


Figure 3: A cycle occurs, assuming districts are fractional and agents are lexicographic. On each edge is the score vector for the candidates induced by this move.

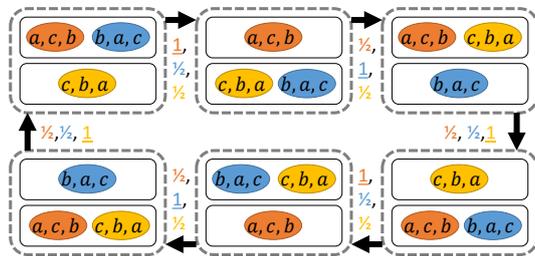


Figure 4: A cycle occurs, assuming districts are fractional and agents are globalist. On each edge is the score vector for the candidates induced by this move.

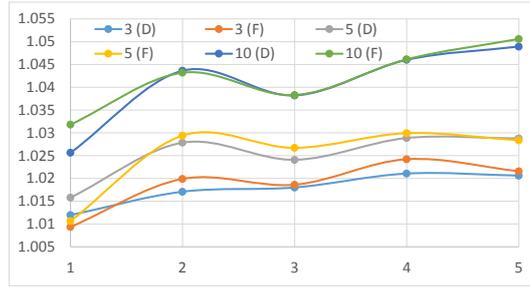


Figure 5: These are the ratio of post-game utility vs. opening position in social welfare for vote-strategic, globalist agents with single-peaked preferences, by the gap constraint (the x -axis), for elections with 5 candidates. The D marks deterministic districts vs. fractional ones (with F). Generally speaking, as the gap constraints are loosened, agents can increase their welfare.

5 Simulations

While the results on convergence are tight, we are interested to see the effects of the decentralized iterative dynamic on the overall welfare of the system. Moreover, as we have multiple parameters of the type of districts and types of agents, we are interested in exploring how (and if) are these parameters affect the efficiency of the process and how this affects the overall social welfare.

While the examples have shown a relatively small number of candidates, when trying to simulate more of them, the question arises of determining the preference order of the voters. We have chosen to run each simulation both with randomized preferences as well as with single-peaked preferences. Single-peaked preferences are relevant when there is an agreed upon ordering of candidates on some axis (e.g., political right to left; location of public parks along a street; etc.). Every voter has a particular location on the axis which is its most favored location, and the farther away an option is from that location, the lower it is in the voter’s preference order. Unlike randomized preferences, which can create quite unrealistic preference orders, there is a case to be made regarding their properties (see [16]) and resembling realistic preferences.

In order to calculate social welfare, we followed the common practice in the literature [17, 11, 5], and used the Borda scoring rule as a proxy for agent utility.

In order to simplify data analysis and comparison, we ran a few experiments with a large variety of voter numbers, but here we will present the extensive simulations we have done of the iterative dynamic with 53 voters (the number was useful due to the district size). We examined the effects of the number of districts (we ran simulations with 2,3,5 and 10), and the size of the gap between maximal and minimal group size $- b^+ - b_-$ (we ran simulations with 1,2,3,4,5).

A simulation setting included a choice of number of the district, a choice of the number of candidates, a choice of the size of gap, whether voter preferences were randomly generated or single-peaked, whether districts were fractional or deterministic, whether agents’ utilities were global or lexicographic, and whether agents were vote-strategic or not. For each of these 320 potential settings, we ran 10,000 scenarios.

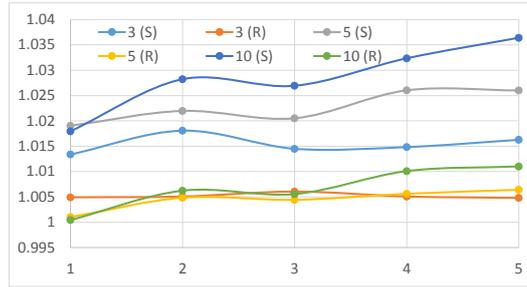


Figure 6: These are the ratio of post-game utility vs. opening position for globalist agents with deterministic districts by the gap constrain (the x axis), for elections with 5 candidates. The S marks single-peaked preferences vs. the randomized ones (with R).

6 Results

Despite the theoretical convergence results, the most striking of the simulation results was that the rate of non-convergence was so small. It was, overall, slightly less than 0.66%. While no setting had a particularly large amount of cycles, the number of districts did slightly increase them, as did using lexicographical utilities and using fractional districts. On average, 57% of initial states were already in equilibrium, but this was highly volatile, and mostly appeared in settings with a small number of districts.

Convergence happened in almost all cases within 13 steps. Single-peaked preferences took markedly less time (presumably, thanks to their more structured form), and global utility converged faster than lexicographic utility.

The flexibility that is given to players to manipulate – the maximal/minimal bounds over district sizes – had a significant effect on the outcome, though it was not necessarily monotonic. As can be seen in Figure 5 (the x -axis is $b^+ - b_-$) the effect of the gap on all cases is almost identical (we believe the dip when the gap was 3 has to do with parity issues) – even when the “amplitude” of the agents’ social welfare is different (mostly dependent on district size), the effect of changing the constraints is fairly consistent across all settings.

Examining and comparing Figures 6 and 7 allows to see that single-peaked voters are improving their social welfare more than those with random preferences. More interestingly, it shows how vote-strategic single-peaked agents are able to significantly improve their social welfare over vote-strategic voters. But random preference vote-strategic ones are not – when very constrained, the vote-strategic agents were on the borderline of not managing to show any gain in utility for all the strategizing, and as the constraints loosened, the non-vote-strategic agents managed to get better outcomes than the vote-strategic ones.

Figures 5, 6, and 7 highlight that while single-peaked voters’ utility increases as the number of districts increases, this is not a significant influence on the social welfare of random preference voters.

7 Discussion

The “reverse gerrymandering” setting, presented in this paper, may sound slightly unnatural at first blush, since people do not usually get to jump between voting districts. However, we believe that when viewed from the perspective of people participating in workplace committees, with their overlapping organizational influence, they do indeed strategize on where they could be more influential, and they move if they find a better position. In a more futuristic outlook, as autonomous systems are more common, the issue of these agents

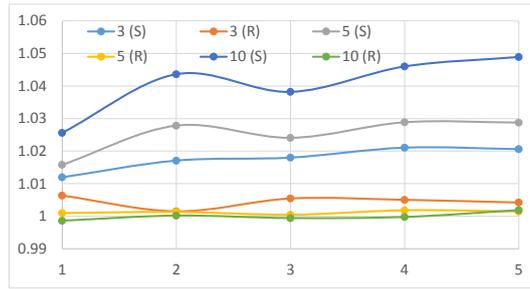


Figure 7: These are the ratio of post-game utility vs. opening position for vote-strategic globalist agents with deterministic districts by the gap constrain (the x axis), for elections with 5 candidates. The S marks single-peaked preferences vs. the randomized ones (with R).

needing to find – on their own – where they are pivotal to help, and when moving would be a waste of resources.

We presented here both theorems on the properties of this dynamic, and also explored it empirically, discovering some key issues on the effect of changing the agent preference model, and the effect of district size on the agent. There is still much to discover – what other preference models work well with this setting; understanding better the effect of the gap constraints on social welfare; and combining various different types of agents in the same simulation, examining how their differences interact.

References

- [1] Haris Aziz, Thomas Kalinowski, Toby Walsh, and Lirong Xia. Welfare of sequential allocation mechanisms for indivisible goods. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*, pages 787–794, The Hague, The Netherlands, August 2016.
- [2] 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.
- [3] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, March 2016.
- [4] 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.
- [5] 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.
- [6] Vijay Krishna. *Auction Theory*. Academic Press, 2002.
- [7] Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative scoring rules. *Journal of Artificial Intelligence Research (JAIR)*, 57:573–591, December 2016.

- [8] 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.
- [9] Alec MacGillis. Go midwest, young hipster. *New York Times*, October 23 2016.
- [10] Reshef Meir. Strong and weak acyclicity in iterative voting. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory (SAGT)*, pages 182–194, Liverpool, Great Britain, September 2016.
- [11] 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.
- [12] 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.
- [13] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [14] 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.
- [15] Everett M. Rogers. *Diffusion of Innovations*. Free Press of Glencoe, 1962.
- [16] Yves Sprumont. The division problem with single-peaked preferences: A characterization of the uniform allocation rule. *Econometrica*, 59(2):509–519, March 1991.
- [17] 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.
- [18] 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.

Omer Lev
 Department of Industrial Engineering and Management
 Ben-Gurion University of the Negev
 Beersheba, Israel
 Email: omerlev@bgu.ac.il

Yoad Lewenberg
 The Rachel and Selim Benin School of Computer Science and Engineering
 The Hebrew University of Jerusalem
 Jerusalem, Israel
 Email: yoadlew@cs.huji.ac.il