# THREE WAYS TO GET YOUR WAY: STRATEGIZE, GERRYMANDER, PARTY

by

Tyrone Strangway

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Computer Science
University of Toronto

# Abstract

Three ways to get your way:
strategize, gerrymander, party

Tyrone Strangway
Doctor of Philosophy
Graduate Department of Computer Science
University of Toronto
2021

Voting is perhaps the most basic aspect of a well functioning society. How one aggregates the preferences of the members of a group into a single decision has been a fundamental issue to human society for millennia. Today many countries use voting at all levels of their society. It can be used to pick political leaders at various levels of government. Members of admission and hiring committees use it to decide which candidates advance to subsequent rounds. Even mundane tasks like deciding what a group should have for dinner often come down to a vote. How one aggregates, or even elicits, the often conflicting preferences of the members of society is thus critical to how that society functions. Thus it should be clear who designs these preference aggregation procedures can greatly influence the direction of society. Furthermore, how the members of society decide to interact with the procedures can cause far reaching and often unintended outcomes.

In this work we will examine voting, and its impacts on society, through the view of both the members of society who have their preferences aggregated and the actors who decide how they are aggregated. What sets this work aside from political science and instead places it in the computer science community is the application of analysis and techniques commonly found in artificial intelligence works. Indeed this work, while written to be broadly approachable and readable for any one possessing a basic understanding of mathematics, is still a work that is intended for the artificial intelligence researcher.

Finally we mention this thesis is a broad look at various topics in voting, and while we go into some depth on each individual topic we stress that entire research careers could be devoted to their individual study. This thesis merely scratches the surface of each of these potential directions.

Dedicated to the memory of Gordon Harold Lewis

## Acknowledgements

Here, I'd like to thank many of the people who made getting to the end of a PhD possible.

First, my supervisors, Allan and Nisarg. You two gave me complete freedom in my research. This was, without a doubt, the reason I was able to make it to the end. Going back to my BSc, I remember one summer while working with Allan he told me publishing was not the only goal. Instead, he believed we should take the time to read and learn about a breadth of topics. I'm eternally grateful that I had the freedom to follow this advice throughout my MSc and PhD. Perhaps the only person to influenced my research as much as my supervisors was Omer Lev. The fact we're co-authors on every single project I've been involved in should speak for itself.

I also want to give special thanks to Andrew Perrault. While we never published together we spent an uncountable number of hours discussing research and other topics. Your contributions to the works in here have been invaluable.

I also should thank (and apologize) to my committee. Ashton and Sasho, sorry for you having to deal with my complete disregard for deadlines. But also, thank you for the many constructive comments you provided in our (slightly fewer than average) committee meetings.

Jérome, thank you for serving as my external reviewer, your comments and insights have greatly improved this work. Your deep knowledge of our field ensured this work is far more complete and well positioned in the body of literature.

I will also mention how important Sam[1] was to my career. In 2012, after I finished your second year data structures course, you encouraged me to apply for a summer research grant. You also encouraged me to consider applying for graduate studies in a few years. You described the MSc as "just a two year deal". I had no idea what research or graduate studies involved, but as I learned more, and have now reached the end, I'm glad you sent me down this path. It is also fitting that you served as one of my final examiners. Your careful reading and questions have made this work much more accessible.

Since I've been a member of the theory group for eight years (including my BSc) and have been at UofT since 2009, I've met a great number of people, many of whom I'm glad to call friends. There are far too many to list here, or even remember at a given time. Each of you made my time that much more enjoyable, and for that I thank you. I do want to provide a special thanks to Toni and Steve, the two of you provided invaluable mentorship towards the end of my BSc. I also want to thank Noah, for among other things, contributing to (and tolerating) our office decor.

Some of my fondest memories were spent in gyms (the Athletic Centre, Goldring Centre and various home gyms), and many of you have joined me while there. Nick, Jimmy, David, Nicolas, Jai, Chris, Sam[2], Demetres and Spencer. Without your help I'd have been crushed under a pile of iron. That probably means not finishing the PhD.

Finally, my parents, Richard and Salina. My brother, Spencer. And partner, Mona. Without your constant support and encouragement I don't know how I could have ever finished this. The greatest thanks belongs to you.

[1] Sam Toueg
[2] Samantha Murray

# Contents

# Chapter 1

# Introduction

> [It's] Dr. Evil, I didn't spend six years in evil medical school to be called mister. Thank you very much.
>
> Dr. Evil
> Austin Powers: International Man of Mystery

In this section we will first provide a brief motivation for and overview of our work. Next, we will discuss the relationship between voting and artificial intelligence. Finally, we will provide an overview of the structure of the remainder of this thesis.

We will cover a broad range of topics related to voting. At a high level voting is simply a mechanism where various groups of agents with aligned or (more often than not) conflicting preferences over various outcomes must aggregate these preferences to pick one outcome. The mechanism which aggregates the preferences and picks an outcome is called the *voting rule*. In our day to day lives we're used to seeing voting rules in the political process. These are the mechanisms that select the winners for various political positions. Broadly we can break the agents into a few groups:

1. There are the agents who submit their preferences to the mechanism for aggregation, we call these the voters. While voters don't necessarily have any input into the design of the mechanism it is their preferences which the mechanism must try to balance and satisfy.

2. There are the agents who constitute the potential outcomes of the mechanism. They are the candidates being elected. Their goal is to get elected, that is the mechanism picking them is their desired outcome. They have no more influence, with respect to their submitted preferences, than any average voter over the mechanism. That being said we will see how they can also serve a complimentary role (for themselves) with our third group of agents.

3. There are the agents who design the mechanism. A benevolent agent would design a mechanism that is good for society, one that can satisfy the various voters. Often it is impossible to satisfy all of the voters, and even a majority may not be achievable[1]. Furthermore, what constitutes a satisfied voter can be ambiguous, especially

[1] When we introduced the Condorcet paradox we will formalize the meaning of a majority being impossible to satisfy.

to a mechanism which only sees an ordinal ranking of their preferences (as all mechanisms in this work do). These agents can also be self-serving. They may design the mechanism to give a higher probability to their preferred outcome, even at a cost to society.

These groups are not necessarily rigid and certainly not mutually exclusive. Indeed, in chapter 4 we will examine how the actors who design the aggregation procedures, who are often the very agents being elected, can influence the process in order to get themselves elected even if this runs opposite to what society wants.

While our work is often inspired and entirely couched in the language of politics in subsequent sections we will discuss how our work can extend beyond this very important, but somewhat limited, setting.

## 1.1 Voting and AI

While our work is in the field of Artificial Intelligence (AI)[2] it often takes inspiration from economics, philosophy and political science. Indeed, many of the fundamental works in voting theory come from economists such as Mark Satterthwaite [100] or philosophers like Allan Gibbard [47][3]. The entire topic of gerrymandering (the focus of chapter 4), one of the most important political issue facing the Unites States electorate, began with the clever political manuevering of one man, Massachusetts governor Elbridge Gerry[4]. So the reader may ask how an AI researcher could benefit from our work and the many previous works on voting. Here we will provide an overview of the two primary reasons our work is considered AI.

The first is that many modern systems have autonomous agents that also must aggregate preferences. For example, a vehicle may have various sensors such as radar, camera and LIDAR that each report on the vehicle's surroundings. If a potential obstacle enters the vehicle's vicinity each of these sensors will report possibly conflicting inputs to a central decision making process. Some may not detect the obstacle while others may register its presence. How the system aggregates these inputs and chooses an action (stop the car, swerve or do nothing) is in essence a vote based on the sensor inputs. As these autonomous agents become more prominent in our world, understanding how they interact and come to their decisions becomes increasingly critical.

The second reason, and in our opinion more relevant to this work, is that modern techniques and analysis fundamental to AI research can be applied to problems in voting. In chapter 4 we approach gerrymandering as an optimization problem using techniques such as integer linear programming (ILP) and simulated annealing (SA) to find solutions to various problems. Theoretical computer science and AI borrow techniques from each other to do worst case analysis[5] in various voting models. We use statistical hypothesis testing in chapter 5 to compare the quality of election outcomes. These applications of modern techniques to problems in voting allow for new insights to be uncovered, advancing both the theoretical underpinnings and patterns emerging from actual voter data.

Sometimes we blend both these directions. In chapter 3 we treat voters as autonomous agents who play a game where each attempts to manipulate the preference aggregation mechanism for their own utility. Using game theory we can better understand the theoretical implications of how these autonomous agents manipulate the mechanism. Furthermore, we borrow techniques from AI to build statistical models which can help give an empirical understanding

[2]All of the components of this thesis were published in various AI conferences.
[3]Their contribution to voting research is the basis of chapter 3.
[4]The contorted and partisan districts governor Gerry drew resembled a salamander, hence gerrymandering was born.
[5]We do worst case analysis throughout the thesis.

of how these manipulations affect the welfare of society. In chapter 5 we follow a similar line of research, but now instead examine how the choices made by the designers of the aggregation mechanism affect the outcome.

## 1.2 Outline

In this section we will give an overview of the remainder of this thesis. In chapter 2 we will provide the technical background and overview of related work necessary to understand our research contributions. In particular, we will provide overviews of:

- The voting model and various voting rules.
- Non-cooperative game theory.
- Our work and the related work relevant to our research.

In chapter 3 we will focus on how strategic agents interact with various preference aggregation systems. In particular, we will look at how rational and strategic voters can manipulate the system by misrepresenting their preferences in an attempt to improve their utility. These manipulations, and subsequent response manipulations, will be formalized with the popular iterative voting model which is concerned with how these manipulations converge to equilibrium. Our contributions are mainly twofold:

- We show for a wide variety of real voting systems and agent strategies, iterative voting has no guarantees of convergence to equilibrium. The contents of this section appeared in our paper.
- We show in practice for many common distributions of preferences the strategic actions of the agents can actually lead to improved societal outcomes under various metrics.

The first section of this chapter was presented at a conference [60] and the later section a workshop [64].

In chapter 4 we will explore how voter location affects election outcomes. In particular, we focus on the problem of gerrymandering, the manipulation of the shape of voting districts in an effort to force partisan outcomes. These manipulations are in effect non-trivial manipulations in the design of the voting mechanism by partisan bodies (often the candidates or their allies). Our contributions are:

- We show for basic models some theoretical bounds on the gerrymanderer's ability to manipulate.
- We develop automated techniques that are able to gerrymander a region forcing a highly partisan outcome.
- For a variety of simple settings we use our automated method to study the potential effects of the increasing urban-rural divide amongst voters in many countries.
- Using our automated methods we show it is possible to find near-optimal (and almost always optimal) partisan solutions for various real world instances.

The first three results were published in a conference [14]. Our goal is to make the final result the basis of a future paper.

Finally, in chapter 5 we will study how political parties influence the decision making process. In particular, we focus on the rules parties use to pick their candidates who compete in general elections (known as the primary process). We will mainly examine how the design of the primary mechanism impacts the utility of the voters who interact with it. Our contributions are:

- We apply the (now) commonly used distortion measure of voter utility to a spatial model of voter preferences.
- In our spatial model we provide various theoretical bounds on the worst case difference in distortion between primary elections and their direct election analogs.
- Using simulations for various elections we examine the average difference in distortion between a primary election and its corresponding direct election.

An earlier version of these results appeared in [15]. Our goal is to turn the expanded results that appear in this thesis into a journal paper.

In chapter 6 we will summarize the main points of this work and suggest some ideas for future work. Finally in chapter 7 we will include plots, some proofs and general details that were necessary for the completion of this work but not required for understanding it.

# Chapter 2

# Background

I'm very good at integral and differential calculus;
I know the scientific names of beings
animalculous: In short, in matters vegetable,
animal, and mineral, I am the very model of a
modern Major-General.

Major-General Stanley
Gilbert & Sullivan's The Pirates of Penzance

In this section we will provide an overview of the previous work and technical background needed to read this thesis. Furthermore, we will review the related work and how our work fits into this body of literature.

## 2.1 Basic knowledge

This thesis only assumes a basic undergraduate level of mathematics, statistics and probability, one that a computer science major would encounter in a standard curriculum. That is we assume a basic understanding of:

- Asymptotic notation: big-Oh and big-Omega.
- Probability: Random variables, expected values and statistical hypothesis testing (P-values).
- Graph concepts: Definition of directed and undirected graphs, connected components of a graph and the partitioning of graph nodes.
- Graph algorithms: Algorithms for finding spanning trees (Kruskal's algorithm) and connected components (breadth first search).

## 2.2 Social choice and elections

Our work is heavily reliant on social choice theory, which studies how the preferences of various agents can be aggregated into a communal decision. This research is extensive within in the field of AI with too many contributions to enumerate (we suggest [17] is an excellent reference text for an overview of the field). Within social choice we have focused on the study of elections. While social choice theory is inherently in the domain of political science and economics its study has become an integral part of computer science, especially within the AI and theory communities, where it is known as computational social choice. Elections in particular have garnered much interest from our communities, after some technical background we will examine how elections have been studied from three different perspectives within computer science.

### 2.2.1 Technical background for voting

First we define what an election $\mathcal{E}$ is:

**Definition 1.** An election $\mathcal{E}$ is comprised of:

1. A set of $n$ voters $V = (v_1, \cdots, v_n)$.
2. A set of $m$ candidates $C = (c_1, \cdots, c_m)$. Depending on the context we may refer to the candidates as *alternatives* $A = (a_1, \cdots, a_m)$ or parties $P = (p_1, \cdots, p_m)$.
3. For each voter $v \in V$ we have a ranking over the candidates. This ranking is a strict linear ordering $\succ_v \in \pi(C)$ where $\pi(C)$ is the set of all linear orders of the candidates in $C$. This ordering represents $v$'s preferences[1] for the election outcome, that is who they prefer win the election.

Given an election instance we have a function, known as a voting rule, that picks a winning candidate. Formally, a voting rule, also known as a voting system, is a function $\widetilde{f} : (\pi(C))^n \to 2^C$. The rule takes a set of $n$ preferences (one from each voter) and selects a set of winners from the set of candidates. While our work is only concerned with single winner elections almost all voting rules can result in ties, that is the set of returned winners contains more than one candidate. In all of our work we will apply some tie-breaking procedure to ultimately pick one winner from this set.

The preferences voter $v \in V$ submits to the voting rule $\widetilde{f}$ need not match the internal preferences (I.E. true preferenes) $\succ_v \in \pi(C)$ for the voters. As we will see later on it is often possible to "manipulate" the voting rule by behaving "strategically", submitting preferences that do not match the internal ones. These internal preferences represent the true preferences of the voter, that is a ranking of the candidates in the order she prefers them. For the rest of this work unless we indicate otherwise we will assume the internal preferences match the submitted ones, that is unless otherwise stated voters are acting "honestly'.'[2]

Often we will consider a single voter changing her ranking; for ease of notation if $\succ = (\succ_1, \cdots, \succ_n)$ we say $\succ_{-i} = (\succ_1, \cdots, \succ_{i-1}, \succ_{i+1}, \cdots, \succ_n)$ is all of the rankings but that of voter $i$ and $\succ = (\succ_i, \succ_{-i})$.

There are many voting rules used in practice; here we examine some common ones we will see throughout the rest of the thesis. First we begin with what are known as scoring rules. Before we define scoring rules we need some related concepts. We say the rank voter $v$ assigns to candidate $c$ is the candidate's location within the voter's submitted ranking. That is if $c$ is the top of $v$'s submitted rankings ($\succ_v$) we say $r_v(c) = 1$ (conversely if $c$ is at the bottom we say $r_v(c) = m$). A voting rule $\widetilde{f}$ is a scoring rule if and only if it is associated with a vector $s \in \mathbb{R}^m$ and the output of $\widetilde{f}$ is $\arg\max_{c \in C} \sum_{v \in V} s_{r_v(c)}$. It should be obvious that to define a scoring rule we only need to specify the vector $s$. These are often considered the simplest rules; here are some common ones we see:

[1] We will often refer to the rankings as preferences; these terms are interchangeable.

[2] The definitions of honest (I.E. truthful), manipulation and strategic behaviour will be made clear later in this chapter.

**Plurality** Here $s = [1, 0, \cdots, 0]$. An individual voter gives exactly one point to exactly one candidate (the top of their submitted ranking). The candidate with the most points, the one who appeared at the top of the most rankings, is the winner.

**Veto** Here $s = [1, \cdots, 1, 0]$. An individual voter rejectes (or vetos) exactly one candidate, giving them zero points (giving the rest of the candidates one point each). The candidate with the most points, the one who was vetoed the fewest times, is the winner.

**Borda** Here $s = [m-1, m-2, \cdots, 1, 0]$. An individual voter gives a decreasing amount of points to each candidate as they go further down their ranking list. Their top ranked candidate receives the most points ($m-1$), the next receives $m-2$ points and so on until the bottom ranked candidate receives 0 points. The candidate with the most points, also known as the candidate with the highest Borda score, is the winner.

While there are other scoring rules, these are the ones we focus on in this thesis. Plurality is perhaps the voting rule with which we are most familiar; most elections around the world are based around a plurality vote (sometimes with a more complex system on top). Plurality and veto are also the simplest from the perspective of the voter, with plurality they need only specify who they are voting for (or with veto who they do not wish to vote for). Borda, introduced by French mathematician Jean-Charles, chevalier de Borda (1733 - 1799), is an example of a ranked ballots system, taking the entire rankings of the voters into account. While it provides more expressive powers to the voters it requires them to provide a much more complex structure (one that they may not fully know in practice).

Any voting rule which cannot be expressed as above is referred to as a non-scoring rule. Often with non-scoring rules we need to know in how many rankings one candidate appear above another, that is how many voters preferred one over the other. For any pair of candidates $c_1, c_2$ let $P(c_1, c_2) = |\{x \in V | c_1 \succ_x c_2\}|$[3]. Some common non-scoring rules we see in the literature are:

**Maximin** For each candidate $c$, define the score of $c$ as $sc(c) = \min_{c' \neq c \in C} P(c, c')$. The winner is the candidate with the maximum score $\arg\max_{c \in C} sc(c)$.

Intuitively, maximin prefers candidates who perform well in their worst pairwise matchup with other candidates.

**Copeland**$^{\alpha}$ For $\alpha \in [-1, 1]$, define the score of a candidate $c$ as:

$$sc(c) = |\{c' | P(c, c') > n/2\}| - |\{c' | P(c, c') < n/2\}| + \alpha \cdot |\{c' | P(c, c') = n/2\}|.$$

The winner is the candidate with the maximum score, $\arg\max_{c \in C} sc(c)$. In this work we always set $\alpha = 0$, so the winner is the candidate who maximizes:

$$sc(c) = |\{c' | P(c, c') > n/2\}| - |\{c' | P(c, c') < n/2\}|.$$

Like maximin the Copeland rule is concerned with how candidates perform in pairwise matchups. Unlike maximin the Copeland rule combines all of a candidate's pairwise matchups when assigning them a score. It prefers candidates who win many, and loose few, pairwise matchups with the other candidates.

**Bucklin** For each $c \in C$, d$sc(c) = \min_{k<m} |\{x \in V | \exists c_1 \neq c_2 \neq \ldots c_{m-k}$ s.t. $\forall c_j \in \{c_1, \ldots, c_{m-k}\} c \succ_x c_j\}| > n/2$. The winner is the candidate with the smallest score, $\arg\min_{c \in C} sc(c)$.

[3]Here we are letting $\succ$ represent the submitted, and not necessarily internal rankings.

Bucklin voting can be though of as a sequence of scoring-like rules, in the first round we use the vector $s = [1, 0, \cdots, 0]$, in the second round we use $s = [1, 1, , 0, \cdots, 0] \cdots$. If in any round a candidate has more than $n/2$ points we pick them as the winner, if no candidate meets the threshold we continue to the next round.

**STV** Single Transferable Vote (STV) is also a rounds based election. In the first round we have a plurality election, if no candidate secures a majority of the support (achieving a score of more than $n/2$) we move to the second round. In the second round we eliminate the candidate who received the lowest plurality score in the previous round (breaking ties in some fashion). Voters who previously had the eliminated candidate as their top choice now "vote" for their second ranked preference. We continue with this procedure until some candidate secures more than $n/2$ votes. Effectively, STV is a series of plurality elections, where candidates with low support are eliminated and voters who supported that candidate shift support to their next most preferred candidate.

**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 SOC chooses $c \in \underset{c \in C}{\arg\max} sc(c)$ s.t. $\sum_{c' : P(c,c') > n/2} sc(c')$ is maximal.

**Ranked Pairs (RP)** 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.

Intuitively, ranked pairs is likely to lock in the idea society prefers candidate $a$ over candidate $b$, if $a$'s margin of victor over $b$ is large. On the other hand if the margin of victory is small it will take longer to lock in society prefers $a$ over $b$. In this later case, when it is time to decide of society prefers $a$ to $b$ it could contradict what has been locked in. In the later case, $a$ is preferred to $b$ is not locked in.

While there is no universal definition of winner quality, a potential criteria is the notion of the Condorcet winner:

**Definition 2.** A Condorcet winner is a candidate who is ranked higher than every other candidate by more than half of the voters. That is $c$ is the Condorcet winner if for every other candidate $c'$ we have $P(c, c') > n/2$. Note, if there is a Condorcet winner then there is exactly one unique one.

The Condorcet winner, named for French mathematician Marie Jean Antoine Nicolas de Caritat, Marquis of Condorcet (1743-1794), is in a sense ideal for a majority of society. Now the question is when are they guaranteed to be chosen as the actual winner of a voting rule, this leads us to the following definition:

**Definition 3.** A voting rule is *Condorcet consistent* if whenever there is such a Condorcet winner, they are picked as the eventual winner.

Despite the high quality of the Condorcet winner, or perhaps because of it, they are not always guaranteed to exist. Condorcet showed that society may not respect the transitivity of the majority preferences, that is there is no Condorcet winner. This results is known as Condorcet's paradox:

**Example 1.** Consider three voters $v_1, v_2, v_3$ and three candidates $c_1, c_2, c_3$. The rankings of the voters are as follows $v_1 : c_1 \succ c_2 \succ c_3$, $v_2 : c_2 \succ c_3 \succ c_1$ and $v_3 : c_3 \succ c_1 \succ c_2$. Clearly the majority of society prefers $c_1$ to $c_2$ and also a majority prefer $c_2$ to $c_3$ but paradoxically a majority prefer $c_3$ to $c_1$. That is the majority of society has the ranking $c_1 \succ c_2 \succ c_3 \succ c_1$.

Amongst all the rules we've seen so far, maximin, Copeland, SOC, and ranked pairs are Condorcet consistent, while Bucklin and STV are not. We note no scoring rule can be Condorcet consistent (see [41] for a proof).

There are a few important concepts related to voting rules and preferences that we now introduce. The first concept has to do with a desirable property of voting rules:

**Definition 4.** A voting rule $\widetilde{f}$ is monotone if for any set of submitted rankings $(\succ = \succ_1, \cdots, \succ_n)$ no voter can change the output of $\widetilde{f}$ by ranking the current winner higher. That is, if $\widetilde{f}(\succ_1, \cdots, \succ_n) = a$ then there is no voter $i$ who can change their submitted ranking to $\succ_i'$, which is the same as $\succ_i$ except that $a$ has been moved higher in the ranking and now the output of the rule is $\widetilde{f}(\succ_i', \succ_{-i}) \neq a$.

Intuitively this means the higher a voter ranks a candidate the better the candidate's chance to win the election. While this may seem like a natural property all voting rules should possess many we use in practice do not. STV, which is used in Ireland and Australia, is not monotone. On the other hand it should be obvious that all scoring rules (where the scoring vectors are arranged in descending order) are monotone.

The second concept is not as formal, more of a principal or observed phenomena, It has to do with the distribution of power in plurality voting systems, also known as first-past-the-post systems. It is due to French sociologist Maurice Duverger (1917-2014) and is known as Duverger's law:

**Definition 5.** Duverger's law states that in first-past-the-post elections (with single representative districts[4]) almost all votes will concentrate about two political parties (that is their candidates).

While there are exceptions to this law, Duverger himself did not claim it must always hold. The United States is perhaps the most prominent example of this, where in most elections almost all the votes go to the Democratic or Republican party. Some notable exceptions include the United Kingdom which includes very strong regional support for particular parties (such as the Scottish Nationalist Party) and Canada which since the year 2000 has had three parties (the Conservatives, NDP and Liberals) hold official opposition status (see [30] for an examination of Duverger's law in various countries).

The final concept we introduce relates to a special class of voter rankings known as single-peaked rankings. So far we've described rankings as totally arbitrary, any voter can have any ordering over candidates, singles-peaked rankings introduce a restriction on voter rankings:

**Definition 6.** A set of rankings $(\succ_1, \cdots, \succ_n)$ is single-peaked if there is:

- A ordering of the candidates of $C = (c_1, \cdots, c_m)$ (all voters agree with this ordering).
- For every voter $v$ if $c_v^*$ is at the top of the ranking $\succ_v$ then $\forall x_i, x_j \in C$:

  $x_i < x_j \leq c_v^* \implies x_j \succ_v x_i$

  and

  $x_i > x_j \geq c_v^* \implies x_j \succ_v x_i$.

[4]In most systems, and those we examine in this work, districts elect a single representative. However, there are some exceptions such as the Republic of Ireland which uses districts with multiple representatives.

Intuitively, for a voter $v$ let $c^*_v \in C$ be the top of that voter's preference ranking. Their preferences would be single-peaked if for any two alternatives to the left of $c^*_v$ in the ordering of $C$ or any two alternatives to the right of $c^*_v$ in the ordering of $C$, the voter's ranking of these two alternatives is determined by their proximity to $c^*_v$ (closer to $c^*_v$ is more preferred). Single-peaked preferences are particularly attractive since they guarantee many useful properties. Perhaps the most useful property is they ensure that there is a Condorcet winner in $C$ if there is an odd number of voters[5]. Unfortunately single-peaked preferences are overly restrictive for modelling real preferences. That is, it is often the case that real preferences violate this property. See [70] for examples of non single-peaked preferences gathered from people.

## 2.3 Strategic voting

In this section we introduce the concepts needed for chapter 3. Specifically, we consider the strategic considerations voters may undertake when deciding how to vote. In all other sections we treat voters as truthful, they report their true preferences to the voting mechanism. In this section we will allow voters to behave strategically and report preferences other than their actual ones. So we can better motivate the idea of strategic voting let us consider this informal example:

**Example 2.** Consider an election where there are 100 voters. Let us assume the election is using the plurality rule and 40 voters have the ranking $a \succ b \succ c$, 45 voters have the ranking $b \succ a \succ c$ and 15 voters have the ranking $c \succ a \succ b$. Assuming these are the rankings that the voters submit we would end up with $b$ as the winner. It would be reasonable to conclude that given the different levels of support for the candidates the 15 voters who voted for $c$ are wasting their vote. They are supporting a candidate with little overall support, a far cry from the 40 and 45 votes that $a$ and $b$ respectively garner. Furthermore each of these 15 voters would prefer that $a$ wins instead of $b$, if they were to abandon $c$ and instead switch to $a$ they could change the winner and end up with a preferred outcome.

We stress that this is just an informal example, we have not specified the process, or processes, that voters go through when making strategic considerations. Defining strategic (also known as tactical) voting can be tricky, later in this section we will look at some models that have been proposed in the literature. Before we continue we consider this quote from Fisher [43]:

> A tactical voter is someone who votes for a party they believe is more likely to win than their preferred party, to best influence who wins in the constituency.

While Fisher's quote only applies directly to plurality (it is not hard to stretch the definition to any scoring or even monotone rule), it highlights one important consideration. The strategic voter acts on some belief they have about the state of the world, how we define this belief is perhaps the most important consideration when defining models of strategic voting. In the above example do we assume voters have complete information, do they know prior to the elections the voting intentions of every other voter. Alternatively, they may only have an estimate of their intentions with some levels of uncertainty. Beyond the available information how do voters evaluate the actions of other voters. If a voter acts strategically do they assume other voters would in turn respond with their own strategizing? That is, do voters behave in a myopic or non-myopic way?

[5] If $|V|$ is even then there need not be a Condorcet winner.

### 2.3.1 Equilibrium and solution concepts

In this section we introduce the game theory concept of equilibrium as a notion of stability amongst voters. For this we need some definitions from game theory. A game $\mathcal{G}$ is defined as follows:

**Definition 7.** A game is comprised of:

1. A set of $n$ agents, where agent $i$ has a finite set of actions $S_i$. Each agent $i$ will play a strategy $\sigma_i$ which is just a probability distribution over $S_i$.

2. A set of finite outcomes $O$, each agent $i$ has a utility for the outcomes $o \in O$, expressed as $u_i : O \to \mathbb{R}$.

3. A function which maps the actions to outcomes $f : S_1 \times \cdots \times S_n \to O$.

If each agent $i$ commits to some strategy $\sigma_i$ we induce a probability distribution over the outcomes $O$, with this distribution we can measure the welfare of all agents as their expected utility.

The reader should notice this is effectively a generalization of the language of elections, the actions of the agents are just the rankings they present[6], the function $f$ is the voting rule, the outcome is just the election winner and utility functions are a cardinal version of the true rankings. We note the concept of games are applicable to a far larger body of work than just elections (for an excellent background text we refer readers to [82]).

Often we restrict the strategy $\sigma_i$ to be deterministic, that is it has the form $\sigma_i = (0, \cdots 0, 1, 0, \cdots, 0)$. We call such a strategy a *pure strategy*. A *mixed strategy* is any probability distribution over pure strategies (pure strategies are a subset of mixed strategies). In the realm of voting we mainly limit ourselves to pure strategies.

Central to game theory is the notion of stability. Stability just means no agent $i$ wishes to change their current strategy $\sigma_i$. That is given the strategy (or their belief of the strategy) of the other agents they can not increase their expected utility by deviating from their current strategy.

Let us return to our above example with 100 voters. Let us examine the strategy where each agent presents their true ranking, that is each voter honestly reports the top of their ranking to the mechanism. If the 15 type $c$ voters were all controlled by one agent[7] then the outcome of everyone reporting their true rankings would not be stable, they would see how close the vote is and instead report $a$ as their top choice[8], thus increasing their utility. This new outcome is stable. The voters who have $a$ at the top of their true rankings are getting the best outcome, they have no reason to switch. The voters who have $b$ at the top of their true rankings can switch and vote for $a$ or $c$, but that would not change the outcome from $a$ winning. And the voters with $c$ at the top of their true rankings could switch back $c$ or switch to $b$, either of which causes $b$ to win which hurts their utility.

Nash provided and proved the existence of what is perhaps the most well known and widely used stability notion, the Nash equilibrium ( [80]):

**Definition 8.** A set of strategies $(\sigma_1, \cdots, \sigma_n)$ is a Nash equilibrium if there does not exist any agent $i$ where the expected utility of playing $\sigma_i' \neq \sigma_i$ is higher than playing $\sigma_i$. That is for every agent $i$ the expected utility of $(\sigma_1, \cdots, \sigma_i, \cdots \sigma_n)$ is as least as high as the expected utility of $(\sigma_1, \cdots, \sigma_i', \cdots \sigma_n)$ for any strategy $\sigma_i' \neq \sigma_i$.

[6] Hence for each agent the action space should contain $m!$ options.

[7] We assume all 15 of the voters are controlled by one agent. If each voter is an individual agent it makes the reasoning about switching votes more difficult. Namely each must weigh the odds of other type $c$ voters switching.

[8] We must stress in this example we assume agents have complete information, the 15 voters for $c$ know the exact vote breakdown. If they instead had no knowledge about the vote breakdown they would not have a compelling reason to change votes.

Essentially a Nash equilibrium is a state of the world where if every agent had complete information no individual would want to unilaterally move. It is a compelling solution concept for many games, and is widely used across many fields. It is useful for capturing the behaviour of rational agents who do not collude together. It is especially attractive since it is guaranteed to exist. Nash showed every normal form game (a finite number of agents and actions) has a Nash equilibrium. It is important to note Nash requires that agents are allowed to play mixed strategies. If we only allow pure strategies the game may not admit a Nash equilibrium (known as a pure strategy Nash equilibrium). Often to understand the properties of a system we ask what are the properties of its Nash equilibria.

In voting we don't really encounter the issue of no pure Nash equilibria. In fact we have quite the opposite problem of far too many pure equilibria, many of which we would consider unrealistic and unsuitable for analysis. Let's return to our informal example from before with 100 voters and 3 candidates. To be more formal we say the 40 voters with $a$ as the top rank are controlled by one agent, the 45 with $b$ as the top rank are controlled by another agent and finally the 15 with $c$ at the top are controlled by a third agent. Furthermore, let us add a fourth candidate $d$, we will place $d$ at the bottom of every voter's ranking, that is $d$ is the least preferred candidate of all 100 voters (or all 3 agents). It should be obvious that all 3 agents (or all 100 voters) reporting $d$ as their preferred choice leads to $d$ winning and this constitutes a (pure) Nash equilibrium. Clearly this is an unreasonable outcome, under no reasonable situation would rational voters decide to vote unanimously for a universally despised candidate.

There are a few ways out of this situation, we could weaken the unilateral requirement in the definition of Nash equilibrium, that is we allow groups of agents to move together. Clearly any pair of agents constitute a majority of the vote, so any two moving could change the winner. This does raise the issue of how the agents coordinate their actions. Furthermore by weakening the definition in this way we lose the guarantee of existence (even for cases where we permit mixed strategies). Instead many approaches we will examine use different notions of stability, all of the approaches vary, but their goal is always to try to capture how voters behave naturally so that we can have and study reasonable stability points. Many of the approaches we will see are built on top of the concept of Nash equilibriums, in fact our approach is essentially a refinement of the set of Nash equilibriums to what can be considered reasonable ones.

### 2.3.2 Only dictatorships can't be beat

At this point the reader may wonder why not just find a voting rule that is immune to the manipulations of individual voters. After all the if we were able to find a way so it's always in the best interests of voters to report their true preferences then the Nash equilibrium would be a perfect basis for analysis. That is, why not construct some voting rule which has the desirable property that tactical manipulations by a voter cannot improve that voter's utility? Here is one potential rule that is immune to tactical manipulations:

**Definition 9.** The dictatorship rule is as follows. Given the preferences of voters select the top ranked candidate of $\succ_i$ for some fixed $i$ as the winner.

Clearly this rule is immune to manipulations, $i$ should submit their true ranking so that their top choice is the winner, and $j \neq i$ can submit any ranking, including the truthful ranking, since they have no impact on the mechanism. Ignoring the immunity to strategizing, this rule is clearly terrible. However, it turns out it is the only rule that is immune to strategizing. This was formalized by Gibbard and Satterthwaite [47, 100]. Their result, known as the Gibbard-Satterthwaite theorem, states:

**Theorem 1.** *If a voting rule is non-manipulable*[9] *and has at least three possible outputs*[10] *it must be the dictatorship rule.*

This result means there is no non-dictatorial rules[11] where we are guaranteed voters are best off reporting their honest preferences.

This impossibility result has spawned various directions of research in the AI community. There has been a large focus on the complexity of manipulations, that is finding an improvement as a voter when faced with a list of other rankings and a voting rule. For many rules it is NP-Hard to find a manipulation, and the best known strategy involves enumerating all $m!$ rankings over $C$. There has also been some research on the complexity of manipulations in games of incomplete information. For more information on the complexity of manipulation we refer the reader to [17].

### 2.3.3 Iterative voting and dynamics

The model we work with belongs to a broad family of models we refer to as iterative models. At a high level an iterative model begins with some initial strategy, often the true rankings being submitted, and in an iterative fashion voters are allowed to change their vote to change the outcome to one they prefer. When no voter feels they can make an improvement step the game has reached a stable state. What level of information a voter has about the actions of other voters will impact the decisions made at each step[12]. We will first introduce the model we work with and then examine some other alternatives that have been studied.

The particular model we work in is known as iterative voting, originally introduced by [74]. Iterative voting is a game of complete information. At every phase all voters know the submitted rankings of every other voter, and hence they can calculate the current winner. If any voter can adjust their submitted ranking and elect a better winner the mechanism will pick one of these voters (possibly at random) and they will submit this new ranking. If a voter has multiple preferred candidates she can make then she will opt to make her most preferred of this set the winner, this is known as the *best response dynamic*. When no individual voter has a manipulation that leads to an improvement, *known as a better response*, then the game has converged to a stable state, that is it has reached a Nash equilibrium.

One issue with iterative voting as we've described it is the best response dynamic voters employ when searching for a manipulation. While it is a natural dynamic it does have issues, primary among them is its complexity to voters. For rules like plurality and veto where the voter just needs to indicate who they wish to vote for (or who they do not wish to vote for), best response is easy to calculate, it only requires looking at $m$ different rankings. However, for rules which make use of the full ranked ballot the calculation can be much more complex, and in many cases it is NP-Hard to find if a better response exists. For rules such as STV to find a better response the best algorithm we know of has voters consider all $m!$ rankings. Even for moderately sized candidate pools this quickly becomes intractable. Because of this we also study heuristics a voter may apply when seeking a better response. The heuristics we introduce result in voters using, what we consider, natural strategies that only consider a subset of potential manipulations.

[9] Voters reporting their true rankings is always a pure Nash equilibrium.
[10] This means the rule can always prevent all but three predetermined candidates from winning.
[11] Even junta based rules which ignore the ranking of all but a fixed subset of voters are covered here.
[12] How they act upon this information is also important. Are they cautious trying to prevent a bad outcome, or are they optimistic seeking a high utility outcome?

### 2.3.4 Previous work on iterative voting

Iterative voting is a well studied model where in most of the previous research authors often examine two questions. The first is how the process proceeds to stability, does it eventually converge to stable state, or will it get stuck in a cycle of the same strategies forever? If it does terminate how quickly does it do so? If it does not always terminate how often can we expect to encounter these cycles? The second question involves the quality of the Nash equilibrium iterative voting reaches. Are the winners in these iterative outcomes an improvement over the initial winner? Are certain candidates more likely to emerge as the winner after iterative voting? Furthermore how focused are these outcomes? That is, do we tend to have a large set of eventual winners or is it only a small group of winners that emerge?

There has been much work on the convergence of iterative voting that investigates if the process always ends in some stable state. The first such work [74] showed plurality elections will always converge if voters pursue best-response and ties in the voting rule are broken according to some linear order. In parallel, [62] and [96] also showed veto converges under best-response and linear-order tie breaking. The remaining question about the convergence of scoring rules were settled by [63], the authors showed that under best-response no other scoring rule will converge. They also showed that maximin will not converge with non-linear order tie breaking.

Some works have shown guaranteed convergence for various voting rules by restricting the best response dynamic. In [48] the authors introduced two dynamics where a voter only considers the current winner, and not the rankings of other voters. They showed under these restricted dynamics all scoring rules, Copeland and maximin will converge. The authors of [67] studied a very restrictive dynamic that also allows Copeland and maximin to converge. In [87] the authors identified various conditions on the properties of dynamics that are sufficient for convergence. Finally, in [92] the authors studied the complexity of deciding if a state is reachable under the iterative voting dynamic. They showed in general it is NP-Hard, but tractable if the voters have a truth-bias or lazy-bias[13].

Finally, [93] varied the model by limiting how much information voters have access to, and showed scoring rules converge with a very restrictive dynamic called k-pragmatism[14].

Many of these works also focus on the quality of iterative voting outcomes [73, 93, 48, 107]. By making use of simulations they show iterative voting often leads to better outcomes when compared to all voters reporting truthfully[15]. For a recent survey on iterative voting we recommend [72].

#### Iterative voting as a solution concept

As a solution concept the set of iterative voting equilibria serve to refine the set of all Nash equilibria. In a sense the iterative voting equilibria represent a "natural" set of pure Nash equilibria, those that myopic agents who behave strategically would arrive at under their own power. If we start from the true rankings the iterative process would avoid the degenerate Nash equilibrium we pointed out before, all voters voting for a unanimously despised candidate. Instead we will end with a candidate who was chosen by some subset of the voters since this new winner is preferred over who would have previously won.

This may not seem like the most realistic model when it comes to actual elections. It assumes voters are able to perfectly discern the decisions made by all other voters and evaluate how a potential change in strategy on their own

[13] A truth-bias voter is a voter who will report her truthful ranking at all times, unless her ranking is pivotal in changing the winner to one she prefers. A lazy-bias voter is the same, except she will abstain from the election unless her vote is pivotal.

[14] Here the voter will move her most preferred candidate amongst the top $k$ front runners to the top of her ballot.

[15] How we define better is an issue since we don't have access to cardinal utility functions.

part would move the outcome. In some settings this is possible, if a vote was taking place with board members in a conference room it would be easy to quickly garner all relevant information. But in an election with tens of millions of participants no voter would possess this power. Instead in these situations we think of iterative voting as reacting to the state of the world, a voter may see information about voting intentions, say through a poll or a poll aggregator, and can conclude who would win the election based on this (for further research on reacting to types of voter information see [94]). If the poll is perfectly accurate, and all voters agree that it is accurate, we have returned to iterative voting. Later on we will see models which can introduce uncertainty about the poll accuracy. Another consideration we must factor in is the impact of a single agent or voter. With any voting rule if we have millions (or even hundreds) of voters the impact a single voter has is almost always meaningless, almost every state is a pure Nash equilibrium. Instead, to model large elections we will have a system with a small number of agents each representing the intentions of a group of like-minded voters having the same preferences and behaving strategically together. This situation is not as unrealistic as one may expect, in Australia which uses STV, voters can indicate the party they prefer act as their agent, the party handles the strategizing for the voter, submitting the ranking as a proxy[16]. Somewhat related to our idea is the notion of proxy voting with autonomous agents [26, 55]. Here voters give some agent the ability to vote in their place, offloading the responsibility of voting and strategizing onto a trusted authority.

### 2.3.5 Our work

Within iterative voting we study two distinct problems. First in section 3.2 we aim to understand the impact that dynamics beyond best response can have on iterative voting. In particular, we focus on the case of non-scoring rules, an area that has been sparingly studied thus far. We note that some non-convergence results have been shown for best-response and some convergence results but only for very restrictive dynamics. We will mainly focus on what we consider less restrictive dynamics, still not as wide as best response, but not as restrictive as previous ones. These dynamics still allow voters to consider the actions of other voters unlike some of the restrictive dynamics we've seen before, they just limit how many of the $m! - 1$ manipulations a voter considers. While these dynamics serve to reduce the computational load on voters it is also interesting to study in its own right. We focus on natural strategies a reasonable voter may employ trying to understand their impact on the quality of the iterative voting outcomes.

Second in section 3.3 we focus on iterative voting in district-based elections. We will also examine district-based elections in chapter 4, but there in the context of the map makers who are the agents manipulating. Here we will consider the implications of the voters manipulating. From the perspective of voter utility, and hence their strategic actions, there is an interesting subtlety introduced by a district system. Voters within a district directly elect a district representative to their legislative body. In parliamentary systems if a majority of the district representatives can agree on a leader this person goes on to become the overall government leader. Now a voter's utility is a function of both who their representative is, that is the person who represents their needs to the government, and who the overall government leader is.

### 2.3.6 Iterative voting with uncertainty

A particularly interesting extension of iterative voting was introduced in [73] and extended in [71]. These papers introduced the notion of uncertainty for plurality elections. Returning to our repeated polls example, our model of iterative voting assumes the poll is perfectly accurate and every agent believes it is accurate. With a radius of

[16] The government body who handles Australian elections, the Australian Election Commission, refers to this as *Above the Line Voting* [1].

uncertainty each voter may believe the poll is somewhat accurate, but there is a viable range of outcomes around what it predicts. Formally we say a voter with radius of uncertainty $r$ believes if a poll predicts candidate $a$ will receive $x$ votes, the voter believes any outcome from $a$ receiving $(1-r)\cdot x$ votes to $(1+r)\cdot x$ is possible[17]. Now with a radius of uncertainty there is a multitude of possible true states of the world a voter believes is viable based on the poll. It is important to note this is not a probability-based model, the voters assume each outcome is equally likely. If we set the radius to be 0 then we arrive back at iterative voting, since the only state of the world the voter believes is viable is the one the poll reports.

Since each voter believes there is now a set $S$ of potential states of the world they can no longer pursue the best-response update, what is a best-response in one state could be a detrimental action in another. A common technique to deal with uncertainty, and one used by the authors, is the notion of non-dominated actions. In the plurality voting setting, we say for a voter $v$, candidate $a$ dominates candidate $b$ with respect to a set of viable states of the world $S$ if there is no state in $S$ where if $v$ votes for $a$ the outcome is strictly worse than if $v$ voted for $b$. Furthermore, there is at least one state of the world $s \in S$ where the outcome is strictly better for $v$ if they vote for $a$ instead of voting for $b$. That is voting for $a$ is always at least as good and in at least one state better than voting for $b$ with respect to the viable states $S$. Instead of best-response a voter will opt to stop voting for candidate $b$ if they are dominated by some other candidate $a$. Now the voter will instead vote for some candidate $c$ who is not dominated by any other candidate[18]. This is still an iterative model, voters move from dominated candidates to un-dominated candidates, and stability means no voter is voting for a dominated candidate[19]. The authors show that this process must converge to a stable point. Further experimental work shows that it often converges to the case where two candidates receive a majority of the vote, just as Duverger's law postulates.

The reader may wonder why we do not work with this particular uncertainty model, after all it seems to have the benefits of the iterative voting model while removing the restriction of voters having complete information. The reason we work with iterative voting is twofold: Firstly, this uncertainty model is a strict generalization of iterative voting, we feel the first steps must be taken with the most basic system, before more advanced considerations can be understood. Secondly, and perhaps more constricting, is the fact this uncertainty model only makes sense for plurality voting and not even any other scoring rules. Our work focuses on non-scoring rules. Recall the uncertainty model says the polls report plurality votes received by a candidate and the voters form a radius of uncertainty around the number of reported vote totals. What exactly is reported in a poll with a non-scoring rule[20]? Furthermore, it is not clear what uncertainty about a non-scoring rule outcome means. With plurality we have a one-to-one map of voting intentions to reported scores, it is easy to say what a small change (or miss-calculation) in submitted rankings does to the outcome. With ranked ballots, and non-scoring rules, what a change to a small number of rankings means for the outcome is no longer obvious.

### 2.3.7 Non-iterative voting models

There have been many non-iterative solutions proposed for strategic voting, here we review some of the more prominent ones. Several models have been proposed where voters use probability to measure the likelihood of outcomes.

[17]The authors also examine other radius of uncertainty measures such as additive ones and the earth-mover distance.

[18]In the case where there are multiple un-dominated candidates the voter will opt to vote for her most preferred in this group.

[19]Note domination is a transitive relationship so if a voter is currently voting for a dominated candidate, there must be some un-dominated candidate they can vote for.

[20]In Australia which uses STV, polls report how many people intend to rank a particular candidate at the top of the ballot, since the full ballot is critical to the outcome this is a rather imprecise tool.

Perhaps the most famous of these is the Myerson and Weber model [77]. Similar to the uncertainty model proposed by [73] this model begins with a set of reported scores for the candidates[21] and voters having some level of uncertainty about the accuracy of these scores. Unlike the work of [73] this uncertainty is associated with a probability. Now a voting profile is stable if each voter is presenting a ranking which will maximize her expected utility with respect to her estimate on the likelihood of the states. While the authors show a stable setting of rankings must exist for a wide variety of rules it is not "easy" to determine when voters are in one. To determine if they are playing a stable strategy voters, must express a cardinal value associated with each outcome, not just an ordinal one, and engage in a non-trivial probabilistic calculation. This is a problem shared with all of the probability based models, they assume something more complex than ordinal rankings and some level of mathematical maturity.

There have also been several works that handle uncertainty in a similar way to [73], where voters do not associate states of the world with probabilities. In [39] an extreme view is taken where a voter considers all states of the world possible. In general we find this type of model too broad, a voter has zero information about the actions of other voters. On the other hand, [27] is closer to the model of [73], here a single manipulating voter must decide if they have a non-dominated action when faced with partial information about the state of the world[22].

### 2.3.8 Strategizing by swapping

Finally, we discuss a recent concept related to strategic voting, vote swapping. In vote swapping voters enter into a contract with other voters where they agree to act strategically on behalf of each other[23]. For a more complete discussion we point the reader towards [51].

A simple example of where vote swapping has potential is in any parliamentary system where more than two parties are competitive. As an example, in Canada the 2019 federal election had 338 districts each elect a member of parliament[24]. These members belong to four official parties, one unofficial party, and one independent member[25]. The more members of parliament a party has the more power they have in government. Because Canada uses the plurality rule within each district, it is often the case voters of non competitive parties are wasting their vote. A voter concerned with their vote being wasted could use swapping to make it have more impact. This would require the voter (call them $a$) to find another voter (call them $b$) in another district who is also concerned with their vote being wasted. Furthermore, each voter would have to be fine with voting as the other voter wants in their own district. That is, $a$ should be fine with $b$'s preferred party representing their district (and $b$ should be fine with $a$'s preferred party representing their district).

While beyond the scope of this thesis it is possible to make this informal example more complete. We would need to extend game outcomes to include concepts like how governments are formed, and overall support for each party.

[21] This model only works for scoring rules.

[22] Unlike iterative models the authors were concerned with the complexity of finding a manipulation and not the process of subsequent manipulations.

[23] Since voting is often a private process the contracts tend to be unenforceable.

[24] Each member of parliament represents one of Canada's 338 districts. We will examine district based elections in chapter 4 and section 3.3. For now it is sufficient to know members of parliament represent a district and are elected by the voters who live in that district.

[25] To achieve official party status a party must have twelve members of parliament elected. The Liberal Party, Conservative Party, Bloc Québécois, and New Democratic Party all achieved official party status. The Green Party did not meet the required threshold. Finally, one member of parliament was elected without any party affiliation.

### 2.3.9 Why iterative voting

We stress again that iterative voting, the model we choose to work in, makes minimal assumptions on the voter's power; we only require them to ask if they prefer an outcome where one candidate wins or one where another candidate wins. Furthermore, as iterative voting allows voters to survey the state of the world, they avoid having clearly unreasonable situations bias their decisions.

## 2.4 Where voters live

In this section we will cover the specific background required for chapter 4. In particular we focus on where voters physically live and vote. While elections are often reduced to the most basic numbers, the number of votes each candidate garnered, a more subtle effect of where people live is at play. In order to understand this we need to examine the role representatives play in elections.

### 2.4.1 The role of representatives and gerrymandering

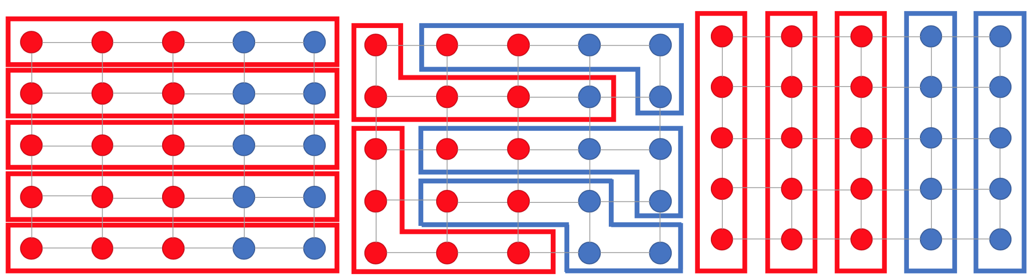

Figure 2.1: Consider this informal election. Each node represents a voter, their colour indicates which of the two parties they intend to vote for (red or blue). In this election we must divide the voters into five equal size connected components, the party which has the most voters in a region will win that region. The goal of the parties is to win as many regions as possible. There are many valid ways of dividing the voters into regions. The first division favours the red party, they win every district with a slight majority of the vote. The second division favours the blue party, they win a majority of the districts, despite having a minority of the vote. The third division is the most proportional possible, the number of districts a party wins is proportional to their vote share.

In Westminster systems (e.g. Canada, the United Kingdom and Australia) voters do not directly vote for a prime minister (the overall leader of the country), instead they vote for a local representative. Voters are partitioned into distinct geographic regions, known as electoral districts or ridings. Each of these regions holds an election to pick a representative (in Westminster systems they are known as members of parliament). The main job of this representative is to advocate for the interests of the residents of their district at a national level. For a myriad of reasons potential representatives often associate with a national political party, perhaps the most important reasons being the financial and logistical support of the party and associating themselves with a unified message.

To form a government a potential prime minister must have the support of a majority of the elected representatives. Also, the more representatives a party has the more power they have in passing legislation into law (a majority or sometimes super-majority of representatives is required to enact law). The goal of a political party[26] is not to garner

[26]Despite the fact potential representatives are individuals and not necessarily synonymous with their party for ease of notation, and consistency with other sections, we will treat them as the same entity referring to both the party and its representative as the candidate.

the most votes, but to have the most representatives (equivalently their goal is to win the most district elections).

Since the goal is to win as many districts as possible the partitioning of voters into these districts is incredibly important. As Figure 2.1 shows there are three valid but vastly different outcomes. It should be clear that who draws the districts and what constraints they must factor in are highly determinative of the election outcome. In the United States the state level governments are often left in charge of drawing the districts every ten years[27]. Usually their only constraint is each district must have an almost equal number of residents (say a $1\%$ difference between the largest and smallest population[28]) and be geographically contiguous. These somewhat lax requirements combined with the partisan goals of governments (who wish to stay in power and advance their own agenda) have led to the drawing of *unnatural* districts. This process is known as gerrymandering – partitioning voters into districts, possibly oddly shaped, for partisan gain.

We should note that the near equal population requirement is somewhat less strict in other countries. In Canada, redistricting is controlled by the independent government agency Elections Canada. Here, population bounds are allowed to deviate by up to $25\%$. Even larger deviations are permitted if required to keep communities of interest together or keep the geographic size of a district manageable [21].

Because of the geographic and population constraints on voting districts, the distribution of where people live, or more precisely where voters of specific parties live, is clearly tied to gerrymandering. We will return to this point later, but first we examine some of the historic, and more recent, work done on gerrymandering.

### 2.4.2 The social-science view on gerrymandering

Much of the work in the political and social scientist concerns gerrymandering in the United States. In particular, since the Voting Rights Act of 1965 the role of minority voters has been of particular interest[29]. Because of more sophisticated techniques and willingness to gerrymander in recent years there has been a focus on the negative impact of gerrymandering [59, 81]. The main criticism of gerrymandering tends to be the election outcomes do not represent the will of the people, and instead only advance partisan interests, these outcomes are artificial and hurt public welfare [86].

One fundamental area of research is understanding what constitutes a "good" or "fair" districting scheme (see [108] for a summary of potential measures). One reasonable goal could be district compactness, do we want districts that have no odd protrusions and resemble convex shapes? What about population homogeneity, do we want voters within a district to have a shared community such as a racial or economic one? Understanding how to balance these sometimes at-odds goals can be quite tricky, in general it seems that there may be no one-size-fits-all solution [109].

### 2.4.3 The technical view on gerrymandering

Prior to recent interest in gerrymandering from the computer science community there was a large amount of research on a related graph problem. Finding a division of a map into electoral districts is effectively a planar graph problem.

[27] The reason for redrawing districts every ten years is to best match the latest census data and ensure districts represent the will of local communities.

[28] What is exactly meant by "near equal" is itself ambiguous. As [65] points out, anything larger than a $1\%$ deviation in population between districts would normally be considered illegal.

[29] In the United States the only clearly illegal districting, besides those which violate population and geographic constraints, are those which are drawn to disenfranchise minority voters. How one determines what is drawn to disenfranchise them is a matter of intense debate.

Here, precincts[30] are nodes in the graph and edges represent physical adjacency. To partition the voters into electoral districts is the same as partitioning the nodes of the graph into different components. There are restrictions on the districts, often they must be physically connected which is to say our graph partition must form connected components. Furthermore, there is often a population restriction, each district must be nearly equal in population; that is, each connected component must have a near equal number of nodes (or node weight for weighted graphs). In [33] the authors conjectured that finding a nearly equal weight partition of a planar graph into two or more connected components is NP-Hard, they provided some evidence that this is likely true. Further work such as [113] have tried to settle the complexity of this problem, but there have been no major breakthroughs.

There has been much research from an optimization perspective on how to find an optimal, either partisan or proportional, districting scheme (for a summary see [2, 58]). There has also been work on how to detect or decide if a districting is drawn in a fair way or for partisan gain. The work in [108, 49, 91] list various measures and techniques. A novel approach to this problem was introduced in [40]. Here the authors attempt to find the number of districts won by each party in a random districting by using a Markov Chain Monte Carlo (MCMC) algorithm. If an actual outcome significantly deviates from this then partisan gerrymandering may have occurred.

Examining a related, but different, problem in [88] an algorithm is suggested that could be used to prevent partisan gerrymandering. The authors provide a mechanism where two opposing parties, each trying to maximize the number of seats they win, can partition a map into electoral districts. At a high level the first party proposes a partition of the graph, the second party then chooses one of the proposed districts and fixes it in place, the parties then swap roles where the second party proposes a partition of the non-fixed portions of the graph and the first party fixes one of those districts in place. They continue this process until all regions of the graph have been fixed. If the parties play optimally the final state of the game has some desirable properties[31].

In a somewhat related work, the authors of [11] relate the problem of gerrymandering to vote swapping. Taking an axiomatic approach to vote swapping, they find a relation between vote swapping and a restricted form of gerrymandering. Their research if significantly different from ours and the other cited ones, they are mainly interested in finding rules which are immune to vote swapping. A previous work [10] by the same authors is more relevant to our research. They look voting rules which are immune to manipulations by gerrymanderers. They find, under some assumptions, the only rule which is immune to gerrymandering is the constant one, where a predetermined candidate always wins.

### 2.4.4 Voter distribution and gerymandering

As we mentioned, since gerrymandering is just clever partitioning of voters based on their geographic location, these two concepts, voter distribution and gerrymandering, are intrinsically linked.

Perhaps the most prominent and direct example of shifting voter distribution is what's known as the urban-rural divide in the United States: for a variety of reasons, individuals are choosing to reside in such a way that they end up living near people with similar party affiliation. In particular, supporters of one party cluster in urban areas while surrounding rural regions fill up with supporters of the other party. The fact voters for each party do not distribute uniformly at random, and instead follow this increasing geographic split, was made prominent by the book "The Big Sort" [12]. There has been followup work which has corroborated this observation [23, 22]. This increasing urban-rural divide of voters is directly relevant to those in charge of drawing districts. Where people live, and vote, directly

[30] A precinct is a central location where residents of a neighbourhood physically vote. This is the most fine grain level of voter data, each one contains only a few hundred voters spread over a small geographic region.

[31] To understand them we need some knowledge from game theory which is too far removed from this work.

impacts the gerrymandering abilities of partisan groups. Already, many pundits [35, 98] are arguing this increasing divide is contributing to the problem of gerrymandering. Some commentators have even argued the urban-rural divide is more of a concern than gerrymandering. In the book "Why Cities Lose: The Deep Roots of the Urban-Rural Political Divide" [97], Rodden argues gerrymandering is only partially responsible for for the Democratic party's underrepresentation relative to their vote share. Instead, he believes a more pressing issue facing the Democrats is the urban-rural divide.

Finally, we note that the urban-rural divide is not just a phenomena in the United States. In [83] the author notes in many European countries there is an increasing political and economic divide between city and country. They argue rural regions feel left behind in terms of their economy, relative to their urban counterparts. This disparity leads to rural voters increasingly embracing positions which differ from ones held by their urban counterparts. These positions include, support for the United Kingdom leaving the European Union and the popularity of the far-right politician Marine Le Pen in France.

Here we will look at two recent papers, from the AI community, that address the impact of voter distribution on gerrymandering, although they don't address the urban-rural divide directly.

### Gerrymandering on a plane

A recent work by [66] directly addresses the impact of voter distribution on gerrymandering. First the authors formulate gerrymandering as a decision problem:

**Definition 10.** $gerrymandering_{\widetilde{f}}$ is a tuple containing:

- A set of $m$ candidates $C$.
- A set of $n$ voters $V$, where each voter $v \in V$ has a preference $\succ_v \in \pi(C)$ and a location in $\mathbb{R}^2$.
- A set of $z$ potential ballot boxes $B$, where each $b \in B$ has a location in $\mathbb{R}^2$.
- Target parameters $l, k \in \mathbb{N}$, where $l \leq k \leq z$.
- A target candidate $p \in C$.

The decision problem is as follows: Is there a set of $k$ ballot boxes $B' \subset B$, where if every voter in $V$ votes at their closest box in $B'$ and each ballot box picks a winner from $C$ using $\widetilde{f}$, and candidate $p$ wins at least $l$ of the districts.

This isn't quite gerrymandering on a planar graph as we have described it, but more of a continuous analog. By selecting the $B'$ ballot boxes and forcing the voters to their nearest box they are implicitly partitioning the region into a Voronoi diagram, where a point in $\mathbb{R}^2$ is assigned to ballot box $b \in B'$ when $b$ is closer than any other one. We note that they don't include what we consider one of the most important constraints, that each district in the partition contains a near equal number of voters. In fact, in their setup it is possible all but one district contain a single voter while the rest of the voters are assigned to one district. Despite these issues they do show $gerrymandering_{plurality}$ is NP-Hard.

In addition to their hardness results the authors examine gerrymandering from a non-worse-case view by using experimental data. On a 2D grid they place several cities then assign voters to different cities placing them according to a normal distribution located at the city centre. Each city has a canonical ranking over the candidates and the closer the voter is to their city centre the closer her preferences match her city's rankings. As a voter is placed further away from her city centre her ranking becomes less similar to her city's ranking[32]. The authors use a greedy heuristic that attempts to gerrymander this region for various candidates. They find as voters become more spread out from their assigned city centre (the variance of the normal distribution is increased) the ability of the small candidates (those who are not ranked at the top of many voter profiles) to win more seats than any other candidate becomes stronger.

As was the case with their theoretical results we must stress that their algorithm can produce districts which are vastly unequal in population size. In some outcomes the difference in population size is up to $6500\%$ a far cry from the $1\%$ that is demanded in many jurisdictions. Furthermore, by limiting their districting to Voronoi diagrams they place a somewhat unrealistic restriction on the ability of those drawing the districts. In our view if the goal is to understand how voter distribution impacts the ability to gerrymander it would be better to broaden the set of permissible district shapes. That being said, Voronoi diagrams do have nice properties as they tend to produce what we would consider compact districts[33] without odd protrusions.

### Gerrymandering on a (arbitrary) graph

Another related work [25] is closer to what we study in some respects, but further away in others. As with our work the authors assume the voters are embedded in some graph instead of a euclidian plane. Here voters are nodes, each with their individual preference over candidates. Furthermore, like our work they do not place any restrictions on the shapes of districts. Unlike much of the previous work, and also unlike our work, which assumes a planar graph, they allow any graph structure. This change essentially means they are no longer studying a geographic distribution of voters and an edge between two nodes does not necessarily represent physical adjacency, but instead some other form of connectedness. Formally they define the following decision problem:

**Definition 11.** $GraphGerrymandering_{\widetilde{f}}$ is a tuple containing:

- A set of $m$ candidates $C$.
- A set of $n$ voters $V$, where each voter $v \in V$ has a preference $\succ_v \in \pi(C)$.
- A set of $e$ undirected edges $E$, where every edge $(u, v) \in E$ is a pair of voters, $u, v \in V$.
- Target parameters $l, k \in \mathbb{N}$, where $l \leq k$.
- A target candidate $p \in C$.

The decision problem is as follows: Is it possible to remove edges from $E$ so that there are $k$ connected components, where each component holds an election and picks their winner using $\widetilde{f}$ and the rankings of the voters within the component, and candidate $p$ wins at least $l$ of the elections.

They show the $GraphGerrymandering_{plurality}$ is NP-Hard, even in the case where they restrict themselves to planar graphs. As in [66] they ignore balanced population constraints. They may create many districts with a single voter each and one district with the rest of the voters[34].

[32]The ranking of the voter is drawn using a Mallow's model where the ground truth is the ranking of their city and the distortion is their distance from the mean of the city distribution.

[33]What compact means is also a matter of debate.

[34]By ignoring population constraints they also avoid dealing with the issue of the unsettle complexity of balanced partition problems on graphs that we mentioned earlier.

The authors also provide a greedy algorithm for gerrymandering over graphs, their algorithm is basically an adaptation of the one found in [66]. As with the previous algorithm this method can produce highly unbalanced connected components. To mitigate the issue the authors include a constraint that the algorithm can never create two components with more than an $\alpha$ multiplicative factor difference in their population. For all of their simulations they set $\alpha = 5$, or up to a $500\%$ population difference (again far removed from the $1\%$ demanded in practice)[35].

Using their greedy algorithm the authors also study the impact of voter distribution on gerrymandering. As we mentioned, the notion of voter location and distribution is a bit different with non-planar graphs since edges no longer represent physical adjacency, at least not representable in a 2D euclidian world. Still edges are used to indicate some notion of connectedness. The authors focus on random graphs using the Erdős-Renyi and Barabasi-Albert models to generate the data. They further augment the models by increasing the probability an edge forms between two voters who are close to each other[36]. The model assumes a hompophily level for the graph. The existence of an edge between nodes is random process determined by the homophily level of the graph and the similarity of the nodes. The authors' main goal is to study how the homophily level impacts gerrymandering abilities.

They find that their algorithm struggles to gerrymander as the graph homophily increases. They suspect as the groups become more separated, with fewer connections between voters of different preferences, the opportunity to gerrymander goes away[37].

### 2.4.5 Our work

Here we will describe our contributions to the body of work on gerrymandering. These mainly come in the form of two distinct, but related works. The first aims to better understand a specific phenomenon we see cropping up in population data across the world, while the second aims to develop optimization techniques which can be used on real world data.

#### The urban-rural divide

As we mentioned, one of our goals is to study the urban-rural divide and its impact on gerrymandering, neither of the two papers which address gerrymandering and voter distribution directly examine this particular effect. To reiterate the urban-rural divide refers to a situation with two political parties[38], one with strong support in small densely populated urban regions and the other with strong support in the surrounding sparsely populated rural regions. Clearly [25] is not examining an urban-rural divide since it is not considering geographic relationships, instead the authors are interested in the case where voters with similar preferences end up connected. While this is a bi-product of the urban-rural divide it also covers situations which are far removed from reality. On the surface [66] comes closer to examining the situation we are interested in, but there is one important difference. In their setting all parties are urban parties, each with a different urban centre as their power base. The rural regions are where voters have almost random preferences, they support some urban party, just not necessarily the urban party for the city they belong to. Furthermore, we feel the issues we pointed out with both mechanisms make them unsuitable for studying population level effects on gerrymandering.

[35] The lowest they can set $\alpha$ would be 2, otherwise the algorithm would fail to converge.

[36] To generate voters and candidates they each are uniformly placed as a point on the $[0, 1]$ line. A voter will vote for the closest candidate to them on this line. To determine how close two voters are the authors consider their proximity to each other on the line.

[37] They note that as the number of target districts increase the performance of the algorithm does somewhat improve.

[38] This is largely inspired by the United States which is effectively a two party system.

Because of this our goal is twofold: We first require a model of voter distribution which captures, and only captures, the urban-rural divide, we also require a new mechanism which is able to produce partisan districts which obey population constraints. In our work we propose a solution to both of these items. Briefly, we mention that the model we use to study voter distribution is not one using real voter data, although it is inspired by it. Instead it is an easy to understand, but sufficiently complex, model of how voters distributed themselves. We made this simplification because of the technical limitations of existing optimization techniques and to remove any confounding factors present in real data[39].

#### Gerrymandering optimally

Inspired by the limitations of publicly available gerrymandering techniques, primarily their lack of performance guarantees, we also set a goal to develop optimal or near optimal automated techniques for gerrymandering. Using publicly available voter data we present a general optimization method that is able to nearly match, and sometimes exceed, the results presented in the 538-Atlas of Redistricting project [104]. We consider this result impressive for two reasons. Firstly, to our knowledge no published work has matched the performance of 538, which claims to have found the optimal results. Secondly, is 538's results are entirely hand crafted, they did not use any optimization techniques instead opting for the laborious process of "hand-drawing" districts. Our results are entirely automated with no human input and thus can generalize to any new situation.

## 2.5 The political party

In this section we will focus on the material necessary to understand chapter 5. As we saw in the previous section elections are not always a competition of individual candidates, but can be more of a competition between large political entities to which potential candidates belong. Interestingly, the party has not always been a mainstay of our political culture. Many of the framers of the United States constitution, including Thomas Jefferson, believed political parties would have a negative impact on the health of the nation [50]. Despite this political parties are incredibly prominent, Jefferson himself was forced to start his own party. Jefferson's party, the Democratic-Republican Party, had quite a successful run, wining many offices and outlasting opposition parties [111].

In this section we will explore, from an AI perspective, the role a political party plays in elections. We will focus our attention on one of the most important roles a political party has, the selection of candidates for elections. The previous section focused on the role a party's representatives played, but before a party puts a potential representative forward for a public vote there is another, internal, process that must take place. To pick a candidate to represent themselves parties engage in some internal selection process. Here, individual party members who wish to represent their party in the final election first compete internally to win the backing of the party. Using various mechanisms the party selects a special group of voters who hold a pre-election to pick the candidate that will go on to represent the party. This is perhaps the most important function a political party has, it ensures the candidate it puts forward for consideration to the general public represents the will of the party members. Furthermore, this internal mechanism must ensure that the candidate chosen is electable by the general populace. A candidate may perfectly represent her party, but if she is unappealing to the general populace she will not be elected. That is the views of her party will go

[39] These factors include, but are not limited to, the effects of city size and location in different states, and incomplete data.

unrepresented. Because of this whatever mechanism a party employs must balance two goals which may be at odds with each other.

There are a wide variety of mechanisms of which political parties make use of. Historically, it was the case some of, if not all of, the voters are current politicians within the party[40]. Recent developments have seen an increase in the power of individual voters, not just necessarily party power brokers, in this selection process. It has become quite common for registered party members (or even non-registered party members) to have the final say in the selection of party candidates [28]. Since the 1970s political parties in the United States have modified their internal mechanisms, now candidate selection is almost entirely handled by registered party members holding a vote for all types of candidate at all levels of government [24]. This system, known as the *primary*-system, has shifted the political power to voters who participate in it. To run with the support of a party in a general election a potential candidate must appeal to the primary voters in that party. It's not clear what kind of primary winner will go on to win the general election. There have been some works [84, 103] which study historic voting outcomes of primaries in an attempt to predict who will win an overall election; that is they try to predict how a "regular" voter will act based on the primary voters.

### 2.5.1 A spatial model for political parties

Motivated by our setting of political parties we adopt a slightly different model for voters and candidates than we have in the rest of this work. Recall, in standard voting literature voters are entirely described by an ordinal ranking of the candidates. In the previous section we extended this to include where voters physically reside, but preferences over candidates were still only ordinal. For our work on political parties we adopt a slightly more complex model, this model is motivated by the fact political parties often adopt views which differ from those of their competitors. These differences in social ideologies are often framed as a left-right divide [54]. Issues such as gun-control and access to abortion have become a wedge dividing the two major parties in the United States. The party position on these issues is a reflection of its member's position on the issues, that is we can describe a party's position as the positions of its voters and candidates.

To model the political positions of agents we will embed voters and candidates as points in $\mathbb{R}^d$ for some $d \in \mathbb{Z}^+$. We can think of each dimension of the space as an issue that parties take a position on. Where someone, a voter or candidate, is located along a dimension indicates their alignment with it, such as opposition or approval. For each voter $v$ we can take the straight line distance to each candidate $c$ to represent how dissimilar $v$ and $c$ are, this distance induces a ranking, where voters rank candidates in descending order of distance. While this is a restriction on voter preferences we feel it is not overly restrictive. It allows for preferences more complex than the common restriction of single-peaked preferences[41].

This model is advantageous for two reasons. Firstly, it provides a flexible model for distributing voters and candidates and defining political parties. By adjusting these distributions we are able to study how elections play out under various circumstances such as the members of different parties being separated in a euclidian, and hence political way. Secondly, in addition to giving the implicit ordinal ranking we now have cardinal values, we can measure how well a candidate represents the positions of a party and society at large by her distance to the other agents. This is a measure of agreement that we can use to measure the quality of an election outcome. Furthermore, this agreement measure is a by-product of the model, it is not something voters need express to the voting rule, or even calculate internally. We

[40]The Conservative party in the United Kingdom selects their leader by first having their MPs pick two candidates who are then put forward to a nation wide members vote.

[41]If $d = 1$ then the induced rankings will be single-peaked, but for higher dimensions the rankings need not have this property.

stress that all decision making processes, the voting rules that pick winners, will not make use of the cardinal values. We only ever study rules which take the ordinal rankings as inputs. Instead we will use the cardinal rankings as a way to measure candidate quality.

### 2.5.2 The distortion metric

With cardinal values we now have tools to analyze the quality of election winners. We can think of the utility of a candidate as the sum of the values each voter assigns them. In particular, we use the notion of distortion, introduced by [90]. Informally, distortion is the ratio of the utility of the election winner compared to the utility of the socially optimal candidate. If the cardinal scores for voters are almost arbitrary, the only restriction being that they are normalized so the value any voter assigns to all of the candidates must sum to one[42], then almost any rule based around the ordinal rankings[43] has unbounded distortion. By placing a somewhat strong restriction on the utilities they are able to come up with better upper-bounds.

### 2.5.3 Previous work with spatial models

Models that embed voters and candidates in some space have previously been examined in the social sciences. A summary of such work is provided in [102]. In general these works use both theoretical models and real election data to understand which candidates will win and how the political power will be distributed in various spatial models. Another question they seek to study is how party positions change over time, that is how the demands of an electable candidate and the demands of party ideology impact party policy within various systems.

Directly related to our approach are the works of [3, 4, 105]. These mainly focus on analyzing the distortion of various voting rules when voters and candidates are embedded in a metric space. With this embedding they are able to get around the lower bounds provided in [90]. A similar problem is examined in [38], but now with strategic voters. The question now is, is it possible for a mechanism to pick the socially optimal candidate. The existence of such a mechanism is dependent on the metric space and information provided to the voting mechanism.

### 2.5.4 Previous work with candidate elimination

There have been several lines of research that focus on which candidates partake in the final election. A few of these are from the optimization view. The first involves strategic candidates, those who have preferences over the outcomes of elections. These works [61, 31, 32, 19, 89] seek to understand when candidates decide to enter an election or sit it out. Another type of research focuses on a central authority deciding which candidates run in an attempt to control who will win the election (see [9] for its introduction). This control-based approach seeks to understand the complexity of identifying a minimal set of candidates who can be added to (or removed from) the election ensuring a particular candidate wins[44].

[42]We note that in our metric space embedding utilities are all non-positive, since we are measuring how much a candidate disagrees with a voter.

[43]Perhaps obviously the order in the ordinal rankings must respect the order in the cardinal ones.

[44]More generally, one can draw parallels between control problems and other parts of this thesis. For example, some control problems examine the partitioning of voters into different groups, this is somewhat related to gerrymandering which we study in chapter 4 (see [36] for more information). Other control problems deal with bribery, here an agent manipulates a voter to act strategically. Unlike our work on strategic voting (chapter 3), where voters choose when and how to manipulate on their own, bribery work is concerned with an external agent picking a voter to manipulate. For a general background on control, we recommend the chapter on it in [17].

Understanding the impacts of the mechanisms parties use to select the final set of candidates has mainly come from the political science literature. In [57] the authors provides an overview of the various methods used by various political parties. The authors of [29] focuses on democratic methods for candidate selection, that is where the deciders include party members and not just the politicians, as noted in [28] this is the trend in many western countries. There has been quite a bit of work on country specific analysis, [53] examine the United Kingdom, [110] studies Belgium and [52] examines Israel. The United States is a particularly interesting example since political parties have been a standard there since its founding [111]. Recently [24] examines how power players within a party try and impact the selection process by influencing the decisions of party members.

### 2.5.5 Our work

Our work aims to study the *primary* voting system, where various parties hold a pre-election with a subset of voters to whittle down the set of candidates they put forward. We will compare this system to the *direct* voting system, where there are no parties, and all voters vote over all candidates. To compare the systems we will use the distortion metric, comparing how much a primary helps or hurts the overall candidate quality when compared to its direct counterpart. Also motivated by the party system we will embed voters and candidates in an $\mathbb{R}^d$ space which helps model the positions a party and their members (voters and candidates) take. As mentioned, this space will implicitly define the cardinal, and hence ordinal, preferences of voters. This work differs from the previous work on distortion in metric spaces that we have cited, those focus on one shot voting rules, and do not consider the effects of narrowing the field of candidates with a primary-like mechanism.

# Chapter 3

# Strategic voting

> To improve is to change, so to be perfect is to have changed often.
>
> Sir Winston Leonard Spencer Churchill
> KG OM CH TD FRS PC

In this chapter we will focus on the topic of strategic voting, in particular we will work with the *iterative voting* model (first introduced in [74]). Recall, 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, voters reassess the outcome, and if any voter wishes to change their vote they may do so, and potentially a new winner replaces the previous one. The process terminates when no voter wishes to change their vote. Iterative voting is both a decision making tool that utilizes the ability of voters to behave strategically, hopefully arriving at a better outcome for society, and a model which captures the natural behaviour of voters who do not wish to "waste" their vote.

We will focus on two question related to iterative voting. First (in section 3.2), for various non-scoring rules we study how non best-response dynamics change the process. In particular we look at several restrictions of best-response that we consider natural, these are heuristics that voters may employ to simplify their strategizing. Second (in section 3.3), we examine a setup similar to the one we examine in chapter 4, where voters are partitioned into districts and vote for a district representative who then goes on to vote for the overall leader. These district-based elections are non-scoring rule elections[1], the districting system adds a level of complexity that has not been examined in the iterative voting literature. These *district-based elections* pose an interesting conflict for voters, now they must balance who their representative is, the person who advocates on their behalf to the main legislative body, and who the overall winner is. Perhaps surprisingly these objectives are not always aligned; for a voter it is possible to ensure a preferred party is the overall winner by choosing an undesirable party as the winner in their district.

## 3.1 The iterative voting model

The iterative voting model can be encapsulated by the concept of *dynamics*. A dynamic is a binary relation $\mathcal{D} \subset \pi(C)^n \times \pi(C)^n$. We call a (possibly finite) sequence of profiles $(\overrightarrow{\succ}_1, \overrightarrow{\succ}_2, ...) \in \pi(C)^*$ a *profile sequence* and a

[1]The overall election for a leader is also a non-scoring rule based election even if we use simple scoring rules like plurality.

(possibly finite) sequence of voters $(v_1, v_2, ...) \in V^*$ a *voter sequence*. A profile sequence $(\overrightarrow{\succ}_1, \overrightarrow{\succ}_2, ...)$ for which $\overrightarrow{\succ}_1$ 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 $\forall i(\overrightarrow{\succ}_i, \overrightarrow{\succ}_{i+1}) \in \mathcal{D}$. We will limit ourselves to profile sequences that differ only on a single element, i.e.,

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

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 $(\overrightarrow{\succ}_1, \overrightarrow{\succ}_2, ..., \overrightarrow{\succ}_k)$ will be called an *equilibrium* if there is no $\overrightarrow{\succ}_{k'}$ such that $(\overrightarrow{\succ}_k, \overrightarrow{\succ}_{k'}) \in \mathcal{D}$.

We call $\mathcal{D}(\widetilde{f}_t)$ the dynamic $\mathcal{D}$ which uses voting rule $\widetilde{f}$ with tie breaking rule $t$. Again this is just a binary relation on preferences profiles. That is, it is a binary relation $\mathcal{D}(\widetilde{f}_t) \subset \pi(C)^n \times \pi(C)^n$. For a dynamic $\mathcal{D}$ and voting rule $\widetilde{f}$ with tie breaking rule $t$, let $\mathcal{I}(\mathcal{D}, \widetilde{f}_t) = \{s | s \text{ is a valid profile sequence for } \mathcal{D}(\widetilde{f}_t)\}$. We will say that *iterative-$\widetilde{f}$ converges* under $\mathcal{D}$ if every element of $\mathcal{I}(\mathcal{D}, \widetilde{f}_t)$ is finite. Otherwise, we will say that iterative-$\widetilde{f}$ under $\mathcal{D}$ *cycles* or does not converge (or may not converge). Notice that, as defined, the semantics of convergence are agnostic to any scheduler which decides the path the profile sequence follows; that is $\mathcal{I}(\mathcal{D}, \widetilde{f}_t)$ converges if *every* element is finite. Our only restriction on scheduling, which we have mentioned before, will be only one voter may move at a time. Furthermore, for this chapter we will assume that the tie-breaking rule $t$ is simple lexicographical ordering of candidate names.

The dynamics we shall consider will be motivated by the truthful preferences, i.e., a dynamic in which a voter's vote changed must have increased the utility of that vote. Most of the relevant literature focuses on two main dynamics (fist studied in [74]). 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 $\overrightarrow{\succ}^{(1)}, \overrightarrow{\succ}^{(2)}$ and a voting rule F, $(\overrightarrow{\succ}^{(1)}, \overrightarrow{\succ}^{(2)}) \in$ BetterResponse iff:

$$\exists i \in V \;\; \text{s.t. } \overrightarrow{\succ}^{(1)}_{-i} = \overrightarrow{\succ}^{(2)}_{-i} \text{ and } F_t(\overrightarrow{\succ}^{(2)}) \succ_i^{tr} F_t(\overrightarrow{\succ}^{(1)}).$$

Such an $i$ is called the *manipulator* and $\succ_i^{tr}$ is their underlying truthful preferences, $\succ_i^{(2)}$ is called the *new vote*, and $\succ_i^{(1)}$ 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, $(\overrightarrow{\succ}^{(1)}, \overrightarrow{\succ}^{(2)}) \in$ BR iff:

$$\exists i \in V \;\; \text{s.t. } \overrightarrow{\succ}^{(1)}_{-i} = \overrightarrow{\succ}^{(2)}_{-i} \text{ and } F_t(\overrightarrow{\succ}^{(2)}) \succ_i^{tr} F_t(\overrightarrow{\succ}^{(1)})$$

and

$$\forall \succ'' \in \pi(C) \;\; \text{s.t. } (\overrightarrow{\succ}^{(1)}_{-i}, \succ''), F_t(\overrightarrow{\succ}^{(2)}) \succeq_i^{tr} F_t(\overrightarrow{\succ}^{(1)}_{-i}, \succ'').$$

The above description clearly defines a game. 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.

## 3.2 Dynamics beyond best response

As mentioned, our first topic of study is examining various dynamics beyond the best response dynamic. Furthermore, we focus on applications of these various dynamics to non-scoring rules.

Recall, as discussed in chapter 2 iterative voting literature has mainly focused on the best response dynamic with plurality. For non-scoring rules, however, it is not immediately clear if best response is appropriate. Indeed, in some cases, like STV, it is NP-complete to decide if there is a better response. Motivated by this complexity we present several dynamics that may serve as natural heuristics for a potential voter. As we discussed in chapter 2 there have been dynamics designed with the express purpose of ensuring convergence, as in $k$-pragmatism, M1, and M2 [93, 48]. However, we propose the following dynamics as more natural correspondences to the strategic behaviour 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. That is, a pair of sequence profiles is valid if the winner changed from $a$ to $b$ and the voter who manipulated submits a ranking with $b$ at the top of her ranking in the later profile.[2] In many of the voting rules we consider (and in particular any weakly-monotone rule), this dynamic is a subset of the best-response dynamic (i.e., $TOP(\pi(C)) \subset BR(\pi(C))$), and, indeed, it generalizes the dynamic used in [74].

**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. That is, a pair of sequence profiles is valid if the winner changed from $a$ to $b$ and the voter who manipulated submits a ranking with $b$ at the top of her ranking and $a$ at the bottom of her ranking in the later profile.[3] While in many scoring rules (e.g., plurality and veto) this is a subset of best response moves (generalizing those used in [62]), 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[4] from the current vote is chosen.

**SWAP:** This dynamic, inspired in part by notions from the literature on bribery (see, e.g., [34, 18]), is quite restrictive. It restricts manipulations to a single adjacent swap (called a 'shift' in the bribery literature), that is, changing to a vote within Kendall-Tau distance of one from the current vote (a 'swap' in the bribery nomenclature).

Dynamics like, TOP, TB, and SWAP all reduce the space a strategic voter needs to search. Instead of considering all $m!$ strategies they only need to consider $(m-1)!$, $(m-2)!$, and $m$ strategies respectively. TOP and TB also rely on the intuition of monotonicity of voting rules[5]. To make a candidate the new winner, set them as high as possible in the submitted rankings. To remove a winner, set them as low as possible in the submitted rankings. The idea behind SWAP and KT is to not change the current ranking too much. For SWAP, this means picking a ranking which is nearly identical to the current one. For KT, out of all the best responses pick the one which is the smallest deviation from the current one.

[2] The ordering of the remaining candidates in her ranking can be arbitrary.

[3] As with TOP the ordering of the remaining candidates can be arbitrary.

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

[5] Even for non monotone rules it is often the case increasing a candidate's rank is not a detriment for them.

### 3.2.1 Worst case convergence

In this section we consider the convergence of iterative voting for six 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-Hard [8, 7, 112].

*A note on reading the examples that follow*: 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 (after ties are broken). The i-th row 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. To better facilitate the reading of this chapter we will only present a few of the many proofs , all missing proofs can be found in section 7.1. For all of our examples we will assume there is a fixed linear ordering over the candidates, and ties are always broken according to this ordering.

#### Maximin

Similar to plurality and veto, maximin changes gradually. That is 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 after major restrictions on the admissible moves.

**Theorem 2.** *Maximin with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT, and SWAP ( We only include the example for BR, see section 7.1 for missing proofs.)*

*Proof.*

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

□

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.

#### Copeland

**Theorem 3.** *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 $\alpha$.*

*Proof.* Since the number of voters in all our examples is odd, they hold for Copeland$^\alpha$ for any $\alpha$. We will only show the example for the TOP dynamic (see section 7.1 for missing proofs):

$$\begin{tikzcd}[row sep=tiny]
DABC & DABC \arrow[r] & ACBD & ACBD \arrow[lll, bend right] \\
BDAC \arrow[r] & BACD & BACD \arrow[r] & BDAC \\
CDBA & CDBA & CDBA & CDBA \\
\mathbf{D} & \mathbf{B} & \mathbf{A} & \mathbf{B}
\end{tikzcd}$$ ∎

**Bucklin**

**Theorem 4.** *Bucklin with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT, and SWAP (we will only show the example for the TB dynamic see section 7.1 for missing proofs).*

*Proof.*
$$\begin{tikzcd}[row sep=tiny]
ADBC & ADBC \arrow[r] & ACBD & ACBD \arrow[lll, bend right] \\
CBDA \arrow[r] & DBCA & DBCA \arrow[r] & CBDA \\
CADB & CADB & CADB & CADB \\
DBAC & DBAC & DBAC & DBAC \\
\mathbf{A} & \mathbf{D} & \mathbf{A} & \mathbf{C}
\end{tikzcd}$$ ∎

**STV**

**Theorem 5.** *STV with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT, and SWAP (we will only show the example for the KT dynamic see section 7.1 for missing proofs).*

*Proof.*
$$\begin{tikzcd}[row sep=tiny]
DBAC & DBAC \arrow[r] & BDAC & BDAC \arrow[lll, bend right] \\
ACBD \arrow[r] & CABD & CABD \arrow[r] & ACBD \\
CDBA & CDBA & CDBA & CDBA \\
DABC & DABC & DABC & DABC \\
\mathbf{D} & \mathbf{C} & \mathbf{B} & \mathbf{A}
\end{tikzcd}$$ ∎

**Second Order Copeland**

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

See section 7.1 for proof.

**Ranked Pairs**

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

**Theorem 7.** *Ranked pairs with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT, and SWAP.*

See section 7.1 for proof.

### 3.2.2 Empirical analysis

In order to analyze the qualitative effects on outcome of iterative voting, we turn to empirical simulations. What makes one outcome better than another is a subtle question as there is no agreed-upon measure of quality. Furthermore, voting rules are defined with different goals in mind. For example, Maximin ensures that the core number of supporters a candidate has, against any other, is maximal (an objective not shared by other rules).

As we wish to see general properties of the interaction of voting rules and dynamics, we focused on a particular setting: 10 voters and 4 candidates. Profiles are generated by either sampling from a uniform distribution or a single-peaked one. For each voting rule, response dynamic, and distribution we sample 1000 different games, and because of the nondeterministic nature of iterative voting each of these games is repeated 100 times, each time with a different order of voter responses. Thus for each pair of game and dynamic we have up to 200,000 different executions. Iterative voting is executed until an equilibrium is reached, a cycle is detected, or some maximum number of iterations have elapsed. Though many sampled profiles start in equilibrium, we are interested in the effects of the iterative process, and focus on profiles where iterative voting occurred.

For both our voting rules and response dynamics, ties are broken in a deterministic fashion. In the case of a tie in a voting rule, out of all the potential winning candidates the lexicographical first is selected. For response dynamics that encounter ties, the first profile that was discovered is chosen.

One may ask why bother with iterative voting simulations, considering we have just shown they are not guaranteed to converge. However, despite these proofs, we did not encounter a single cycle in our millions of simulations (fewer than 6000 runs were stopped after reaching the cut-off number of 10,000 steps, and may have turned out to be cycles, but that still is a very low share). This indicates the relevance of examining iterative voting properties, even for voting rules that are not guaranteed to converge.

#### The truthful winner

One way of measuring quality, used in [107, 60], is how often the **truthful winners** emerges after iterative voting. While there is no guarantee that the truthful winners will emerge as the overall winners from iterative voting, it is often the case that they do (albeit in a non-truthful profile). As truthful winners are, in a sense, what the mechanism designers wanted the voting method to achieve, it is desirable that using iterative voting, they will be the rule's outcome. Frequently the initial truthful profile was an equilibrium under the chosen dynamic and voting rule. Surprisingly when the truthful profile was not an equilibrium and we had iterative voting the truthful winner often ended up emerging as the eventual winner in equilibrium (albeit from an non-truthful profile).

Approximately 78% of all sampled truthful profiles were an equilibrium (with single-peaked profiles almost 20% more likely to be an equilibrium than uniform profiles). Unsurprisingly, the more restrictive response dynamics had a higher ratio of truthful profile equilibrium. Hence, best response and Kendall-Tau had fewer truthful equilibria than SWAP. A similar disparity in the fraction of truthful profile equilibria is seen when examining the voting rules: the Condorcet consistent rules, Maximin, Copeland, etc., were more often initially in equilibrium than the non-Condorcet consistent rules, STV and Bucklin. Since the initial profiles are truthful, the Condorcet consistent rules will initially pick the Condorcet winner, if one exists. Moreover, as will be noted below, because of the Condorcet winner's appeal over each of the other candidates, it is less likely that iterative voting would lead to a "better" winner (for some metrics detailed below).

When iterative voting does occur, there does not appear to be a relation between how often the truthful winner emerges and the response dynamic. Instead the voting rule seems to be the more significant factor in determining how likely it is for the truthful winner to be chosen. For any dynamic and preference type none of the non-Condorcet consistent rules selected the truthful winner more than 45% of the time (all but one selected the truthful winner less than 40% of the time). However, in some regard, we will show the iterated winner is often an improvement over the truthful winner. On the other hand, Copeland, under any dynamic and preference type, selected the truthful winner at least 55% of the time. More generally with single-peaked preferences, the Condorcet consistent rules, except for maximin, select the truthful winner over 50% of the time under any combination of dynamic and profile type.

**Convergence to equilibrium**

In games in which iterative voting did take place, most dynamic and voting rule combinations converged, on average, within 10 steps, and except for SOC, reached fewer than 15 overall equilibria states. In general, while each voting rule is different, we mainly noticed significant differences between the dynamics. When using the KT dynamic, the pace to convergence was significantly longer than other dynamics in all voting rules except Bucklin (and for SOC, with uniform distribution, KT along with BR took far longer to converge than the rest). With KT for all rules, except Bucklin, the number of different equilibrium states reached was significantly higher. See Figure 3.1. For example, Copeland with KT averaged more than 25 steps to convergence with single-peaked preferences (and over 10 steps with uniform). Almost all cases of runs that had to be cut-off after 10,000 steps were the KT dynamic (for Copeland, just under 800 runs).

KT's behavior might be a bit surprising, since it is, fundamentally, a best response dynamic with a different tie-breaking rule—favoring votes close to one another (instead of lexicographic, pre-determined ordering). It would seem that for most voting rules, bias towards smaller, more local, changes when manipulating has a significant adverse effect on the convergence properties of the iterative game.[6]

A somewhat connected issue is the difference between SOC and Copeland, which differ in their tie-breaking rules. Unlike Copeland (which only had this with KT), all SOC dynamics had cases that did not converge after 10,000 steps (single-peaked struggled more than uniform ones).

This subtlety with tie breaking can be hard to pinpoint. With many voting rules, especially non-scoring ones, how a profile is set up in the short term can have a substantial impact in the long term. We speculate what may seem like an optimal move now may make certain candidates currently ranked lower more viable down the line. With more complex tie breaking rules, in addition to optimizing for the current winner, there is now a secondary tie breaking condition that may not be explicitly optimized against, but that can come into play in the long term. Candidates that would be vanquished under a predetermined tie breaking procedure could be lifted up by these more fluid rules and continue to compete in the long run, greatly affecting the convergence properties.

**Voter utility**

While the social welfare of the voters would be a compelling measure of the quality of an outcome, we, naturally, do not have access to the voters' utility functions. However, as has been suggested in previous research [107, 73], we can use the Borda score on the truthful preferences as a proxy for utility. Here the utility for each voter of a chosen candidate $c \in C$ which the voter ranks in place $i$ is $m - i$. The Borda score for a set of voters is the sum of the individual utilities. We study how iterative voting affects the Borda score of the winning candidate (see Figure 3.2).

[6] Conversely, TB influences most voting rules towards a faster and more focused, fewer equilibrium states, convergence.

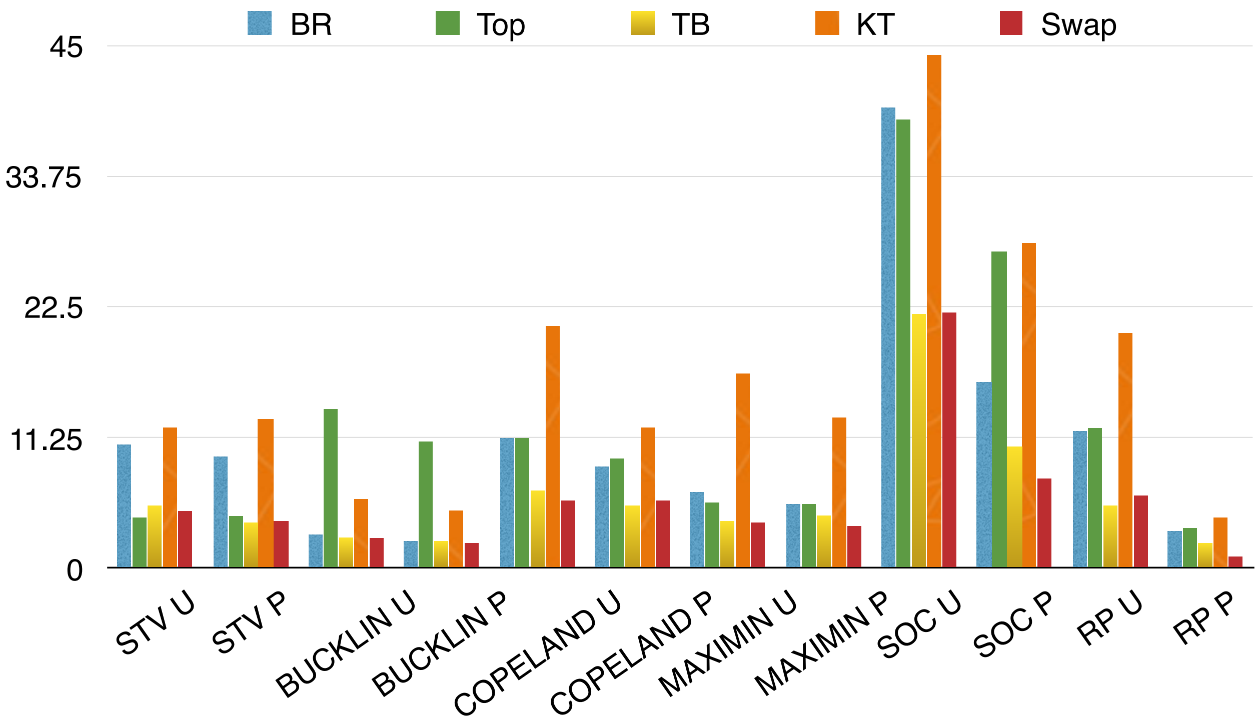


Figure 3.1: The average number of equilibrium states reached in 100 runs of each profile. KT can be obviously seen to be above most other dynamics (apart from in Bucklin), while TB is below them (U indicates uniform distribution; P indicates single-peaked one).

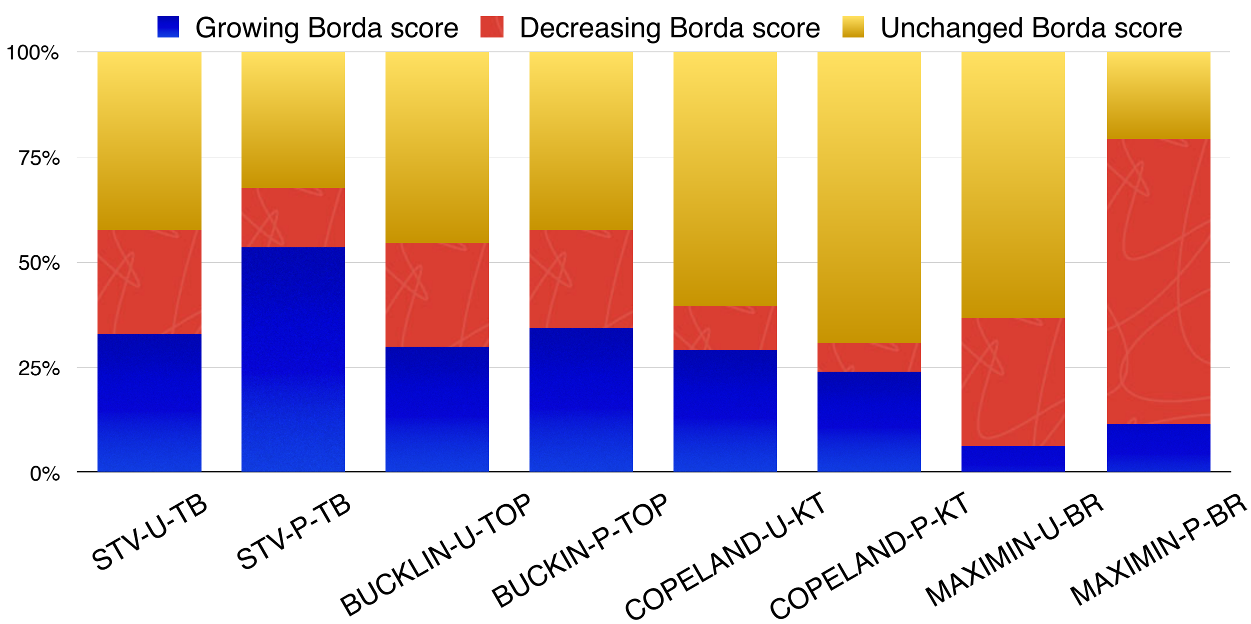


Figure 3.2: The ratio of winners in select voting rules, dynamics, and distributions, that had the same Borda score as the truthful winner, a higher Borda score, or a lower one.

Generally, the effect of dynamics on the Borda score of winners seems minimal. STV consistently showed significant improvements to the Borda scores under iterative voting (with single-peaked preferences doubling this effect). This effect was much less pronounced with Bucklin, although single-peaked again had a larger improvement. But with Maximin, winners' Borda scores went down more often than not, and again the effect was more pronounced under single-peaked preferences. Since Copeland frequently had the truthful winner emerge, its Borda scores were largely unchanged (especially under single-peaked preferences). It seems that when a Condorcet winner exists, when going through the iterative process Condorcet-consistent rules are less likely to find a candidate with higher Borda scores. Intuitively, this is because Condorcet winners commonly have a high Borda score, so it is harder to improve on the Borda score of the winner in these rules.

## 3.3 Iterative voting in districts

As mentioned, our second line of research in iterative voting is concerned with how it interacts with district-based elections.

As we discuss in detail throughout this work (especially in chapter 4) a particularly interesting subset of preference aggregation 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 of the district-based election research, even its computational parts [6, 66, 14] 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.

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 behaviour 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. Often it is the case that voters are concerned with a mixture of these outcomes, that is who represents them locally and who the overall leader is.

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.

For the remainder of this section we will make some additional assumptions for the set of voters $V$ of size $n$ and a set of candidates (or options) $C$ of size $m$. We now divide our voters between a set of districts $D = \{D_1, D_2, \dots, D_k\}$. We shall assume the district sizes are equal, so $n$ is divisible by $k$.

### 3.3.1 Convergence in district-based elections

Recall, in non district-based elections the following convergence results are known for various voting rules:

**Theorem 8** ([75] 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 9** ([63] Theorem 3)**.** *Iterative veto (with deterministic tie-breaking), when voters are myopic and pursue a best-response strategy, will always converge to a Nash equilibrium.*

**Theorem 10** ([63] Theorem 4)**.** *For any scoring rule except plurality and veto, iterative voting, even with deterministic tie-breaking, and even when voters are myopic and pursue a best-response strategy, have cases in which they will not converge to a Nash equilibrium.*

**Theorem 11** ([60] Theorems 1-6)**.** *Even with deterministic tie-breaking, and when voters are myopic and pursue a best-response strategy, iterative maximin, iterative Copeland, iterative Bucklin, iterative STV, iterative second-order-Copeland, and iterative ranked pairs*[7] *will all have cases in which they will not converge to a Nash equilibrium.*

As we alluded to earlier we are interested in studying two types of voters, those who care about who the overall winner is and those who care about the local representative is. Formally we define them as:

**Global** These voters, as in the "regular", non-district, version of iterative voting, wish to make the overall winner as favourable to them as possible. Hence if they are pursuing a better response dynamic, they will choose a vote that, regardless of what happens in their own districts, makes the global winner someone they prefer better over the current global winner.

**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.

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, that is one they prefer less than their current local winner, as it would effect the overall winner's identity in a way they prefer (see Example 3).

**Example 3.** 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 third 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

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

than $c$ which the previously voted for), making the overall winner candidate $b$, which they prefer over the previous overall winner, candidate $a$.

For local voters, all previous convergence results (Theorems 8, 9, 10, 11) 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 10, 11) 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 12.** *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 13.** *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 [75] 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 3.3.

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.

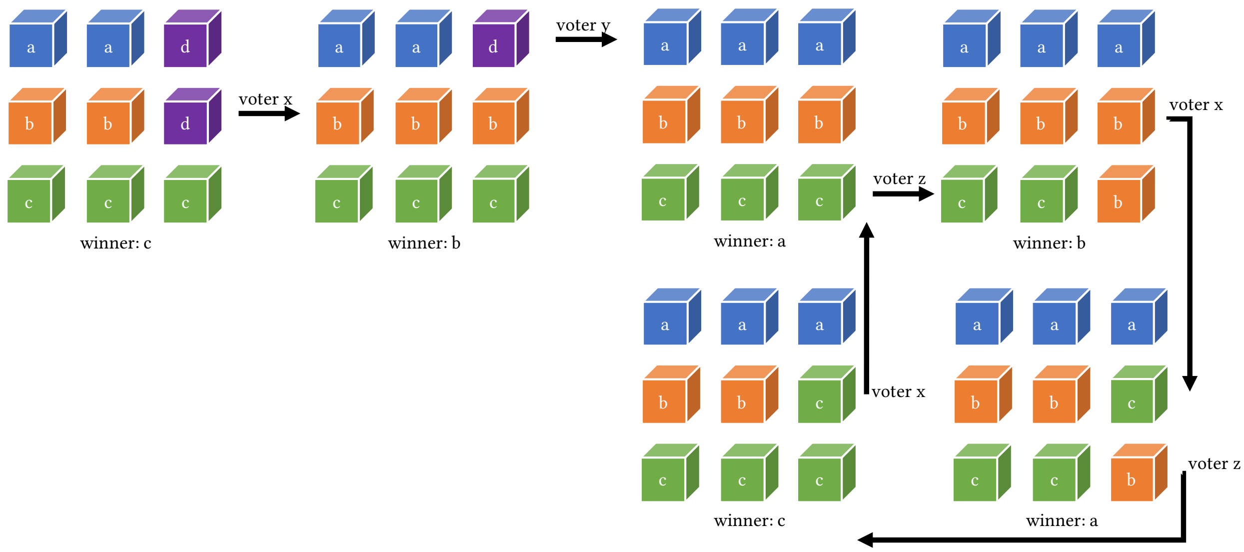


Figure 3.3: An overview of the voters causing the cycle for both plurality and veto in iterative voting when using district-based elections. Each box is a district, with the letter on it representing the district winner. An arrow coming out of a district indicates that district changing its winner because the voter indicated by the arrow (in that district) changed their vote. 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.

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 3.3), only containing slightly different voters to allow for similar strategic changes (which district votes for which candidates). Namely, 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:

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 13 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.

### 3.3.2 Empirical analysis

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[8] 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).

We examine the simulation results in two regards: how the iterated process proceeded, 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 from the perspective of mechanism design, seeing which voting systems are most suitable for a district-based system.

#### 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 12 and 13 that shows convergence is not guaranteed, in almost all cases, convergence was reached. In the few where it was not, it was because number of states exceeded our threshold limits. The cases where our

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

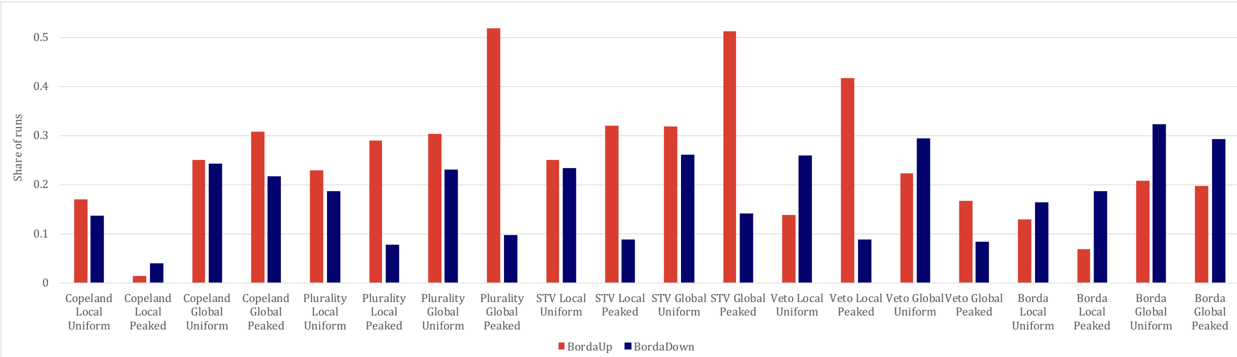


Figure 3.4: Examining the change in the Borda score of the overall winner, and whether it increased or decreased (from the truthful winner to the iterated winner).

limit was exceeded were quite rare. For almost all fixed election setting, a voting rule, voter type and preference, significantly less than $1\%$ of cases failed to converge. There was one exception, for Borda with globally minded voters and uniformly distributed preferences about $5\%$ of the runs timed out after 1000 moves. This is consistent with what was observed in the previous section. In general, Borda based systems struggled most with convergence. Here $20\%$-$40\%$ of election settings had at least one run of 200 in total which failed to converge. Other voting rules had far fewer issues, and 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. Despite the difference in convergence time, in all cases the iterative process was relatively focused. On average, for a fixed set of preferences and election setting, less than half of the candidates ever end up as an iterated winner. Overall, voters with single-peaked preferences tended to have fewer winners in each setting. We conjecture that since single peaked preferences guarantee a Condorcet winner in each district the winning opportunities for other candidates was limited.

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 winner, but their moves may change the winner's identity. However, global voter moves (see Figure 3.5) 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.

### 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.

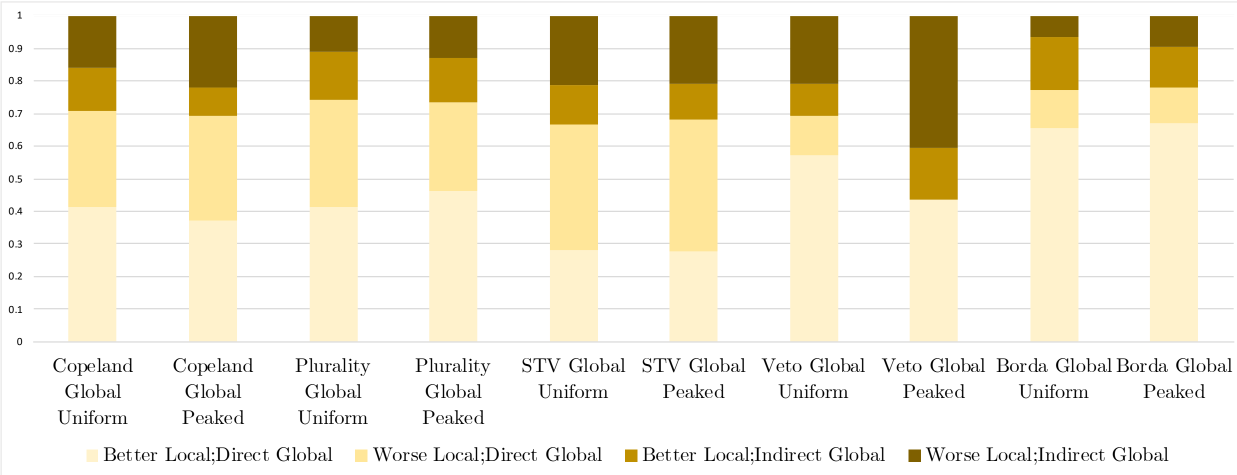


Figure 3.5: The fraction of move types of globally minded voters.

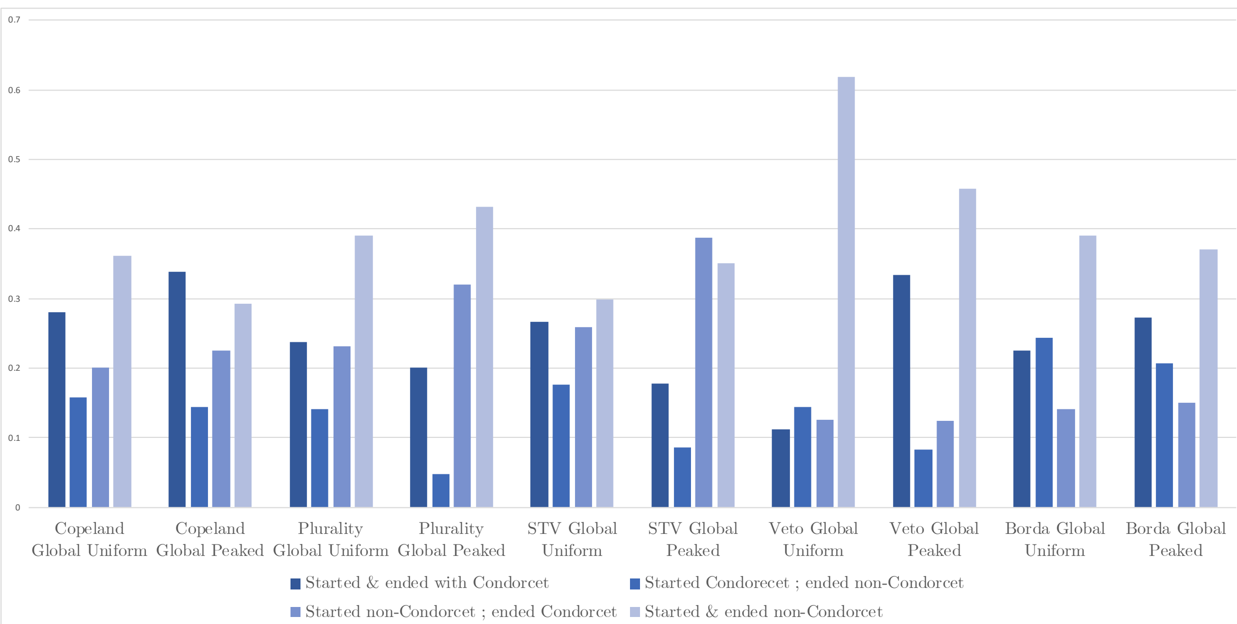


Figure 3.6: How the Condorcet winner emerges or disappears in runs where iterative voting took place and a Condorcet winner existed.

As we did in subsection 3.2.2 we can measure the iterative voting winner quality by how often the truthful winner emerges. Recall, our justification then was that 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 preference type (veto was something of an exception). This is consistent with what we saw in the previous section, which found the voting rule tended to determine how often the truthful winner was the iterated. 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 (which we also used and introduced in section 3.2). As before 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 3.4). 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 scores[9] (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[10] (where 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 3.6).

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 3.6, 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 move from the Condorcet winner than to find them.

## 3.4 Conclusion

In this chapter we investigated how strategic voting can impact the outcome in various elections. Firstly (in section 3.2), we examined how the well-studied model of iterative voting interacts with various non-scoring rules. In the first line of work we were mainly interested in better understanding how various voter dynamics beyond best response influence the process. Secondly (in section 3.3), we studied district-based elections behave when voters follow the best-response dynamic. In contrast to chapter 4 where we focus on how the strategic goals of those who draw the districts influence

[9] It is not as obvious the winner is not necessarily the candidate with the highest Borda score, in to the district-based system. It was shown in [6] 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.

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

election outcomes, here we instead focus on the strategic goals of the voters who partake in the election. Specifically we focus on how they must balance the utility they garner from their elected representative and the utility from the overall election winner.

While both these topics are distinct lines of research[11] there were some common themes we observed. Unlike plurality based elections where convergence guarantees are promised we found no such guarantees for any setting examined in this chapter. Neither the case of restricted voter dynamics for complex rules or plurality inserted into district-based elections guaranteed convergence in a finite number of steps. Despite these negative results our empirical analysis yielded simulations free of cycles[12], and often converged to an equilibrium in a few steps. Furthermore, these equilibriums reached via voter strategizing often led to more positive outcomes for society according to various metrics. Perhaps the most surprising theme (or at least one we least anticipated) was the impact of the Condorcet winner. Often it was the case voters were unable to move away from the Condorcet winner, either once they had been made the iterated winner or if they were the initial winner. Furthermore, we saw discovering them (or moving away from them) was highly correlated with an increase in voter utility (or a drop in voter utility). In retrospect these effects make sense, the Condorcet winner is a very strong candidate who should be hard to strategize against, but nevertheless it was a surprise to us.

[11] Both were undertaken years apart and appeared as different publications.

[12] We should note that some cases ran for quite awhile before we terminated them, but they never recached a previously seen state.

# Chapter 4

# Gerrymandering

> The evils we experience flow from the excess of democracy.
>
> Elbridge Thomas Gerry

In this chapter we will focus on how political districts are drawn and how that impacts the outcome of elections. As a reminder, in district-based elections the region (henceforth a state) where voters live is partitioned into non-overlapping segments called districts. The voters within a segment hold an election to pick a representative. Voters can still vote for political parties, but now they vote for a candidate who was picked by a political party to contend in that district. The more representatives a party gets elected, that is the more districts they win, the more power they have in the legislative body. Thus it is possible for a party to have more power than what is proportional to their vote share.

In many U.S. states the boundaries of the electoral district are determined by the government in power. These lawmakers can use clever map drawing techniques to strengthen their party's position, this practice is named gerrymandering. Oftentimes they can stretch a minority of the vote share into a majority of the representatives.

While this work is highly inspired by the gerrymandering that has run rampant across the United States (in fact we often talk about dividing states) we note these ideas can be applied to any political body where its members represent geographic regions. Because we are highly inspired by the political situation and gerrymandering in the United States we will focus on simplified elections with just two parties[1] and the plurality rule[2].

We will explore two distinct, but as it turns out somewhat related, directions for gerrymandering. First (in section 4.2), we will focus on how the urban-rural divide impacts the ability of a party to gerrymander. Briefly, the urban-rural divide is a pattern present in many western countries where supporters of one political party tend to congregate into densely packed urban centres and supporters of the opposition party spread out in the surrounding rural regions. For this direction we will develop an optimization technique based around integer-linear programming and a model of the voter distribution for the urban-rural divide on simulated graphs. Second (in section 4.3), we will focus on algorithmic ways to optimally gerrymander on real voter data from various U.S. states. In this direction our objective will be slightly different, we will consider totally automated methods and we will compare against a supposedly

[1] In the United States there are only two political parties with representation at the federal level.

[2] Recall Duverge's law says for plurality we will end up converging to a two party system.

optimal baseline which manually incorporates human expertise.

## 4.1 The model

A state in the United States must be divided every ten years into a pre-determined number of political districts. For convenience we won't work with actual state maps, but instead with a graph theoretic formulation. In our graph theoretic formulation a state is represented by a graph $G$, where the vertex set $V(G)$ contains a vertex for every precinct[3], and the edge set $E(G)$ contains an edge between every pair of precincts that share a physical boundary. Because the map of a state is two-dimensional, we assume that $G$ is planar. Let $n_v$ denote the number of voters in vertex $v$. For simplicity, we assume that voters are divided between two major political parties, $P_1$ and $P_2$. For $P \in \{P_1, P_2\}$, let $n_v^P$ denote the number of voters of party $P$ in vertex $v$, and let $N^P = \sum_{v \in V(G)} n_v^P$ denote the total number of voters of party $P$. Let $N = N^{P_1} + N^{P_2}$ denote the total number of voters. We use $\alpha^{P_1} = N^{P_1}/N$ and $\alpha^{P_2} = N^{P_2}/N$ to denote the proportional vote shares of the two parties.

Given a desired number of districts[4] $K \in \mathbb{N}$, the *districting problem* is to partition the graph $G$ into $K$ vertex-disjoint subgraphs $G_1, \ldots, G_K$ (called *districts*) that satisfy a number of constraints. In this work, we focus on two constraints that exist widely in practice.

1. *Contiguity.* For each district $k \in [K]$, $G_k$ must be a connected subgraph of $G$.

2. *Equal Population.* The total number of voters in each district should be approximately equal. Formally, given a tolerance level $\delta$, we need that for each $k \in [K]$,

$$1 - \delta \leq \frac{\sum_{v \in V(G_k)} n_v}{N/K} \leq 1 + \delta.$$

We say that a districting is valid if it satisfies both these constraints. Let $\mathcal{R}$ denote the set of valid districtings. There are additional criteria that districting should satisfy such as compactness of districts, preservation of existing political communities, and racial fairness. However, we overlook these criteria in this work, as there is still work to be done on formulating a consensus on their quantitative definitions.

There may of course be many solutions to the districting problem satisfying the contiguity and equal population constraints. The goal of gerrymandering is to find the districting that maximally favours one party. In this work, we focus on partisan gerrymandering (henceforth, simply gerrymandering), where the goal is to maximally favour a given party.

How one defines what is partisan can be surprisingly difficult. Is our goal, as partisan map-drawers, simply to draw as many districts as possible where our preferred party has a majority of the votes? We could also take a probabilistic view, where we say voters behave similar to their historic behaviour but not exactly. Now our goal could be to ensure we win many districts with high probability. In subsequent sections we will explore various ideas like these.

[3]A precinct is the smallest geographic unit for which we have voter data in the United States, it represents a neighbourhood of a few hundred people who vote in the same locations.

[4]We again stress, the number of districts is fixed. Often the number of districts is a simple function of population.

## 4.2 Voter geography and gerrymandering power

In this section we are particularly interested in how voter distribution affects a partisan map-maker's ability to gerrymander. In particular, we are interested in a measure of ability to gerrymander we call *gerrymandering power*.

We define the gerrymandering power of party $P$ to be $\max_{R \in \mathcal{R}} \sigma^P(R) - \alpha^P$, i.e., the maximum boost the party can get through gerrymandering above their proportional share of the districts. Note that negative gerrymandering power implies that the voters are distributed in such a way that the party falls short of its proportional share of the districts even with maximum gerrymandering. Given a districting $R \in \mathcal{R}$, we say that party $P$ wins district $k$ if it has a majority in the district: $\sum_{v \in V(G_k)} n_v^P > (1/2) \cdot \sum_{v \in V(G_k)} n_v$[5]. Let $K^P(R)$ denote the number of districts won by party $P$ in districting $R$, and let $\sigma^P(R) = K^P(R)/K$.

For this section we also make some simplifying assumptions. First, we assume that graph $G$ is an $n \times n$ grid. Grids are among the simplest planar graphs that still present non-trivial challenges. Second, we assume that each vertex of the grid has an equal number of voters: let $n_v = T$ for each $v \in V(G)$, for a sufficiently large constant $T$. Third, we mandate that all districts be of equal size, i.e., we set $\delta = 0$ in the equal population constraint. Finally, we assume voter preferences to be fixed. While these assumptions drag our model a bit farther from reality, they allow us to focus on the dependence of a party's gerrymandering power on its vote share and the geographic distribution of its voters.

### 4.2.1 A worst case view

Our goal is to study the effect of voters' geographic distribution on the gerrymandering power of the parties. In this section, we take a worst-case point of view: *How does the gerrymandering power of a party change with its vote share when its voters are distributed in the worst possible way?* Formally, given a party $P$ and its vote share $\alpha^P$, we want to analyze the maximum fraction of districts the party can win in the worst case choice of $\{n_v^P\}_{v \in V(G)}$ that satisfies $0 \leq n_v^P \leq T$ for each $v \in V(G)$ and $\sum_{v \in V(G)} n_v^P = \alpha^P \cdot N$. For the grid graph, $N = n^2 \cdot T$ is the total number of voters.

We begin by making an interesting observation in the large-graph limit. Imagine the $n \times n$ grid embedded in a bounded convex region. As $n \to \infty$, one can treat the graph as a continuous convex region in $\mathbb{R}^2$ endowed with two measures $\mu^{P_1}$ and $\mu^{P_2}$ that represent how the voters of the two parties are distributed across the region. In this case, we show that there is a sharp transition in which a party can win every district or no district depending on whether it has a majority or a minority vote share.

The idea of the proof is as follows. When $\alpha^P < 1/2$, party $P$ clearly wins no districts if its voters are uniformly spread, i.e., if it has $\alpha^P$ fraction of the voters in each individual vertex. When $\alpha^P \geq 1/2$, we invoke a generalization of the popular Ham Sandwich Theorem [106, 56], which states that given $d$ measures in $\mathbb{R}^d$, there exist $K$ interior-disjoint convex partitions that divide each measure equally. Applying this to measures $\mu^{P_1}$ and $\mu^{P_2}$ in $\mathbb{R}^2$, we get a valid districting in which party $P$ has a majority in every district.

**Theorem 14.** *Suppose an n x n grid is embedded into a bounded convex region. As $n \to \infty$, for every $K \in \mathbb{N}$, party $P$ (which controls the districting) can guarantee winning every district if its vote share is $\alpha^P \geq 1/2$, and wins no districts in the worst case if its vote share is $\alpha^P < 1/2$.*

When $n$ is finite and $\alpha^P < 1/2$, uniform voter distribution still remains a worst case for party $P$ regardless of the number of districts $K$, and prevents the party from winning any district. However, the case of $\alpha^P \geq 1/2$ becomes

[5]We break ties in favour of the party which controls the districting procedure

more fine-grained. For a constant $K$, increasing the graph size (i.e., increasing $n$) gives the party more gerrymandering power. We illustrate this using the case of two districts ($K = 2$). For $n = 2$ (i.e., in a $2 \times 2$ grid), it is easy to show that a party needs $75\%$ vote share to win both districts in the worst case.

**Proposition 15.** *For $n = K = 2$, For party $P$, which controls the districting, the following holds in the worst case for them.*

1. *If $\alpha^P \geq 3/4$, the party wins both districts.*
2. *If $1/2 \leq \alpha^P < 3/4$, the party wins a single district.*
3. *If $\alpha^P < 1/2$, the party wins no districts.*

*Proof.* Let us number the precincts in the $2 \times 2$ grid as $\{1, 2, 3, 4\}$ so that the number increases from left to right and from top to bottom. For $i \in \{1, 2, 3, 4\}$ and $P \in \{P_1, P_2\}$, let $n_i^P$ denote the number of voters of party $P$ in precinct $i$. Suppose party $P$ is gerrymandering. Note that there are two ways to district: either each row is a district, or each column is a district.

Suppose $\alpha^P \geq 3/4$, i.e., $\sum_{i=1}^{4} n_i^P \geq 3N/4$. We want to show that party $P$ holds a weak majority either in each row or in each column. Suppose this is not the case. Without loss of generality, suppose it holds a minority in the top row ($n_1^P + n_2^P < N/4$) and in the left column ($n_1^P + n_3^P < N/4$). Summing the two equations, and subtracting from the total vote share of party $P$, we obtain $n_4^P - n_1^P > N/4$, which is impossible since each precinct has $N/4$ voters.

Next, suppose $1/2 \leq \alpha^P < 3/4$. It is easy to see that regardless of the districting, party $P$ must win at least one district because $\alpha^P \geq 1/2$. For any $\epsilon > 0$, we want to show an instance with $\alpha^P = 3/4 - \epsilon$ in which party $P$ cannot win both districts. One such instance is given by $n_1^P = n_4^P = N(1/4 - \epsilon/2)$, $n_2^P = N/4$, and $n_3^P = 0$.

Finally, let $\alpha^P < 1/2$. If the voters are uniformly spread, the party trivially cannot win a majority in any district, regardless of the districting. $\square$

However, as $n$ increases, we can show that the required vote share for winning both districts quickly converges to the $50\%$ limit indicated by Theorem 14. In the next result, we only consider even $n$ because creating two districts of equal size is impossible when $n$ is odd.

**Theorem 16.** *For even $n$ and $K = 2$, if party's vote share is at least $1/2 + 1/n$ it can find a districting where it wins both seats.*

*Proof.* Consider an $n \times n$ grid. Suppose party $P$ has vote share $\alpha^P \geq 1/2 + 1/n$. We want to show that there exists a valid districting in which the party wins both districts.

To take care of the contiguity and equal population constraints, let us impose a specific structure on the districting. We assign the top row consisting of $n$ vertices to district 1, and the bottom row consisting of $n$ vertices to district 2. This leaves $n$ columns of height $n - 2$ each, which we call *strips*. Note that every solution in which $n/2$ strips are assigned to each district gives a valid districting. We want to show that one such assignment results in party $P$ winning both districts.

Suppose this is not true. Consider the assignment that maximizes the minimum vote share of party $P$ across the two districts. Without loss of generality, suppose party $P$ wins district 1, but loses district 2. Let $n_1^P$, $n_2^P$, and $n_t^P$ denote the number of voters of party $P$ in district 1, district 2, and a strip $t$, respectively. Recall that the total number of voters is $N$.

Since party $P$ loses in district 2, which has $N/2$ voters, we have $n_2^P < N/4$. Hence, there exists a strip $t$ in district 2 such that $n_t^P \leq n_2^P/(n/2) < (N/4)/(n/2) = N/(2n)$.

On the other hand, we have $n_1^P = \alpha^P N - n_2^P > \alpha^P N - N/4$. Even after discounting the top row which has $N/n$ voters, there must exist a strip $t'$ in district 1 such that

$$n_{t'}^P \geq \frac{\alpha^P \cdot N - N/4 - N/n}{n/2}. \tag{4.1}$$

Let us consider the (valid) districting obtained by exchanging strips $t$ and $t'$ between the two districts. We observe that party $P$ still wins district 1 because by losing strip $t'$, it loses at most $N/n$ of its own voters, and

$$n_1^P - \frac{N}{n} > \alpha^P \cdot N - \frac{N}{4} - \frac{N}{n} \geq \frac{N}{4},$$

where the last inequality follows because $\alpha^P \geq 1/2 + 1/n$. On the other hand, district 2 now has strictly more voters of party $P$ because it loses at most $n_t^P < N/(2n)$ such voters, but gains at least $n_{t'}^P$ such voters. From Equation (4.3) and the fact that $\alpha^P \geq 1/2 + 1/n$, it readily follows that $n_{t'}^P \geq N/(2n)$. We have created a partition with a higher minimal voting share for party P, contradicting our construction. □

While the party with a majority vote share can easily gerrymander large graphs when $K$ is fixed, it is much more difficult to do so when $K$ is large as well. At the extreme, when $K = n^2$, it is easy to show that party $P$ wins $\max(0, 2\alpha^P - 1)$ fraction of the districts in the worst case. This fraction is zero for $\alpha^P \leq 1/2$, and linearly increases to 1 as $\alpha^P$ goes to 1. This is in sharp contrast to Theorem 16, where the fraction jumps from 0 to 1 when going from $\alpha^P = 1/2$ to $\alpha^P = 1/2 + 1/n$.

While our results are for the extreme cases (Theorem 14 holds as $n$ goes to infinity and Theorem 16 holds for $K = 2$), the worst-case viewpoint lends us to a key insight: the gerrymandering power of a party significantly depends on the relationship between $n$ and $K$. While large graphs are easy to gerrymander, a large number of districts make it hard to gerrymander.

### 4.2.2 Optimizing for gerrymandering

We now conduct an empirical study of the gerrymandering power of political parties. Instead of the worst case partisanship distribution we considered in the previous section, we adopt a more realistic model based on the urban-rural divide referenced in the introduction. We also use grid graphs with a less extreme ratio of the graph size to the number of districts (in fact, we use numbers that are similar to some states within the American Congressional system).

### An urban-rural model

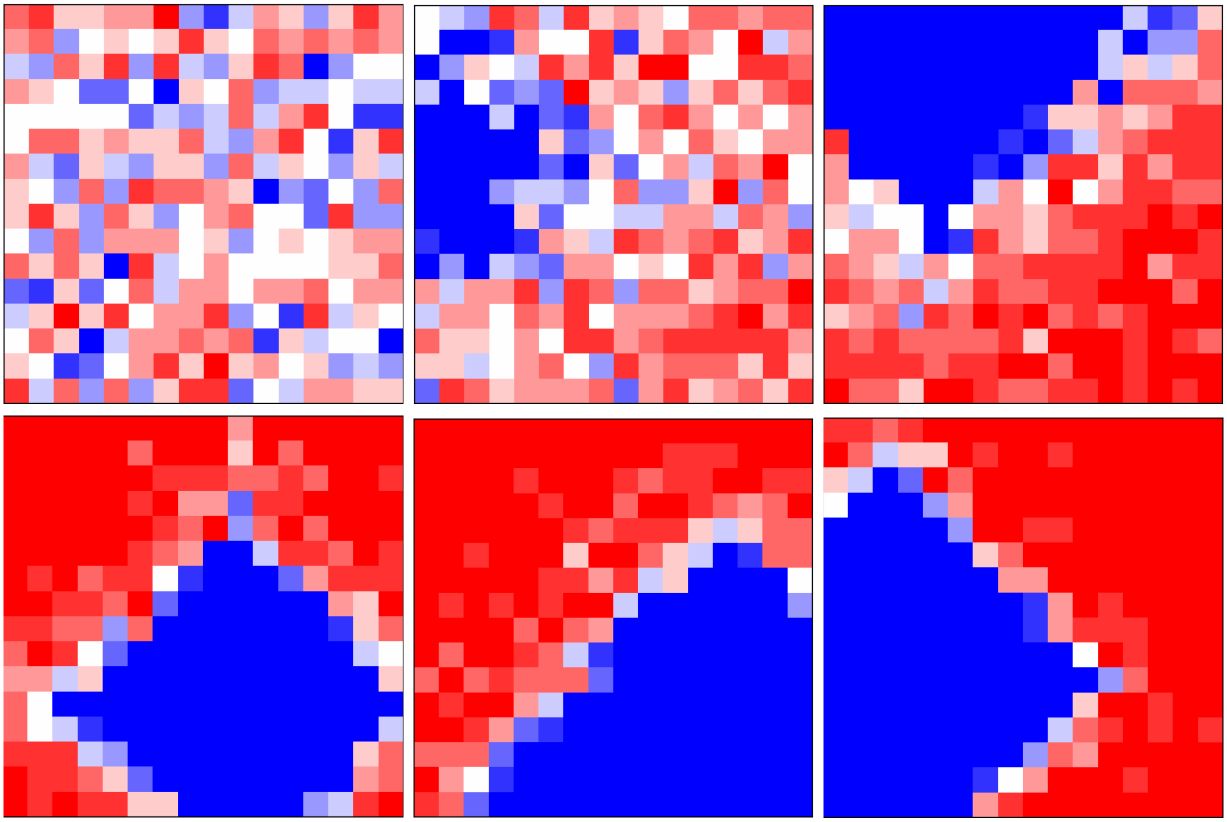

Figure 4.1: Partisanship distribution for $\alpha^U = 0.45$ and various values of $\phi$ ($\phi = 0, 1, 3$ in the top row from left to right and $\phi = 5, 7, 9$ in the bottom row from left to right). Blue/red represents a majority of $U/R$ voters, and colour intensity increases with the majority strength.

To model an urban-rural divide on a graph $G$, we use two parameters. The fraction of the urban party $U$'s voters in $G$, $\alpha^U \in [0, 1]$, and the strength of an urban-rural divide $\phi \in \mathbb{R}_{\geq 0}$ . Given $G, \phi$ and $\alpha^U$, we use the following process:

1. Set all voters in $G$ to be for the rural party $R$.
2. Pick a set of urban centres $C \subset V$ randomly. For $v \in V$, let $d(v)$ be the minimum distance of $v$ to any $c \in C$.
3. Pick a node $v$ (with at least one $R$ voter left) with probability proportional to $\frac{1}{1+(d(v))^{\phi}}$.
4. Convert one of its $R$ voters into a $U$ voter.
5. Repeat steps 3 and 4 until the fraction of $U$ voters in $G$ is at least $\alpha^U$.

See Figure 4.1 for sample heat maps generated by this process. Note that in step 3, we pick a node with a probability that decays polynomially in $d(v)$; we also conducted experiments with exponentially decaying probabilities and did not notice a qualitative difference in our results.

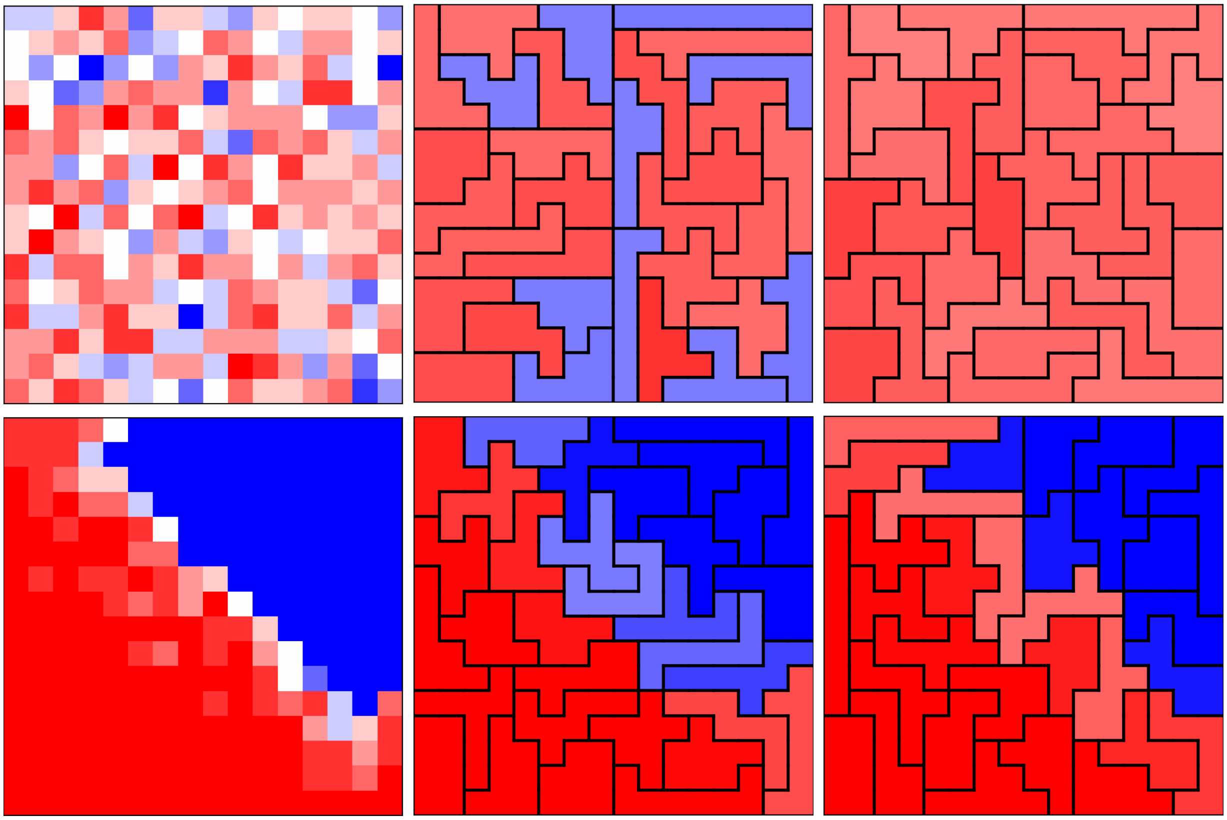

Figure 4.2: An example of the gerrymandered solutions $B^+$ found when $\alpha^U = 0.4$. Top row $\phi = 0$ bottom row $\phi = 10$. First column is partisan distribution of voters (similar to Figure 4.1). Second column is the gerrymandered solution for $U$, third column is the gerrymandered solution for $R$. Colour intensity within a district increases with the strength of victory.

**An algorithm to gerrymander**

Our starting point for an algorithm for optimal gerrymandering is to formulate a Mixed Integer Linear Program (MILP), which uses network flow constraints to ensure connectedness of the districts. For full details of our program see subsection 7.2.2 in the appendix.

Unfortunately, this program does not scale well, and takes hours on grids with a hundred nodes. Let us call the MILP approach algorithm $A$. We devise a bottom-up algorithm $B$, which uses $A$ as a subroutine to optimally solve small sub-problems with at most $\beta$ nodes. To divide $G$ into $K$ components in favor of party $P$, algorithm $B$ works as follows:

1. Find an arbitrary division of $G$ into $K$ connected components $(G_1 \cdots G_K)$ of equal or near-equal size.
2. Randomly pick two adjacent components $G_i$ and $G_j$.
3. Merge them into a new component $G_M = G_i \cup G_j$. If $|V(G_M)| \leq \beta$, use algorithm $A$ to optimally gerrymander $G_M$ into $K'$ districts, where $K'$ is the number of districts in $G_i$ and $G_j$. Otherwise, let the districting of $G_M$ be dictated by the districting of $G_i$ and $G_j$.
4. Repeat steps 2 and 3 until there is one component left.

Finally, we chain algorithm $B$ with itself by feeding the districting found in one execution of $B$ to step 1 in the next execution of $B$, and repeating until there are no improvements. We call this algorithm $B^+$. While the algorithm is not guaranteed to find an optimal gerrymandering, we see (see Section 4.2.2) that it finds highly gerrymandered districting on large instances; in contrast, algorithm $A$ simply fails to work for large instances. There are a few options for finding a districting for step 1 in the first execution of algorithm $B$:

1. The first option is using our mixed integer program with no objective, that is it returns the first districting it finds. This does work for smaller problem instances. Unfortunately, we found for problems close to the size we study this technique was too slow.
2. The second option would be using an existing districting, that is take the plan currently used and "improve" upon it by feeding it into $B$. Since we mainly focus on simulated data in this chapter we will not use this option.
3. The third would be to start with an arbitrary districting that is hand drawn. Because of the constraints, population and connectedness, and the difference in geography between various maps, this option is only viable for small and simple settings.
4. The fourth, and final option, uses Algorithm 1 which we will formally introduce in section 4.3. At a high level this algorithm draws a random spanning tree (using Kruskal's algorithm), it then cuts $k-1$ edges of the tree (somewhat at random) which divides the tree into $k$ connected components. If these components satisfy population constraints then we use them as our districting, otherwise we draw another random spanning tree and repeat the procedure.

Once we find one valid districting, we can find more for different executions of $B^+$ using an iterative process $I$, where we take a pair of adjacent districts $G_i$ and $G_j$, find one node from each district such that exchanging them gives another valid districting (if possible), and repeat this for a number of steps.

**Simulation setup**

For all of our experiments we use a $16 \times 16$ grid graph (i.e., 256 nodes) with 10 voters per node, and divide it into 32 equally sized districts. This problem size is about the same as Vermont's state senate (270 precincts and 30 districts). For the urban party vote share, we use $\alpha^U \in \{0.40, 0.45, 0.48, 0.5, 0.52, 0.55, 0.6\}$, and for the strength of the urban-rural divide, we use $\phi \in \{0, 0.5, 1, \ldots, 9.5, 10\}$. Using our urban-rural model, we generate $40$ graphs $G$ for each combination of $\alpha^U$ and $\phi$, each with a randomly chosen urban centre (more centres would be too crammed for a $16 \times 16$ grid graph). For each $G$, we run $B^+$ 20 times to find the best gerrymandering for each party[6]. To generate the $40$ sufficiently different starting points, we use process $I$. Half the starting points, that is 20, are generated by running the hand drawn starting point[7] through $I$ with 100,000 swaps. The other 20 starting points are generated using the tree-districting and running that through $I$ with 100,000 swaps.

We use IBM CPLEX and GUROBI for solving the MILP in algorithm $A$, and use $\beta = 16$, i.e., we solve instances with at most 16 nodes optimally using algorithm $A$. Overall, when provided with a starting point, algorithm $B+$ was able to solve any of our instances within 2 minutes[8].

[6] The gerrymandering outcomes were fairly consistent for fixed values of $\phi$ and $\alpha^U$. The standard deviation over the 40 maps was always under 1.4; in comparison the number of districts was 32.

[7] Since this is a $16 \times 16$ grid graph we divide each row of the graph into two equal sized districts.

[8] There was no noticeable difference in solve times between the two solvers, all simulations were carried out on laptops with fewer than 8 cores and at most 16 GBs of memory. No system was newer than a 2019 MacBook Pro.

**Some basic results**

Previsouly, we show that for $K = 2$, a party needs at least $50\%$ vote share to guarantee winning at least one district in the worst case. In our simulations with a moderate urban-rural divide ($\phi = 5$), we observe that just $26\%$ vote share allows a party to win one district with $n$ as low as $8$.

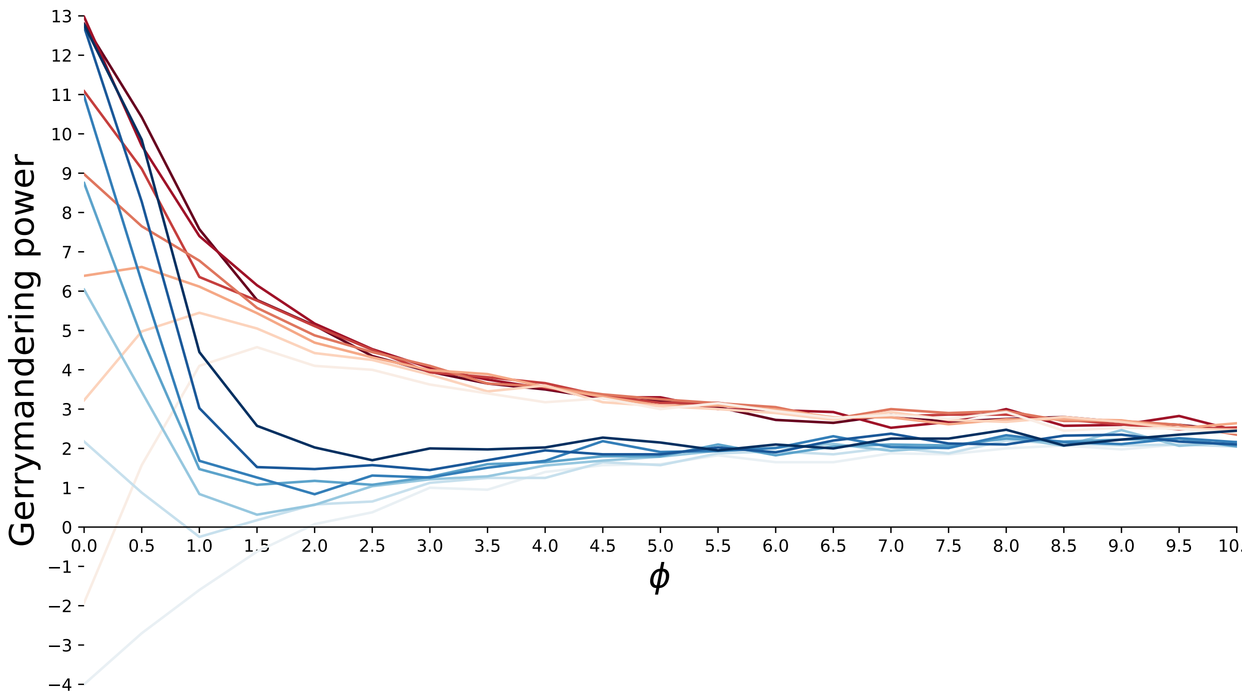


Figure 4.3: The average gerrymandering power in terms of seats won of the parties versus $\phi$. The urban/rural party is in blue/red, and a darker colour represents a higher vote share of the gerrymandering party.

For most combinations of $\alpha^U$ and $\phi$, our approach was able to secure more districts for the gerrymandering party than its proportional vote share, resulting in a positive gerrymandering power. The one exception (which we will elaborate on later) was the case of highly unbalanced vote shares with completely homogenous precincts. See Figure 4.2 and Figure 4.5 for examples of gerrymandered solutions our algorithm found. Furthermore, we also tested our algorithm on several examples (on the grid), where we knew the optimal gerrymandering outcomes for each party.

Overall we noticed no difference when comparing results that stem from the simple and tree-based starting points. In this chapter all our data will be presented with all $40$ runs from both starting points averaged together. For a comparison of outcomes between starting points see subsection 7.2.3.

### 4.2.3 Gerrymandering for the urban-rural divide

We now describe the results of our simulations, and explain several important trends based on three key figures[9]. Figure 4.3 shows the gerrymandering power of the two parties for different vote shares as a function of the urban-rural

[9]For all of the figures see subsection 7.2.4.

divide.

### Highly unbalanced elections

In unbalanced elections with vote share difference of at least $20\%$ (see Figure 4.4 and Figure 4.2), we see an expected trend. At $\phi = 0$, when the voters are spread uniformly at random, the party with a majority vote share holds a majority in most precincts despite the randomness in our generation process. This makes it trivial for the majority party to gerrymander to win almost all districts, but difficult for the minority party to gerrymander well. In fact, the minority party has a negative gerrymandering power, i.e., it wins less fraction of districts than its vote share despite gerrymandering.

However, as the voters of the minority party concentrate, this disparity reduces. The majority party sees a reduction in its gerrymandering power as it can no longer avoid forming districts where the minority party wins due to its concentrated voters. Similarly, the minority party finds it easy to gerrymander to win a larger number of districts. At the extreme, with $\phi = 10$, it is able to win almost half the districts despite being at a $20\%$ vote share disadvantage.

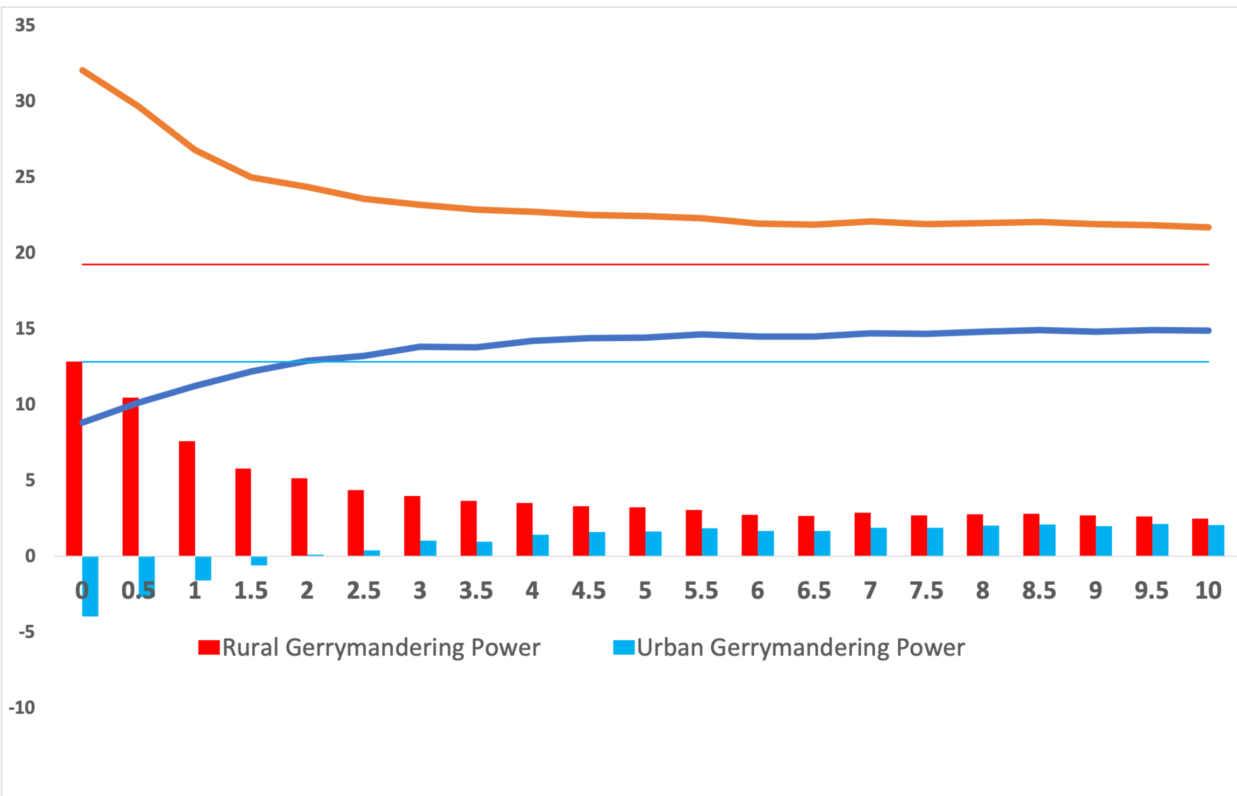


Figure 4.4: The number of districts won (top) and gerrymandering power (bottom) compared to $\phi$ for $\alpha^U = 0.4$. The horizontal blue/red line represents how many seats the urban/rural party would win in an entirely proportional system. The blue/brown line represents how many seats the urban/rural party won when they were gerrymandering. The blue/red bar graph at the bottom represents the gerrymandering power of the urban/rural party.

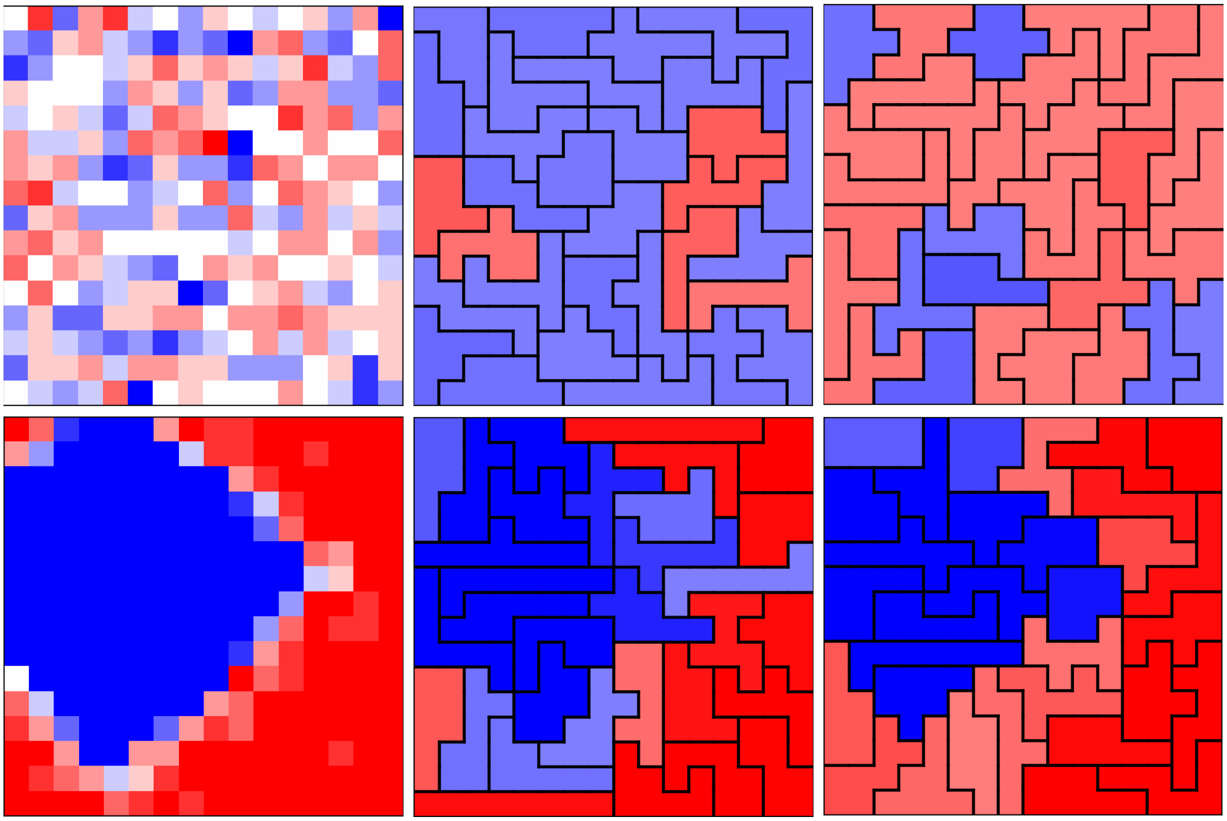

Figure 4.5: An example of the gerrymandered solutions $B^+$ found when $\alpha^U = 0.5$. Top row $\phi = 0$ bottom row $\phi = 10$. First column is partisan distribution of voters (similar to Figure 4.1). Second column is the gerrymandered solution for $U$, third column is the gerrymandered solution for $R$. Colour intensity within a district increases with the strength of victory.

### Close elections

Arguably, the more interesting elections in practice are the close elections with vote share difference of less than $20\%$. The trend is very different in these elections. For instance, consider Figure 4.5 and Figure 4.6 with $\alpha^U = 0.5$.

At $\phi = 0$, the voters are spread uniformly at random, which makes it easy for the gerrymandering party to put precincts with a slight majority together with precincts with a slight minority to form many districts with a slight majority, leading to a high gerrymandering power. Further, this holds for each party due to symmetry.

As the divide strengthens, the rural party witnesses a diminishing gerrymandering power as in the case of unbalanced elections. However, an interesting pattern emerges in the gerrymandering power of the urban party. As $\phi$ increases, we see that the gerrymandering power decreases suddenly till $\phi = 2$, then increases slowly, and finally levels out, forming a *trough*.

We do not believe this trough to be an artifact of our algorithm $B^+$. On a smaller number of instances, we ran the iterative algorithm $I$ for several hours to come up with hundreds of thousands of districting plans, and chose the most gerrymandered of them. A similar pattern emerges, though this approach returns less gerrymandered solutions than $B^+$, making the pattern a bit less emphasized.

We hypothesize that with a moderate $\phi$, there are still many urban voters in deeply rural regions, which constitutes a lot of wasted votes for the urban party as it is unable to put them together with other urban voters and form a district it can win. However, as the concentration further increases, these voters are brought closer to the urban centre, allowing the party to utilize their votes to win a few additional districts.

Due to a similar reason, when the rural party has a minority vote share (say $\alpha^R = 0.45$), we see an inverse pattern with its gerrymandering power initially increasing, and then slightly decreasing, thus forming a *peak*. Again, this is because with a moderate $\phi$, the urban party has a lot of wasted votes within rural regions, which helps the rural party gerrymander well.

#### Concentration leads to fairer districting

Interestingly, Figure 4.3 shows that across all vote shares, the gerrymandering power of both parties converges to $1/16$ (i.e., 2 more districts compared to the vote share with $K = 32$) as $\phi$ goes to 10. In fact, the convergence seems to begin at a fairly low concentration level (around $\phi = 2.5$). That is, at an extreme level of concentration, both parties are able to gerrymander and win about two more districts than their proportional share (dictated by their share of the votes).

Intuitively, at extreme concentration levels, there is a densely packed region of urban voters near the urban centre, a densely packed region of rural voters surrounding it, and a much sharper boundary in between (Figure 4.1). Irrespective of which party gerrymanders, districts near the urban centre are won by the urban party, and districts densely packed with rural voters are won by the rural party. This ensures each party approximately its proportional share of the districts. The only control that the gerrymandering party has is near the boundary, where it can merge its own voters with voters of the opponent, creating districts with a slight majority. This is reflected in the urban-gerrymandered and rural-gerrymandered districting shown in the final columns of Figure 4.2 and Figure 4.5. Since both parties control the same boundary region when gerrymandering, their gerrymandering power becomes identical with such extreme concentration. Further, since the number of vertices on the boundary is a fraction of the total number of vertices, this gerrymandering power is relatively small.

#### A rural advantage

Finally, we observe that the gerrymandering power is not symmetric between the urban and rural parties. The rural party almost always has a higher gerrymandering power than the urban party, even in the case of proportional vote split (Figure 4.6). This asymmetry is not surprising. As Figure 4.1 shows, the distribution of voters is also not symmetric; the urban party's voters congregate together in a tight area, while the rural party's voters surround them on all sides.

## 4.3 Gerrymandering optimally for real data

In this section we will continue to examine gerrymandering in various situations but now instead of examining a specific trend (such as the urban-rural divide) we will focus on gerrymandering on real data. The real election data we will use has been gathered and cleaned by the Metric Geometry and Gerrymandering Group (MGGG) [76] a multi university research group that focusses on geographic election data, and in particular gerrymandering in the United States.

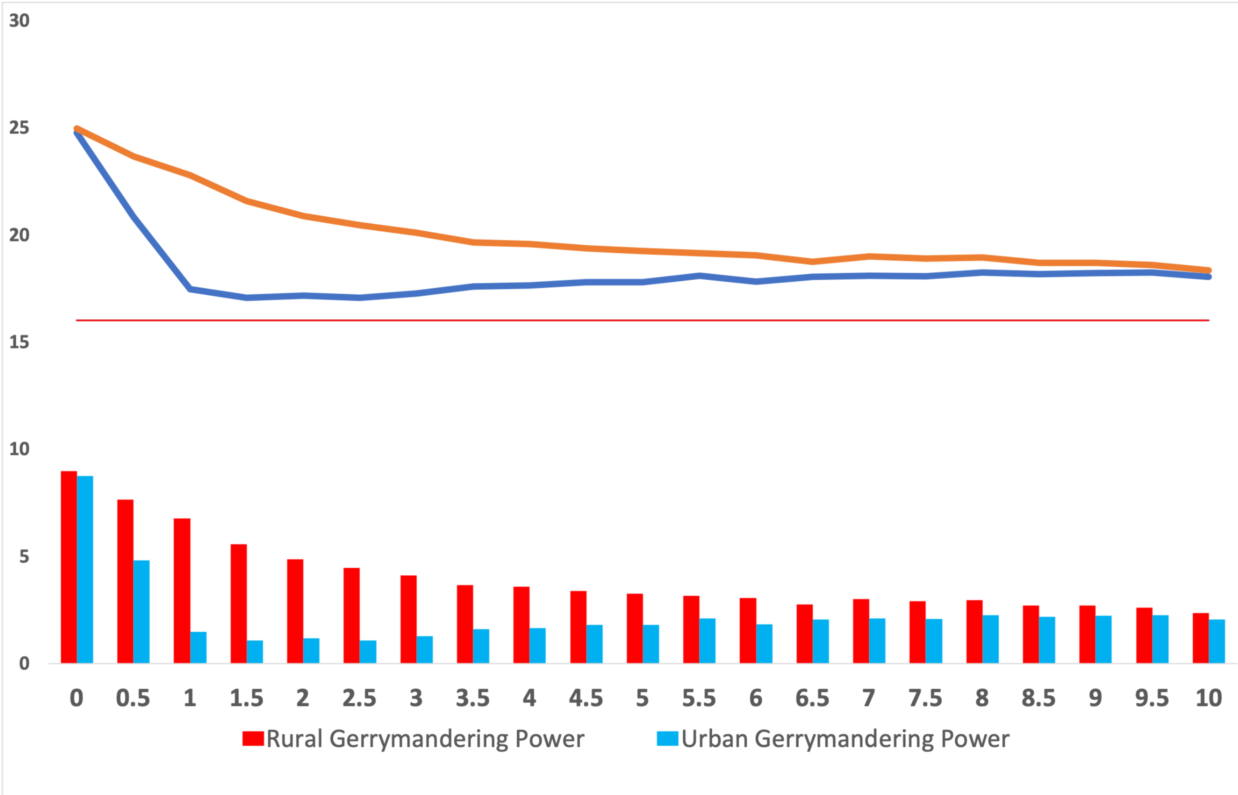


Figure 4.6: The number of districts won (top) and gerrymandering power (bottom) compared to $\phi$ for $\alpha^U = 0.5$. The horizontal blue/red (only one is visible since they overlap) line represents how many seats the urban/rural party would win in an entirely proportional system. The blue/brown line represents how many seats the urban/rural party won when they were gerrymandering. The blue/red bar graph at the bottom represents the gerrymandering power of the urban/rural party.

Since we are working with well-studied data it turns out we will have a (supposedly) optimal benchmark to compare against. In 2018 Nate Silver and his website 538[10] launched their atlas of redistricting as a part of their larger gerrymandering project [104]. Here Nate Silver and 538 claim that they have found the optimal way to gerrymander each of the states with more than one representative[11]. While we will describe in more detail their technique and what they consider optimal, we now note that it differs significantly from what we did in the previous section.

One final note, since we are working with real data we will drop the pseudonym of rural and urban parties and instead directly refer to the two as the Republican and Democratic parties respectively (furthermore we will use $N^D$ and $N^R$ to represent the number of Democratic and Republican voters).

### 4.3.1 A probabilistic model for gerrymandering

Since we wish to compare our optimization techniques against the results presented by 538 we will need to make sure we are comparing against the same metric. At a high level the goal for which 538 optimizes for differs significantly from what we presented in the previous section. Before we aimed to maximize the number of districts won by at least

[10]Perhaps best known for nearly perfectly predicting the 2008 and 2012 presidential elections.
[11]Of the 50 states which send representatives to the lower house of congress all but 7 send at least two members.

one vote, needless to say a district won by the slimmest of margins is not necessarily a win one would consider safe. Even a slight uniform variation across each district compared historic performances could be disastrous for a party. To mitigate this risk a partisan gerrymanderer should adopt margins of safety in their districts, they can even view the problem with a probabilistic lens.

We will now introduce the model 538 uses for their probabilistic gerrymandering, before we do so we note that much of this was not published in detail and had to be reconstructed from examination of their published results.

#### The Cook PVI

The basis of the 538 goal is the Cook partisan voting index (Cook PVI or PVI) published by the Cook Political Report [95] a non-partisan and independent newsletter that analyzes elections and trends in the United States. The PVI is a particular metric which measures how partisan a group of voters, in particular those who form a congressional district, are relative to the average voter in the United States.

To calculate the PVI we need a running value for how partisan the country is as a whole. To calculate this we take the number of votes garnered in the two most recent presidential elections[12] and see what fraction of these votes belong to the Democratic party[13]. The partisan skew expresses the average of the vote fractions for the Democrats in the last two presidential elections. Note, this is an average of averages, it is not weighted by the total votes in each election. This value, call it $\beta_D$, is approximately $0.51$ (see section 7.2.5 for exact values and calculations).

The PVI of a district is then just how partisan that district is relative to $\beta_D$. In a particular district if the total number of Democratic votes is $N_1^D$ and number of Republican votes is $N_1^R$ for the last presidential election and the total number of Democratic votes is $N_2^D$ and number of Republican votes is $N_2^R$ for the presidential election before that, then the PVI is :

$$100 \cdot \left(\frac{\frac{N_1^D}{N_1^D+N_1^R} + \frac{N_2^D}{N_2^D+N_2^R}}{2} - \beta_D\right) \tag{4.2}$$

Equation (4.2) can range from $-\beta_D \cdot 100$ for completely Republican dominated districts, to $(1 - \beta_D) \cdot 100$ for districts with only Democratic voters, or $0$ for districts which match the national average in the last two presidential elections.

Intuitively a district with a very positive PVI should be safely Democratic. Even if there is a uniform swing towards Republican sentiments this particular district should lean Democratic (the same is true for Republicans and districts with a very negative PVI).

Now a potential goal for a gerrymanderer is to draw districts which are safe wins, that is districts with very positive (or very negative) PVI values. But it is not clear what constitutes a safe win in terms of PVI values, in the next section we will explore what 538 considers safe.

#### The probability of a win

Ideally given the PVI value for some potential district we would be able to asses the probability that this district is a win for our target party. There are many reasonable ways one could go about this and unfortunately 538 did not publish their approach. The only reference for going from PVI to probability was in this quote:

---

[12] The presidential election is chosen since they use the same candidate for the entire country and thus are free of any local effects.

[13] We could also see what fraction belong to the Republican party, this would not change any of the subsequent results it would just switch the sign of all the outcomes.

"The probabilities of electing a Democrat or Republican are based on how often seats with a given Cook PVI elected members of each party between 2006 and 2016. They reflect a seat's expected performance over the long run, across a variety of political conditions. They are not predictions for the 2018 election, specifically."

After examination of the data from 538 we strongly believe their method is a sigmoid function which maps the PVI of a district to the probability the Democratic party wins the district. For the exact sigmoid form and how we found it see section 7.2.5. It is approximately of the form:

$$\sigma(x) = \frac{1}{1 + e^{-0.304 \cdot x}} \tag{4.3}$$

where $x$ is the PVI of a district and $e$ is the base of the natural logarithm. Briefly we note the range of the function is bounded strictly between $0$ and $1$ and the output is monotonically increasing in $x$ (for any real valued $x$). It is exactly equal to $0.5$ when the PVI is $0$. If $\sigma(x)$ is the probability the Democrats win a given district (with a PVI of $x$) then $1 - \sigma(x)$ is the probability the Republicans win that district.

#### The optimization goal

The goal of 538 (and the objective we are comparing against) is to maximize the number of districts where the probability of victory for the target party is above a certain threshold $\tau \in [0, 1]$. Intuitively $\tau$ is the threshold for a strong win, that is if the target party's probability of winning some district is above $\tau$ we consider that district safe for them. For their work, 538 goes with $\tau = 0.82$ roughly about a five-in-six chance of victory.

### 4.3.2 A technique for real data

Since we are now working with real data which is significantly more complex than the grid graphs in the previous section, in both geography and size, we will need a more complex optimization technique. The technique we used in the previous section is insufficient. The ILP step, while locally optimal[14], is woefully slow. In cases with hundreds of nodes being recombined it is slow to propose a solution, which makes it inadequate for our real data with thousands of nodes per district.

Instead we will use a similar but faster method that sacrifices locally optimal moves for speed. Our basic approach is built around simulated annealing (SA), a general optimization technique that has found much success in various discrete optimization problems. To build our SA approach we will expand on the Markov Chain Monte Carlo (MCMC) package for redistricting known as *Gerrychain* provided by MGGG. Details of what they provide and our extensions will be expanded on in subsequent sections.

With hill climbing based approaches at every step we have a candidate solution and propose some neighbouring solution, and we must decide if we should move to this neighbour or not[15]. In standard hill climbing methods, like (greedy) local search, only moves which improve the solution are accepted, thus each move must be an improvement step.

At a high level simulated annealing based optimization is essentially a hill climb, but moves are allowed towards inferior solutions. The ability to accept non-improvement steps, that is objectively worse solutions, become increasingly less permissible as the optimization proceeds. The logic with non-improvement steps is that they allow the

[14]By this we mean if it recombines the nodes of $k$ adjacent districts into $k$ new districts it finds the best way to make use of those nodes in any $k$ districts.

[15]The reader will notice that our previous method is effectively a hill climbing procedures where the ILP guides the neighbour proposal.

procedure to escape local optima earlier on in the process. If the space of solutions is especially non-convex (with respect to solution quality) these local optima can act as sinks for procedures which only allow improvement steps. The ability to accept a non-improvement neighbour is controlled by two parameters, the temperature of the system[16] and the difference in quality (also known as energy-difference) of the current and proposed solution.

### Simulated annealing components

Simulated annealing is not the simplest of optimization methods. Here we will introduce its components and how we fit them into our needs.

**Energy:** The first component of a SA based approach is the *energy* of a solution. The energy of a solution is a function which maps a potential solution to a numeric measure of quality, for our work we consider the set of solutions to be all legal districting plans. That is if a graph of a state has node set with $n$ nodes and they must be partitioned into $K$ districts the energy function is:

$$E : [n]^K \to \mathbb{R}_+. \tag{4.4}$$

It is standard for lower energy values to correspond to superior solutions and for zero energy to be the best any solution can take on[17]. For our purposes we will have our energy function be based on the expected number of districts won with one slight modification. Say we are gerrymandering for the Democratic party, if a potential solution $S$ is comprised of $K$ districts called $S_1, \cdots, S_K$ then the energy of that solution is :

$$E(S) = K - \sum_i v_D(S_i), \tag{4.5}$$

where $v_D(S_i)$ is equal to :

$$v_D(S_i) = \begin{cases} \sigma(S_i) & \sigma(S_i) \leq \tau \\ 1 & o.w. \end{cases}$$

If the target party is Republican party we can replace $v_D(S_i)$ with $v_R(S_i)$ [18] which is defined as follows:

$$v_R(S_i) = \begin{cases} 1 - \sigma(S_i) & 1 - \sigma(S_i) \leq \tau \\ 1 & o.w. \end{cases}$$

Recall $\tau$ is the threshold for what we consider a strong win. That is if a district win probability for our target party is above $\tau$, we say that is a safe win for that party. Intuitively our function is aiming to maximize the number of safe wins for the target party. Our method would prefer a solution with several borderline safe wins over a solution with fewer very safe wins. For all of our simulations we copy 538 and use $\tau = 0.82$.

**Proposal:** The second component of the SA based approach is the *proposal* function. A proposal function $P$ takes in a potential solution $S$ and picks a *neighbour* $S'$ of $S$:

[16]Here annealing is an analogy for metallurgical annealing where a metal is rapidly heated and cooled which allows the molecules to initially reshuffle and form new structures.

[17]The optimal solution for a particular instance could have non-zero energy, zero just serves as a lower bound.

[18]Recall, the probability the Republican party wins a district is just one minus the probability the Democratic party wins it.

$$P : [n]^K \to [n]^K. \tag{4.6}$$

There is no fixed definition of what a neighbour is and this can vary from domain to domain, or even within a problem itself. For our work we will use the following recom-proposal function, which was first suggested by MGGG. The recom-proposal is presented in the following algorithm 1:

**Algorithm 1** recom_proposal(G, S, j):

1: Let $S$ be the current solution (districting plan which partitions the vertices into $k$ components).
2: Pick $i \in \{2, \cdots, j\}$ random, but connected, districts from $S$ (where $j \leq k$).
3: Let $R$ denote the precinct nodes in these $i$ districts.
4: **for** $t \in \{1, \cdots, i-1\}$ **do**
5:     Draw a random spanning tree using using only the nodes of $R$. Call this spanning tree $T_t$.
6:     Sample a random edge $e$ (that has yet to be picked) in $T_t$. This divides $T_t$ into 2 connected components.
7:     If $T_t$ beneath $e$ forms a valid district: Make it one of our new districts, remove these nodes from $R$.
8:     Else, if there are edges yet to be sampled and $T_t$ beneath $e$ is not a valid district: Repeat step 6.
9:     If all of the edges of $T_t$ have been sampled and no valid district was ever found in $T_t$: repeat step 5.
10: **end for**
11: Let the remaining nodes of $R$ be the final new district.
12: Let $S'$ be the solution identical to $S$ but where $R$ has been redistricted according to steps $4 - 11$.
13: **if** $S'$ is a valid solution **then**
14:     Return $S'$
15: **else**
16:     Retry the algorithm from step 1.
17: **end if**

When we say a solution or district is valid we mean that it both satisfies contiguity and is population balanced. It should also be noted when using Kruskal's algorithm for drawing spanning trees the drawing of $T_t$ can effectively be done in time linear in the number of nodes left in $R$[19].

Each time through the for loop at step 6 the algorithm may pick several edges in the spanning tree $T_t$ (steps 6 through 8). For each iteration of the loop the first such $e$ must be connected to a leaf of $T_t$ (the first time step 6 is executed for each loop iteration). If the node under $e$ is not a valid district then another random edge is chosen (step 8). This next edge is either another edge connected to a leaf, or the edge directly above $e$ in $T_t$. In subsequent steps (if they are required) the algorithm picks an edge $e$ that has not been previously selected in $T_t$. In addition there are two more restrictions on selecting the new $e$. Firstly, the algorithm requires this edge is either connected to a leaf node or is the direct ancestor of a previously chosen edge. Secondly, the algorithm requires all of the edges under $e$ in $T_t$ have previously been selected. Intuitively, this process works by bubbling up through the various branches for $T_t$, trying edges until a sufficient one is found.

Each time through the for loop the algorithm should find one of the districts we need. It is possible some iteration of the for loop will fail to find a valid division, sampling every edge in $T_t$ (step 9). In this case this iteration of the for loop restarts, finding a new spanning tree, but keeps the districts found up to this point.

It is also possible that at some iteration of the for loop no spanning tree can lead to a valid districting. That is, it will just draw new spanning trees forever. We are unaware of any method that can detect this scenario, short of sampling every spanning tree. As a heuristic solution we put a time limit on the algorithm. We found the algorithm

[19] To draw a random spanning tree we find the minimum spanning tree after randomly assigning each edge a weight.

tends to find solutions within 20 seconds for the most complex instances we work with. If after 1000 seconds we don't have a solution we restart the entire algorithm. This was an addition we made to the functionality provided by MGGG.

We note that the recombination method proposed by MGGG only worked for recombining two districts at a time, whereas we extended it to work for any number. In the step where we pick $i$ random districts for recombination we do so by sampling uniformly at random from the set of all sets of connected districts up to size $j$. The intention of the MGGG method seems to be the same (for $j = 2$), but their code shows that they pick districts by uniformly sampling from all edges which cross district boundaries. This will favour picking pairs of districts which share large boundaries (in terms of nodes).

**Temperature:** The third part of the SA approach is the *temperature*, which acts as a control for how likely negative moves are at a given state of time. Generally the temperature is a decreasing function of the number of steps taken so far in the optimization. While there are many temperature functions (see [85] for a comparison of various methods) and choosing the ideal one is somewhat of a black-box in optimization, we've found the following temperature function works well (here $(s)$ is the step counter):

$$T(s) = 10000 \cdot (0.99^s) \tag{4.7}$$

This is known as the exponential cooling schedule. From the initial temperature of $10,000$ at every step we retain 99-percent of the remaining heat until we eventually cool to a temperature of $0$.

### The simulated annealing method

The simulated annealing method is as follows for a graph $G = (V, E)$ which is to be partitioned into $K$ districts:

**Algorithm 2** simulated_annealing_for_gerrymandering(G):

Let $S_0 = recom\_proposal(G, None, K)$.
$i = 0$
**while** $i \leq s_{max}$ **do**
    $S' = recom\_proposal(G, S_i, j)$
    **if** $E(S_i) \geq E(S')$ **then**
        $S_{i+1} = S'$
        $i = i + 1$
    **else**
        Let $\Delta E = E(S_i) - E(S')$
        Let $r$ be a value drawn uniformly at random from $[0, 1]$.
        **if** $\exp \frac{\Delta E}{T(i)} \geq r$ **then**
            $S_{i+1} = S'$
            $i = i + 1$
        **end if**
    **end if**
**end while**

In the first step $None$ refers to the districting which makes no assignments. To find the initial partition we do not need to provide the sub-routine with a valid districting since we are recombining all of the nodes. Intuitively the algorithm will always move to a lower energy solution and will move to a higher energy solution with high probability if the increase in energy is not too high and the temperature is not too cool. By tuning our method we found the process had stopped making improvement steps around $1000$ steps and we could reasonably recombine up to $4$ districts quickly. Thus we set $s_{max} = 10,000$ and $j = 4$.

It is entirely possible that the procedure will eventually be caught in a local optimum (or even the global one) it cannot move away from with reasonable probability. This is especially true later on as the temperature cools. If this is the case the main loop will, with very high probability, make no progress to completion. Because of this we often set a hard time limit and cut off the procedure after this point. In general with SA, or any random algorithm, one needs to run many parallel executions of the procedure, and each of these will iterate over many potential solutions. The best of all iterated solutions will be chosen as the returned solution.

**Running the simulations**

Our goal was to produce gerrymandered solutions, where the distance from the ideal population size ($\delta$) was at most half a percent, $\delta = 0.005$, or each district being within a percent of population of one another. Using this $\delta$ we did encounter some sub-optimal solutions. We suspect this is because the validity check in Algorithm 2 would reject many more solutions restricting the neighbourhood of each potential solution. With these restricted neighbourhoods there are far fewer paths in the search space to the optimal solution. To mitigate this for each state and gerrymandering target we optimized for two $\delta$ values, $\delta = 0.005$ and $\delta = 0.02$. For the looser population bounds, after finding the optimal solution, we tightened the population parameter by running an iterative process which shifted one node at a time on the boundary of one district to another district. This process would shift nodes that both improved the average distance from the ideal population and did not reduce the number of strong wins for the target party. If we can tighten the population bound into the ideal range of $\delta = 0.005$ we keep the solution, otherwise we do not use it. To reiterate, the hyper-parameters and settings of our process were:

- Each chain took a maximum of $10,000$ steps.
- In Algorithm 2 we recombined up to $4$ adjacent districts at once.
- Our initial temperature was $10,000$.
- We used an exponential cooling schedule which sheds $1\%$ of the current temperature at each step.
- Our cutoff for strong wins was $\tau = 0.82$, same as 538.

For each combination of state, party and $\delta$ we ran 50 chains for 20 hours on a AMD Ryzen Threadripper 2990WX with 128 GB of RAM which supports up to 64 threads in parallel[20]. Any solution for the looser $\delta$ which outperformed the tighter $\delta$ in terms of strong wins was further refined for population on $5$ threads in parallel using the process we described above. The results of our optimization (and 538's) can be seen in Table 4.3.

[20]Because this is a shared resource we did not take all the virtual threads at any given time.

### 4.3.3 The 538 optimal and real data

| State | Democrats | Republicans |
|---|---|---|
| **MA** | 0% | 0% |
| **MD** | 36% | 24% |
| **NC** | 0.3% | 0.1% |
| **PA** | 0.3% | 0.3% |
| **WI** | 0% | 0% |

Table 4.1: The fraction of the vote missing for each party in different states, for the 2012 and 2016 presidential elections.

In the 538 atlas of redistricting there was no description of the technique they used to optimize for each party, apart from stating that their districts were drawn by hand[21]. While we don't see any guarantees that their districting is optimal from a partisan perspective, they have achieved a more gerrymandered outcome for real data than any other work we're aware of. That being said introducing humans into the loop of any process generally slows it down. Using their method for studying questions like the urban-rural divide and its impact on optimal gerrymandering would be impossible. Therefore we wish to take our automated method and compare it against the 538 optimal. By showing our automated method is comparable to their handcrafted one we can use it method to study interesting questions.

Unfortunately, 538 did not publish their districting at a precinct level, instead they published the geographic boundaries of each district in the shape file format. Luckily, MGGG provides a tool called MAUP which can be used to map large geographic regions (the districts) onto smaller ones (the precincts). Using this tool we were able to build a pipeline which maps the 538 districts to our precinct level graphs, assigning each node its district. This isn't perfect, the 538 shape file seems to be somewhat coarse with overlapping boundaries of districts. Furthermore, the non-overlapping boundaries can divide individual nodes in our precinct data[22]. As a result some of the 538 outcomes were technically invalid on our graphs since they violate contiguity. We should also note, 538 did have a more rigorous population bound. They took "as equal as possible" to mean a $1000$ person difference between the largest and smallest district. For our plans we used a $1\%$ deviation limit (which seems to be an accepted limit [65]).

Because of the sparsity of available data we are unable to compare all $43$ states 538 gerrymandered, from our research we found publicly available data for $5$ states[23]. All of the data we used was gathered and made available by the researches at MGGG. We had a few minor issues. The 538 model was built with absentee voter data[24] that we could not find[25], and so their optimal outcome may not necessarily be optimal for our data. For most of the states this was not an issue as we have almost all of the vote (see Table 4.1) and this did not seem to have a large impact on the outcome (see Table 4.3) for the 538 optimal solution between data sets. The one exception was Maryland which seems to have extensive absentee voting that highly favours the Democrats. In this case the 538 optimal solution for the democrats differed fairly significantly. Furthermore, the graphs for these states were not always connected, these separate components were the result of natural water boundaries. To reconnect the components we added in as few edges as possible, and always between the physically closest precincts, this means to reconnect a graph with $k$

[21] They did use some software to help visualize and evaluate their hand drawn districts.
[22] In this case, the district which contains the most of the node's volume is assigned that node.
[23] Recall we require the vote breakdown for the previous two presidential elections aggregated at a prescient level.
[24] This can include early voters and those away from the state who vote by mail.
[25] For some states some of the absentee data seems to be included, it's not clear from our data which states include what absentee data.

components we added in $k-1$ edges[26].

As we mentioned real data is considerably more complex than the grid graphs we considered in the previous section. Recall, the grid graphs were $16 \times 16$ graphs with the exact same number of voters in each node. Furthermore, each person was a voter, thus when balancing the population constraint with packing and cracking voters a certain degree of complexity was lost. On the other hand the data we now work with is significantly more complex, for a summary of this information see Table 4.2.

| | Number of nodes | Total population | Standard deviation of population |
|---|---|---|---|
| MD | 1809 | 5774552 | 1711.99 |
| MA | 2151 | 6728169 | 819.77 |
| NC | 2692 | 9535483 | 2459.18 |
| PA | 9255 | 12702379 | 926.21 |
| WI | 6634 | 5686986 | 627.57 |

Table 4.2: The table shows the a summary of the data for each state. The first column is the total number of nodes in the induced graph. The second is the population of the state (in the 2010 census). The third column is the standard deviation of the population (each data point is the population of a node in the graph).

### 4.3.4 The optimal outcomes

| State | Total seats | Our D | 538 D (their data) | 538 D (our data) | Our R | 538 R (their data) | 538 R (our data) |
|---|---|---|---|---|---|---|---|
| MD | 8 | 7 | 8 | 5 | 4 | 4 | 4 |
| MA | 9 | 9 | 9 | 9 | 0 | 0 | 0 |
| NC | 13 | 7 | 8 | 8 | 11 | 10 | 10 |
| PA | 18 | 8 | 9 | 8 | 13 | 13 | 13 |
| WI | 8 | 5 | 5 | 5 | 6 | 6 | 6 |

Table 4.3: The number of strong wins in each state (where $\tau = 0.82$), for various objectives and data sources. Each row is a state. The first column is the total number of seats up for grabs. The second, third and fourth columns are what our algorithm found (on our data), what the 538 optimal was (using all the vote data) and what the same 538 solution was for our data, respectively, when gerrymandering for the Democrats was the goal. The fifth, sixth and seventh is the same but now when gerrymandering for the Republicans.

In this section we will analyze our algorithm's performance and compare it against the 538 optimal. We will discuss general trends and focus on a few specific states. For the complete set of all of our maps (and the 538 ones) see subsection 7.2.6. As Table 4.3 shows, our results are mostly in line with the optimal results found by 538 (when using our data). In general the algorithm performs quite well and in all but two case matches 538. The first exception is North Carolina when gerrymandering for the Democrats, here our optimal solution found 1 fewer strong wins. Maryland was also an exception. Our results for Democrats were far better than what 538 found, but we suspect this is due to the high number of absentee ballots we did not have access to. Indeed, in Maryland both of these results, our gerrymandering

[26] It's not clear what method 538 used to reconnect the graph components, but from our examination of their produced districting it seems they considered any two nodes bordering the same body of water as connected.

the 538 one applied to our data, did not match the 538 one applied to the full data set (which contains many more Democratic than Republican votes).

Figure 4.7 shows an example of the outcome our algorithm found (along with the one 538 presented) for the state of Wisconsin. While our plans aren't exactly equal to the 538 ones they are fairly similar. In Figure 4.7 we can see for either plan the Democrats only strongly win districts in the southern half of the state (where the major urban centres are). There are some differences, for our algorithm one of the Democratic safe wins stretches into the north of the state, but it still largely corresponds to a district in the 538 plan. In general our plans have more jagged boundaries than the 538 ones. While 538 did not make smoothness of the boundary an explicit goal for their partisan plans, they were able to find optimally partisan outcomes without requiring it. On the other hand our algorithm does not care about the shape of the boundaries, only that they lead to highly partisan outcomes. We suspect if we were to introduce some measure of smoothness we could "clean up" our boundaries with post-processing like we did for the population constraint.

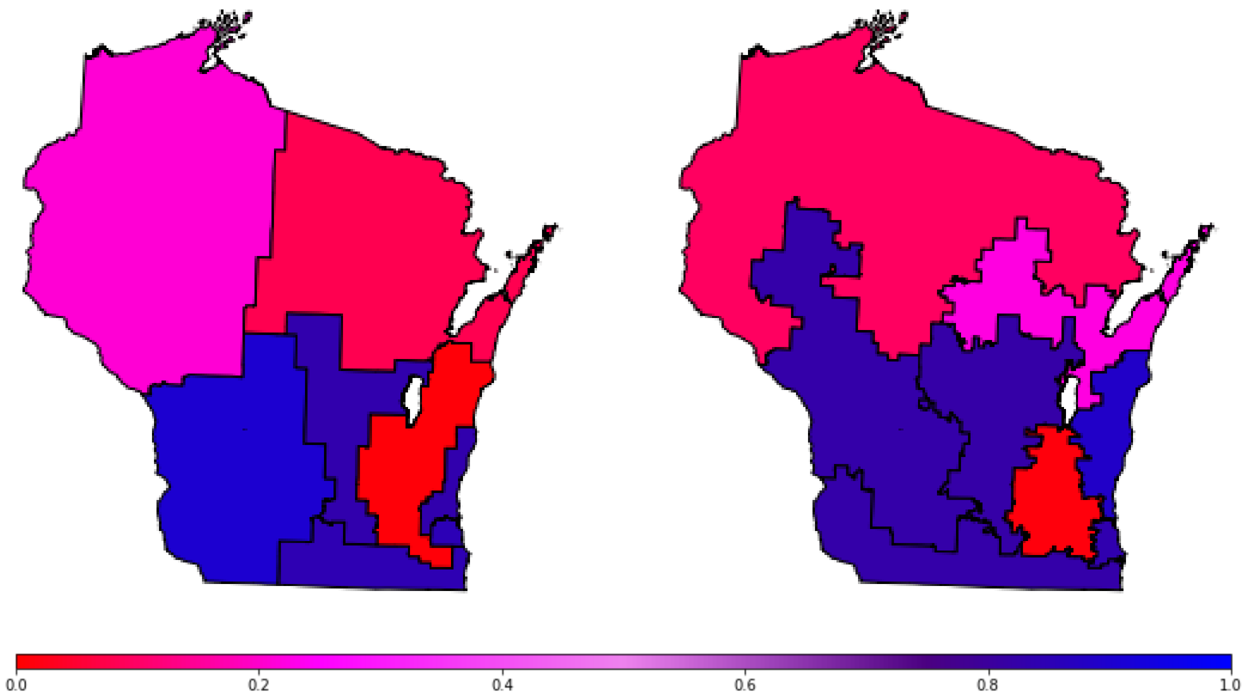


Figure 4.7: The left figure is how 538 optimally gerrymandered for the Democrats, the right figure is how our algorithm gerrymandered for the Democrats in Wisconsin. The bottom bar shows the scale for the probability the Democrats win the district.

### 4.3.5 The (almost) optimal outcomes

As we mentioned in Table 4.3, in North Carolina when gerrymandering for the Democrats our algorithm was one district away from being optimal. As Figure 4.8 shows, our outcome differs more significantly from the ones seen in Figure 4.7. The main issue is in the western and southern part of the state we don't have strong wins where 538 has them. Instead, in the middle of the state we have a strong win for the Democrats where 538 instead has a strong win for the Republicans. This middle district where we differ from 538 seems to act as a sort of a sink that our optimization objective (Equation 4.5) is drawn to. It's possible with more parallel runs and a higher initial temperature the method may be able to avoid this point, but for now it is the only (minor) blemish in an otherwise optimal method.

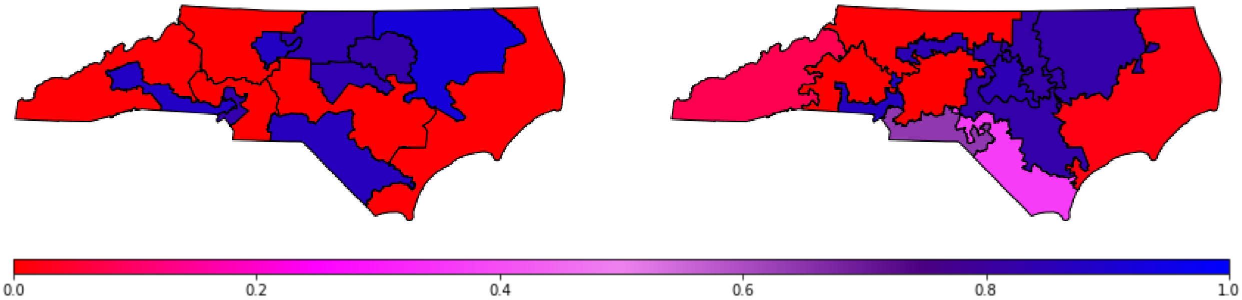


Figure 4.8: The left figure is how 538 optimally gerrymandered for the Democrats, the right figure is how our algorithm gerrymandered for the Democrats in North Carolina. The bottom bar shows the scale for the probability the Democrats win the district.

## 4.4 Conclusion

In this chapter we examined how the agents who draw political boundaries can influence the outcome of elections. This process, gerrymandering, can greatly distort election outcomes away from what many would consider fair. In particular we've focused our research on how automated methods can greatly amplify the power of a minority of voters.

Initially (in section 4.2) we studied the urban-rural divide, an ongoing trend where supporters of one party congregate in urban centres while the other party's supporters spread out in the surrounding rural area. We looked at this trend in the restricted setting of grid graphs. At first we provided some theoretical bounds, both best and worst case, on the ability to gerrymander in the very simplified settings of two districts and infinite graphs. To better understand more realistic voter distributions, like the urban-rural divide, we moved beyond the best and worst case and instead looked at various simulated distributions. To better study these simulated data points we developed a hill-climbing optimization approach based around an optimal integer linear program. We found that as the partisan concentration increased the ability of either party to gerrymander greatly diminished, and in the end they could only make gains in regions on the border of the urban and rural sections.

Next (section 4.3) we studied techniques to gerrymander on real data. This significantly more complex setting required a more advanced approach based around simulated annealing. While this part of our research is still in its infancy we have developed a near optimal technique that is almost always able to match districts which were hand crafted by experts at 538. In the near future we will refine this technique further and hope to apply it to various questions with respect to complex and real data.

# Chapter 5

# Primaries

> If I could not go to heaven but with a party, I would not go there at all.
>
> Thomas Jefferson

In this chapter we will focus on the mechanism political parties use to pick the candidates they run in elections, this process of filtering out potential candidates is often called the primary system. In a primary system party members first hold an internal vote over a set of potential party candidates, the winner of this internal election goes on to represent the party at the general level. At the general level all of the voters, those who partook in a party primary and those who abstained from voting in the primary, hold another election over all of the candidates at the general level. The general level contains all of the candidates who won the primary they competed in and those who enter the general without going through a primary themselves.

From the perspective of a party holding the primary the mechanism they use to pick their representative serves two goals, their representative should align with the party politically and have some degree of electability in a general election. Often these goals can be conflicting, and how a party should balance them is not easily answerable. That being said this is not the primary focus of this chapter (pun intended), despite it being an interesting question worth further investigation, instead we will focus on the utility to the voters who partake in the primary and general elections.

Using an embedding of voters and candidates into a metric space, which allows us to compare each candidate's social utility in terms of its total distance to the voters, we compare the primary election systems with their analogous direct system (one where all candidates enter the general election using the same voting rule the primary uses). We make the evaluation objective formal using the notion of *distortion* as advocated by a recent line of research [90, 16].

From a theoretical perspective we provide various bounds which show how far the primary and direct systems can differ in terms of distortion in the worst case. We then turn to simulations to see how the various parameters of the election can influence the quality of the outcome and again compare which is superior, primary or direct in terms of voter utility.

## 5.1 Model

For $k \in \mathbb{N}$, define $[k] = \{1, \ldots, k\}$. Let $V = [n]$ denote a set of $n$ *voters*, and $A$ denote a set of $m$ *candidates*. We assume that voters and candidates lie in an underlying metric space $\mathcal{M} = (S, d)$, where $S$ is a set of points and $d$ is a distance function satisfying the triangle inequality and symmetry. More precisely, there exists an *embedding* $\rho : V \cup A \to S$ mapping each voter and candidate to a point in $S$. For a set $X \subseteq V \cup A$, we slightly abuse the notation and let $\rho(X) = \{\rho(x) : x \in X\}$. Also, for $x, x' \in V \cup A$, we often use $d(x, x')$ instead of $d(\rho(x), \rho(x'))$ for notational convenience.

In this work, we assume that voters and candidates additionally have an affiliation with a political party. Specifically, we begin with a setting of two parties, denoted $-1$ and $1$. For now, we shall assume that *every* voter and candidate is affiliated with exactly one of the two parties. In Section 5.6 we will show that our results continue to hold even when some of the voters are independent and unaffiliated with either party. The *party affiliation* function $\pi : V \cup A \to \{-1, 1\}$ maps each voter and candidate to the party they are affiliated with. For $p \in \{-1, 1\}$, let $V_p = \pi^{-1}(p) \cap V$, $A_p = \pi^{-1}(p) \cap A$, $n_p = |V_p|$, and $m_p = |A_p|$. We require $n_p \geq 1$ for each $p \in \{-1, 1\}$.

Collectively, an *instance* is the tuple $I = (V, A, M, \rho, \pi)$. Given $I$, our goal is to find a candidate $a \in A$ as the winner. The *social cost* of $a$ is its total distance to the voters, denoted $C^I(a) = \sum_{i \in V} d(i, a)$. For party $p \in \{-1, 1\}$, let $C_p^I(a) = \sum_{i \in V_p} d(i, a)$. Hence $C^I(a) = C_{-1}^I(a) + C_1^I(a)$. For $X \subseteq V$, we also use $C_X^I(a) = \sum_{i \in X} d(i, a)$. Given an instance $I$, we would like to choose a candidate $a_{\text{OPT}} \in \arg\min_{a \in A} C^I(a)$ that minimizes the social cost. We shall drop the instance from superscripts if it is clear from the context.

However, we do not observe the full instance. Specifically, we do not know the underlying metric $M$ or the embedding function $\rho$. Instead, each voter $i \in N$ submits a *vote*, which is a ranking (strict total order) $\succ_i$ over the candidates in $A$ induced by their distance to the voter. Specifically, for all $i \in N$ and $a, b \in A$, $a \succ_i b \Rightarrow d(i, a) \leq d(i, b)$. The voter is allowed to break ties between equidistant candidates arbitrarily. The *vote profile* $\overrightarrow{\succ}^I = (\succ_1, \ldots, \succ_n)$ is the collection of votes. Given an instance $I$, its corresponding *election* $E^I = (V, A, \overrightarrow{\succ}^I, \pi)$ contains all observable information. Note that requiring voters' preferences to be defined by the underlying metric space they are in does not constrain the possible preference sets – any preference order, for any number of voters and candidates can be expressed in a metric space with a large enough dimension.

In the families of instances that we consider, we fix the number of candidates $m$ and let the number of voters $n$ to be arbitrarily large. This choice is justified because in many typical elections (e.g., political ones), voters significantly outnumber candidates. Let $\mathcal{I}^{\alpha}_{m,\mathcal{M}}$ be the family of instances satisfying the following conditions:

- Each party has at least an $\alpha$ fraction of the voters affiliated with it, i.e., $n_p \geq \alpha \cdot n$ for each $p \in \{-1, 1\}$. Note that $\alpha \in [0, 0.5]$: $\alpha = 0.5$ is the strictest (exactly half of the voters are affiliated with each party), while $\alpha = 0$ imposes no conditions; in the latter case, we omit the superscript $\alpha$ in $\mathcal{I}^{\alpha}_{m,\mathcal{M}}$.

- The number of candidates is at most $m$.

In particular, we shall focus on a few cases of $\mathcal{M}$ :

- $\mathcal{M} = \star$: This allows $M$ to be any arbitrary metric space.

- $\mathcal{M} = \mathbb{R}^k$: The metric space should be $M = (\mathbb{R}^k, d)$, where $d$ is the standard Euclidean distance. In particular, $\mathbb{R} = \mathbb{R}^1$ denotes the real line.

- $\mathcal{M} = \text{sep-}\mathbb{R}^k$: This means the embedding $\rho$ must be such that $\rho(V_{-1} \cup A_{-1})$ and $\rho(V_1 \cup A_1)$ are linearly separable.[1] In this case we also take the metric to be $M = (\mathbb{R}^k, d)$ with $d$ as the standard Euclidean distance. In plain words, the voters and candidates affiliated with each party reside in a certain part of the metric space, separate from those affiliated with the other party. In the single dimension, this means there exists a threshold on the line such that voters and candidates affiliated with one party lie to the left of it, while those affiliated with the other party lie to the right. Note that that this choice of $\mathcal{M}$ restricts the embedding $\rho$ based on the party affiliation $\pi$.

These families of instances are related by the following relation. For all $k$,

$$\begin{array}{ccccc} & & \mathcal{I}^{\alpha}_{m,\mathbb{R}^k} & & \\ \mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}^k} & \begin{array}{c}\subset\\ \subset\end{array} & & \begin{array}{c}\subset\\ \subset\end{array} & \mathcal{I}^{\alpha}_{m,\mathbb{R}^{k+1}} \subset \mathcal{I}^{\alpha}_{m,\star} \\ & & \mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}^{k+1}} & & \end{array}$$

### 5.1.1 Distortion and affiliation-independent rules

Since distortion is a worst-case notion, we have that when $\mathcal{I} \subseteq \mathcal{I}'$, $\phi_{\mathcal{I}}(f) \leq \phi_{\mathcal{I}'}(f)$ for every voting rule $f$.

For all the voting rules we examine in this chapter we refer the reader back to subsection 2.2.1. Standard voting rules choose the winning candidate independently of party affiliations. These include rules such as plurality, Borda, and STV. More concretely we call a voting rule *affiliation-independent* if $f(E) = f(E')$ when elections $E$ and $E'$ differ only in their party affiliation functions. Since an affiliation-independent voting rule $f$ ignores party affiliations, we have $\phi_{\mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}^k}}(f) = \phi_{\mathcal{I}^{\alpha}_{m,\mathbb{R}^k}}(f)$. All of the above-mentioned rules, in addition to being affiliation-independent, share the property of being *unanimous*, i.e., they return candidate $a$ when $a$ is the top choice of all voters.

### 5.1.2 Stages and primaries

Given an affiliation-independent voting rule $f$, voting systems with primaries employ a specific process to choose the winner, essentially resulting in a different voting rule $\widehat{f}$ that operates on a given election $E = (V, A, \overrightarrow{\succ}, \pi)$ as follows:

1. First, it creates two *primary elections*: for $p \in \{-1, 1\}$, define $E_p = (V_p, A_p, \overrightarrow{\succ}_p, \pi_p)$, where $\overrightarrow{\succ}_p$ denotes the preferences of voters in $V_p$ over candidates in $A_p$, and $\pi_p : V_p \to \{p\}$ is a constant function.

2. Next, it computes the winning candidate in each primary election (primary winner) using rule $f$: for $p \in \{-1, 1\}$, let $a^*_p = f(E_p)$.

3. Finally, let $E_g = (V, \{a^*_{-1}, a^*_1\}, \overrightarrow{\succ}_g, \pi)$ be the *general election*, where $\overrightarrow{\succ}_g$ denotes the preferences of all voters over the two primary winners. The winning candidate is $\widehat{f}(E) = \text{maj}(E_g)$, the majority candidate, preferred by a majority of voters. This is what most voting rules become when dealing with only 2 candidates[2].

This setting resembles systems employed by the main US, Canadian and other countries' parties, in which a party's members vote on their party's candidates to select a winner of their primary. In other systems, such selection could be

[1] Two sets of points are linearly separable if their convex hulls are disjoint, or equivalently, if there exists a hyperplane that contains each set in a distinct open halfspace.

[2] If one of the parties has no affiliated candidates, then the primary winner of the other party becomes the overall winner. In a setting with more than 2 parties, or where each party nominates several candidates, the general election can use $f$ to determine the outcome (or use some other voting process).

a multi-stage process. While we assume for now that every voter votes in the general election, in Section 5.6 we show that our results hold as long as at least a constant fraction of the voters participate in the general election.

Given an affiliation-independent voting rule $f$, the goal of this chapter is to compare its performance under the *direct system*, in which $f$ is applied on the given election directly, to its performance under the *primary system*, in which $\widehat{f}$ is applied on the given election instead. Formally, given a family of instances $\mathcal{I}$ and an affiliation-independent voting rule $f$, we wish to compare $\phi_{\mathcal{I}}(f)$ and $\phi_{\mathcal{I}}(\widehat{f})$ (henceforth, the distortion of $f$ with respect to $\mathcal{I}$ under the direct and the primary systems, respectively).

## 5.2 Small primaries are terrible

As defined above, in a family of instances $\mathcal{I}^{\alpha}_{m,\mathcal{M}}$, we require that at least $\alpha$ fraction of voters be affiliated with each party, i.e., $n_p \geq \alpha n$ for each $p \in \{-1, 1\}$. In other words, each primary election must have at least $\alpha n$ voters.

We first show that when a primary election may have very few voters ($\alpha = 0$), every reasonable voting rule has an unbounded distortion in the primary system, even with respect to our most stringent family of instances $\mathcal{I}_{m,\text{sep-}\mathbb{R}}$.

**Theorem 17.** *For $m \geq 3$, $\phi_{\mathcal{I}_{m,\text{sep-}\mathbb{R}}}(\widehat{f}) = \infty$ for every affiliation-independent unanimous voting rule $f$.*

*Proof.* Consider an instance $I \in \mathcal{I}_{m,\text{sep-}\mathbb{R}}$ in which voter 1 is located at 0 and affiliated with party $-1$, while the remaining $n-1$ voters are located at 1 and affiliated with party 1. All $m$ candidates are affiliated with party $-1$: one is at 0, and the rest are at 1.

The candidate $a^*$ at 0 becomes the primary winner of party $-1$, and trivially becomes the overall winner. Its social cost is $C(a^*) = n - 1$. In contrast, an optimal candidate $a_{OPT}$ at 1 has social cost $C(a_{OPT}) = 1$. Hence, $\frac{C(a^*)}{C(a_{OPT})} = n - 1$. Since the number of voters $n$ is unbounded, $\phi_{\mathcal{I}_{m,\text{sep-}\mathbb{R}}}(\widehat{f}) = \infty$. $\square$

Theorem 17 continues to hold even if we require that at least a constant fraction of candidates be affiliated with each party: we could simply move a constant fraction of the candidates from 1 to 3 and assign them to party 1, and the proof would still hold.

On the other hand, if we require that at least a *constant fraction of voters* be affiliated with each party, the result changes dramatically.

## 5.3 Large primaries are never much worse than direct elections

For every affiliation-independent voting rule $f$, we bound the distortion of $\widehat{f}$ in terms of the distortion of $f$ *for every instance*. Note that this is stronger than comparing the worst-case distortions of $f$ and $\widehat{f}$ over a family of instances.

Given an instance $I = (V, A, M, \rho, \pi)$ and party $p \in \{-1, 1\}$, we say that $I_p = (V_p, A_p, M, \rho_p, \pi_p)$ is the *primary instance* of party $p$, where $\rho_p$ and $\pi_p$ are restrictions of $\rho$ and $\pi$ to $V_p \cup A_p$. The primary election $E_p$ of party $p$ is precisely the election corresponding to instance $I_p$.

**Theorem 18.** *Let $I = (V, A, M, \rho, \pi)$ be an instance. For $p \in \{-1, 1\}$, let $I_p$ be the primary instance of party $p$, and $n_p = |V_p| \geq \alpha n$. Then,*

$$\phi(\widehat{f}, I) \leq 3 \cdot \frac{1 - \alpha + \max(\phi(f, I_{-1}), \phi(f, I_1))}{\alpha}.$$

*Further, for a socially optimal candidate $a_{OPT} \in \arg\min_{a \in A} C^I(a)$, we have a bound depending only on the distortion of the primary election of its party:*

$$\phi(\widehat{f}, I) \leq 3 \cdot \frac{1 - \alpha + \phi(f, I_{\pi(a_{OPT})})}{\alpha}.$$

For each family of instances $\mathcal{I}$ that we study, it holds that for every instance $I \in \mathcal{I}$, both its primary instances, if seen as direct elections, are also in $\mathcal{I}$ (since the party division has no effect on the direct election distortion). Hence, we can convert the instance-wise comparison to a worst-case comparison.

**Corollary 19.** *For every $\alpha > 0$, $k \in \mathbb{N}$, family of instances $\mathcal{I} \in \{\mathcal{I}_{m,\star}, \mathcal{I}_{m,\mathbb{R}^k}, \mathcal{I}_{m,\text{sep-}\mathbb{R}^k}\}$, and affiliation-independent voting rule $f$,*

$$\phi_{\mathcal{I}}(\widehat{f}) \leq 3 \cdot \frac{1 - \alpha + \phi_{\mathcal{I}}(f)}{\alpha}.$$

Since $\phi_{\mathcal{I}}(f) \geq 1$ by definition, we can write $\phi_{\mathcal{I}}(\widehat{f}) \leq \frac{6}{\alpha} \cdot \phi_{\mathcal{I}}(f)$. In other words, for every affiliation-independent voting rule $f$, its distortion under the primary system is at most a constant times bigger than its distortion under the direct system, with respect to every family of instances that we consider.

To prove Theorem 18, we need two lemmas. The first lemma moves the relation of two candidates, from being connected by the distortion of a rule $f$ in a primary instance $I_p$, to being connected by the general electorate costs. The intuition is that when voters in $V \setminus V_p$ drive up the social cost of $a_p^*$ (i.e., when they are far from $a_p^*$), they must do so for every candidate in $A_p$.

**Lemma 20.** *Let $a_p^*$ denote the primary winner of party $p$. Then*

$$C(a_p^*) \leq \frac{1 - \alpha + \phi(f, I_p)}{\alpha} \cdot \min_{a \in A_p} C(a).$$

*Proof.* Let $\theta = \phi(f, I_p)$ (hence $\theta \geq 1$). Fix an arbitrary $a \in A_p$. Then, $C_{V_p}(a_p^*) \leq \theta \cdot C_{V_p}(a)$. Now,

$$\begin{aligned} C(a_p^*) &= C_{V_p}(a_p^*) + C_{V \setminus V_p}(a_p^*) \\ &\leq \theta \cdot C_{V_p}(a) + C_{V \setminus V_p}(a) + (n - n_p) \cdot d(a, a_p^*) \\ &\leq \theta \cdot C(a) + (n - n_p) \cdot d(a_p^*, a), \end{aligned} \tag{5.1}$$

where the second transition follows due to the triangle inequality. We also have $d(a_p^*, a) \leq d(a_p^*, i) + d(i, a)$ for any $i \in V_p$. Thus,

$$\begin{aligned} d(a_p^*, a) &\leq \frac{C_{V_p}(a_p^*) + C_{V_p}(a)}{n_p} \\ &\leq \frac{1 + \theta}{n_p} \cdot C_{V_p}(a) \leq \frac{1 + \theta}{n_p} \cdot C(a). \end{aligned} \tag{5.2}$$

Substituting Equation (5.2) into Equation (5.1), and using $\frac{n - n_p}{n_p} \leq \frac{n(1-\alpha)}{\alpha n} = \frac{1-\alpha}{\alpha}$ we end up with

$$C(a_p^*) \leq \theta \cdot C(a) + (n - n_p) \cdot d(a_p^*, a) \leq \theta \cdot C(a) + (n - n_p) \cdot \frac{1+\theta}{n_p} \cdot C(a)$$
$$\leq (\theta + \frac{1-\alpha}{\alpha}(1+\theta))C(a) = \frac{1-\alpha+\theta}{\alpha} C(a)$$

□

Our next lemma shows that the primary winner that wins the general election is not much worse than the primary winner that loses the general election. We defer the proof of the lemma until Section 5.6, where we prove a more general result (Lemma 26).

**Lemma 21.** *Let $a_{-1}^*$ and $a_1^*$ be the two primary winners, $a^* \in \left\{a_{-1}^*, a_1^*\right\}$ be the winner of the general election, and $\widehat{a} \in \left\{a_{-1}^*, a_1^*\right\} \setminus \{a^*\}$. Then,*

$$C(a^*) \leq 3 \cdot C(\widehat{a}).$$

We are now ready to prove our main result.

*Proof of Theorem 18.* Recall that $a_{-1}^*$ and $a_1^*$ are the primary winners. Let $p \in \{-1, 1\}$ be such that $a_p^*$ is the winner of the general election. Let $a_{OPT} \in \arg\min_{a \in A} C(a)$ be a socially optimal candidate. We consider three cases.

- *Case 1: $a_{OPT} \in A_p$.* That is, the optimal candidate and the winner are affiliated with the same party. In this case, Lemma 20 yields

$$\phi(\widehat{f}, I) \leq \frac{1-\alpha+\phi(f, I_p)}{\alpha} = \frac{1-\alpha+\phi(f, I_{\pi(a_{OPT})})}{\alpha}.$$

- *Case 2: $a_{OPT} \in A_{-p}$, $a_{OPT} = a_{-p}^*$.* That is, the optimal candidate and the winner are affiliated with different parties, and the optimal candidate is a primary winner. In this case, Lemma 21 yields

$$\phi(\widehat{f}, I) \leq 3.$$

- *Case 3: $a_{OPT} \in A_{-p}$, $a_{OPT} \neq a_{-p}^*$.* That is, the optimal candidate is affiliated with a party different than that of the winner, and is not a primary winner. In this case, we use both Lemmas 20 and 21 to derive

$$\begin{aligned} \phi(\widehat{f}, I) &= \frac{C(a_p^*)}{C(a_{OPT})} = \frac{C(a_p^*)}{C(a_{-p}^*)} \cdot \frac{C(a_{-p}^*)}{C(a_{OPT})} \\ &\leq 3 \cdot \frac{1-\alpha+\phi(f, I_{-p})}{\alpha} \\ &= 3 \cdot \frac{1-\alpha+\phi(f, I_{\pi(a_{OPT})})}{\alpha}. \end{aligned}$$

Thus, in each case, we have the desired approximation. □

Notice that we do not use the assumption that both parties use the same voting rule $f$ in their primaries. The theorem extends easily to allow the use of different voting rules, with the distortion under the primary system still being bounded in terms of the maximum of the distortions in the two primary elections. Additionally, we show in

Section 5.6 that two other assumptions we made so far — every voter participates in one of the two primaries and every voter participates in the general election — can be relaxed without significantly affecting our results.

## 5.4 Large primaries are not better without party separability

In the previous section we showed that a voting rule does not perform much worse under the primary system than under the direct system. Now we show that it does not perform any better either, at least in the worst case over all instances with at most $m$ candidates. The result continues to hold even if we require each party to have at least a constant fraction of the voters.

Note that this result is weaker than Theorem 18 because it is a worst-case comparison instead of an instance-wise comparison. However, it still applies to all voting rules $f$. It applies to any metric that does not require separability of parties, in particular to $\mathcal{I}_{m,\star}$ and $\mathcal{I}_{m,\mathbb{R}^k}$.

**Theorem 22.** *For $\alpha \in [0, 0.5]$, $k \in \mathbb{N}$, $\mathcal{M} \in \{\star, \mathbb{R}^k\}$ (i.e. when the metric space does not require party separability), and affiliation-independent voting rule $f$, we have $\phi_{\mathcal{I}^\alpha_{m,\mathcal{M}}}(\widehat{f}) \geq \phi_{\mathcal{I}^\alpha_{m,\mathcal{M}}}(f)$.*

*Proof.* We shall denote $\mathcal{I}^\alpha_{m,\mathcal{M}}$ as $\mathcal{I}$. We want to show that for every instance $I \in \mathcal{I}$, there exists an instance $I' \in \mathcal{I}$ such that $\phi(\widehat{f}, I') \geq \phi(f, I)$. This would imply the desired result.

Fix an instance $I = (V, A, M, \rho, \pi) \in \mathcal{I}$. Let $a_{OPT} \in A$ denote an optimal candidate in $I$, and $a^* = f(E^I)$. Note that $\phi(f, I) = \frac{C^I(a^*)}{C^I(a_{OPT})}$. Construct instance $I' = (V', A, M, \rho', \pi')$ as follows.

- Let $V' = V \cup \widetilde{V}$, where $\widetilde{V}$ is a new set of voters and $|\widetilde{V}| = |V|$.
- Let $\rho'(x) = \rho(x)$ for all $x \in V \cup A$, and $\rho(x) = \rho(a_{OPT})$ for all $x \in \widetilde{V}$. That is, $\rho'$ matches $\rho$ for existing voters and candidates, and the new voters are co-located with $a_{OPT}$.
- Let $\pi'(x) = -1$ for all $x \in V \cup A$, and $\pi'(x) = 1$ for all $x \in \widetilde{V}$. That is, all existing voters and candidates are affiliated with party $-1$, while all new voters are affiliated with party 1.

First, let us check that $I' \in \mathcal{I}$. Since $I$ has $m$ candidates, so does $I'$. Further, in $I'$, we have $|V'_{-1}| = |V'_1| = |V'|/2$, which satisfies the constraint corresponding to every $\alpha \in [0, 0.5]$. Hence, we have $I' \in \mathcal{I}$.

Let us apply $\widehat{f}$ on $I'$. One of its primary instances, $I'_{-1}$, is precisely $I$. Hence, the primary winner of party $-1$ is $f(I'_{-1}) = f(I) = a^*$. Because there are no candidates affiliated with party 1, $a^*$ becomes the overall winner.[3]

Next, $C^{I'}(a^*) \geq C^I(a^*)$ because $V \subset V'$. Also, $C^{I'}(a_{OPT}) = C^I(a_{OPT})$ because $a_{OPT}$ has zero distance to all voters in $V' \setminus V$. Together, they yield

$$\phi(\widehat{f}, I') = \frac{C^{I'}(a^*)}{C^{I'}(a_{OPT})} \geq \frac{C^I(a^*)}{C^I(a_{OPT})} = \phi(f, I), \tag{5.3}$$

as desired. □

Our proof establishes a slightly stronger result than stated in the theorem: instead of showing $\phi_{\mathcal{I}^\alpha_{m,\mathcal{M}}}(\widehat{f}) \geq \phi_{\mathcal{I}^\alpha_{m,\mathcal{M}}}(f)$, we actually show $\phi_{\mathcal{I}^{0.5}_{m,\mathcal{M}}}(\widehat{f}) \geq \phi_{\mathcal{I}^\alpha_{m,\mathcal{M}}}(f)$.

[3] Even if we require each party to have at least one affiliated candidate, the proof essentially continues to hold. In this case, we can add one candidate affiliated with party 1 that is located sufficiently far from all the voters, ensuring that $a^*$ still becomes the overall winner. This would show $\phi_{\mathcal{I}^\alpha_{m+1,\mathcal{M}}}(\widehat{f}) \geq \phi_{\mathcal{I}^\alpha_{m,\mathcal{M}}}(f)$ because instance $I'$ may now have $m + 1$ candidates.

## 5.5 The advantages of party separability

The analysis for $\mathcal{I}_{m,\text{sep-}\mathbb{R}^k}$ is not as straightforward as in the previous section. In the proof of Theorem 22, we co-located the new voters affiliated with party 1 and $a_{OPT}$ affiliated with party $-1$. This was allowed because non-separable metrics like $\star$ and $\mathbb{R}^k$ place no constraints on the embedding.

With $\mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}^k}$, we need the voters and candidates affiliated with one party to be separated from those affiliated with the other. Hence, this operation of putting all of one party's voters at the location of $a_{OPT}$ belonging to another party would be allowed only if, in the original instance $I$, $a_{OPT}$ is on the boundary of the convex hull of $\rho(V \cup A)$. While this is not the case for all instances, we only need this in at least one *worst-case instance* for $f$, i.e., for at least one $I \in \mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}^k}$ with $\phi(f, I) = \phi_{\mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}^k}}(f)$. Equation (5.3) would then yield the desired result. More generally, it is sufficient if, given any $\epsilon > 0$, we can find an instance $I$ such that $\phi(f, I) \geq \phi_{\mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}^k}}(f) - \epsilon$ and $a_{OPT}$ is at distance at most $\epsilon$ from the boundary of the convex hull of $\rho(V \cup A)$.

Interestingly, [3] show that this is indeed the case for plurality and Borda voting rule (see the proof of their Theorem 4). Thus, we have the following.

**Proposition 23.** *Let $f$ be plurality or Borda. Then, $\phi_{\mathcal{I}^{0.5}_{m,\text{sep-}\mathbb{R}}}(\widehat{f}) \geq \phi_{\mathcal{I}_{m,\star}}(f)$.*

However, known worst cases for the Copeland rule [3] and STV [105] do not satisfy this requirement. It is unknown if these rules admit a different worst case that satisfies it.

This raises the question if Proposition 23 holds for *all* affiliation-independent voting rules. We shall shortly answer this *negatively*.

More precisely, we construct an affiliation-independent voting rule $f$ such that $\phi_{\mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}}}(\widehat{f}) \ll \phi_{\mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}}}(f)$ for every $\alpha > 0$. That is, with large primaries, $f$ performs much better under the primary system than under the direct system, when voters and candidates are embedded on a line and the separability condition is imposed.

Note that instances in $\mathcal{I}_{m,\text{sep-}\mathbb{R}}$ are highly structured. For instance, it is known that when voters and candidates are embedded on a line, there always exists a weak Condorcet winner [13], and selecting such a candidate results in a distortion of at most 3 [3]. Hence, we have $\phi_{\mathcal{I}_{m,\text{sep-}\mathbb{R}}}(f) = 3$ for every Condorcet-consistent, affiliation-independent voting rule $f$.[4]

Our aim in this section is to construct an affiliation-independent voting rule $f_{\text{fail}}$ that with respect to $\mathcal{I}_{m,\text{sep-}\mathbb{R}}$ has an unbounded distortion in the direct system, but at most a constant distortion in the primary system.

**Definition 12.** Let $f_{\text{fail}}$ be an affiliation-independent voting rule that operates on election $E = (V, A, \vec{\succ})$ as follows. Let $A = \{a_1, \ldots, a_m\}$, and $t = (m + 1)/2$.

- *Special Case*: If $m \geq 9$, $m$ is odd, $n \geq m^2$, and $\vec{\succ}$ has the following structure, then return $a_1$.

  1. For voter 1, $a_1 \succ_1 \ldots \succ_1 a_m$.
  2. For voter 2, $a_m \succ_2 \ldots \succ_2 a_1$.
  3. For voter 3, $a_{t-1}$ is the most preferred, and
     $a_{m-2} \succ_3 a_1 \succ_3 a_{m-1} \succ_3 a_m$.
  4. For voter 4, $a_{t+1}$ is the most preferred, and
     $a_3 \succ_4 a_m \succ_4 a_2 \succ_4 a_1$.

[4] [3] also proved that no affiliation-independent (deterministic) voting rule can have distortion better than 3, even with respect to $\mathcal{I}_{m,\text{sep-}\mathbb{R}}$.

5. For $j \in [m-2]$, for voter $i = 4 + (2j-1)$,
   $a_{j+1} \succ_i a_{j+2} \succ_i a_j$, and for voter $i' = 4 + 2j$,
   $a_{j+1} \succ_{i'} a_j \succ_{i'} a_{j+2}$.

6. For every other voter $v$, $a_t$ is the most preferred.

- If $E$ does not fall under the special case, then apply any Condorcet consistent voting rule (e.g., Copeland).

Note that $m$ being odd ensures that $t$ is an integer, and $m \geq 9$ ensures that $a_1$, $a_3$, $a_{t-1}$, $a_t$, $a_{t+1}$, $a_{m-2}$, and $a_m$ are all distinct candidates. The significance of $n \geq m^2$ will be clear later.

We will now establish that a worst-case instance of $f_{\text{fail}}$ falls under the special case; for this instance, we need to show that $a_t$ is socially optimal; that $f_{\text{fail}}$ returns $a_1$ on this instance; and most importantly, that the structure of $\overrightarrow{\succ}$ ensures that the optimal candidate $a_t$ is sufficiently far from both the leftmost and the rightmost candidates.

We prove this last fact in the following lemma.

**Lemma 24.** *Let $I = (V, A, M, \rho, \pi) \in \mathcal{I}_{m,\text{sep-}\mathbb{R}}$ be an instance for which the corresponding election $E^I$ falls under the special case of $f_{\text{fail}}$. Then the following holds.*

1. *Either $\rho(a_1) \leq \ldots \leq \rho(a_m)$, or $\rho(a_1) \geq \ldots \geq \rho(a_m)$, or $|\rho(A)| = 2$.*

2. *If $|\rho(A)| \neq 2$, $min\{d(a_t, a_1), d(a_t, a_m)\} \geq \frac{d(a_1,a_m)}{4}$.*

*Proof.* Since voter 1 ranks $a_m$ last and preferences are single peaked on the line, $a_m$ is at one edge of the candidate ordering. Similarly, since voter 2 ranks $a_1$ last, candidate $a_1$ is also at the edge of the candidate ordering (i.e., $\rho(a_1) = \max_{a \in A} \rho(a)$ or $\rho(a_1) = \min_{a \in A} \rho(a)$ and $\rho(a_m) = \max_{a \in A} \rho(a)$ or $\rho(a_m) = \min_{a \in A} \rho(a)$). If $\rho(a_1) = \rho(a_m)$, this means voters 1 and 2 are located in an equal distance from all candidates (which means all candidates are located in the same location, or some are at some distance from voters 1 and 2, and the rest are at the same distance in the other direction from these voters).

Assume $|\rho(A)| > 2$ (this also means $\rho(a_1) \neq \rho(a_m)$ and $\rho(v_1) \neq \rho(v_2)$), we wish to show the order of candidates is as voter 1 ordered them, i.e., $\rho(a_1) \leq \ldots \leq \rho(a_m)$ or $\rho(a_1) \geq \ldots \geq \rho(a_m)$. If voter 1 is further away from all candidates (i.e, if $\rho(a_1) = \max_{a \in A} \rho(a)$, $\rho(v_1) > \rho(a_1)$, and if $\rho(a_1) = \min_{a \in A} \rho(a)$, $\rho(v_1) < \rho(a_1)$), the ordering of the candidates is as voter 1 orders them. Otherwise, let $\ell$ be the smallest index such that $\rho(a_1) \neq \rho(a_\ell)$, then $\rho(v_1)$ may be between $\rho(a_1)$ and $\rho(a_\ell)$. If $d(v_1, a_1) < d(v_1, a_\ell)$, once again, the ordering of candidates is as voter 1 ordered them. If $d(v_1, a_1) = d(v_1, a_\ell)$, for any $\ell' > \ell$, $\rho(a_{\ell'}) \neq \rho(a_1)$, as that contradicts voter 2's vote ($a_{\ell'} \succ_2 a_\ell \succ_2 a_1$). Therefore, $\rho(\ell')$ is either at $\rho(\ell)$, or further away from $\rho(a_1)$, meaning that candidates locations are ordered in the order voter 1's ordered them.

For the second condition, we will show that $d(a_t, a_1) \geq \frac{d(a_1,a_m)}{4}$. By symmetry, we also obtain $d(a_t, a_m) \geq \frac{d(a_1,a_m)}{4}$. Assume $|\rho(A)| \neq 2$, and without loss of generality, let $\rho(a_1) \leq \ldots \leq \rho(a_m)$ from the first condition. We show that either $\rho(a_1) = \ldots = \rho(a_m)$ or $\rho(a_1) < \ldots < \rho(a_m)$. If not, then we can find three consecutive candidates $a_j$, $a_{j+1}$, and $a_{j+2}$ such that either $\rho(a_j) = \rho(a_{j+1}) < \rho(a_{j+2})$ or $\rho(a_j) < \rho(a_{j+1}) = \rho(a_{j+2})$. Both of these options are impossible in our case due to the existence of voters with preferences $a_{j+1} \succ a_j \succ a_{j+2}$ and $a_{j+1} \succ a_{j+2} \succ a_j$.

If $\rho(a_1) = \ldots = \rho(a_m)$ the second condition is trivially true. Suppose $\rho(a_1) < \ldots < \rho(a_m)$. Because voter 3 (resp. 4) prefers candidate $a_{t-1}$ (resp. $a_{t+1}$) the most, we have $\rho(v_3) \in (\rho(a_{t-2}), \rho(a_t))$ (resp. $\rho(v_4) \in$

$(\rho(a_t), \rho(a_{t+2}))$. We can now show

$$\begin{aligned}
d(a_1, a_m) &\leq d(a_1, v_4) + d(v_4, a_m) \\
&\leq 2d(a_1, v_4) \qquad (\because a_m \succ_4 a_1) \\
&\leq 2d(a_1, a_{m-2}) \quad (\because \rho(a_t) < \rho(v_4) < \rho(a_{m-2})) \\
&\leq 2(d(a_1, v_3) + d(v_3, a_{m-2})) \\
&\leq 4d(a_1, v_3) \qquad (\because a_{m-2} \succ_3 a_1) \\
&\leq 4d(a_1, a_t). \qquad (\because \rho(a_1) < \rho(v_3) < \rho(a_t))
\end{aligned}$$

This concludes the proof. □

**Theorem 25.** *For $m \geq 9$ and constant $\alpha \in (0, 0.5]$, $\phi_{\mathcal{I}^\alpha_{m,\text{sep-}\mathbb{R}}}(\widehat{f}_{\text{fail}})$ is upper bounded by a constant, whereas $\phi_{\mathcal{I}^\alpha_{m,\text{sep-}\mathbb{R}}}(f_{\text{fail}})$ is unbounded.*

*Proof.* First, we show that $\phi_{\mathcal{I}^\alpha_{m,\text{sep-}\mathbb{R}}}(f_{\text{fail}})$ is unbounded. Consider the following instance $I = (V, A, M, \rho, \pi)$. Let $V = \{v_1, \ldots, v_{2n}\}$ (where $n \geq m^2$), $A = \{a_1, \ldots, a_m\}$ ($m$ being odd), and $M = (\mathbb{R}, d)$ with $d$ being the Euclidean distance on the line.

The embedding function $\rho$ is as follows. For $\ell \in [m]$, $\rho(a_\ell) = \frac{\ell-1}{m-1}$; that is, candidates $a_1$ through $a_m$ are uniformly spaced in $[0, 1]$ with $\rho(a_1) = 0$ and $\rho(a_m) = 1$.

Fix $\epsilon < 1/m^2$. The voters are embedded as follows.

$$\begin{aligned}
\rho(v_1) &= \rho(a_1) - \epsilon, \\
\rho(v_2) &= \rho(a_m) + \epsilon, \\
\rho(v_3) &= \rho(a_t - 1) + \epsilon, \\
\rho(v_4) &= \rho(a_t + 1) - \epsilon, \\
\rho(v_{4+(2j-1)}) &= \rho(a_{j+1}) + \epsilon & \forall j \in [m-2], \\
\rho(v_{4+2j}) &= \rho(a_{j+1}) - \epsilon & \forall j \in [m-2], \\
\rho(v_j) &= \frac{1}{2} & \forall j \geq 2m+1.
\end{aligned}$$

Finally, in the party affiliation $\pi$, since we are just showing how bad the distortion of direct elections are and $f$ is affiliation-independent, the party affiliation is not important, and the outcome is independent of $\pi$, so we can assign it in any way such that half the voters are of one party and half are of the other. This construction simply ensures (and it is easy to check) that $I \in \mathcal{I}^{0.5}_{m,\text{sep-}\mathbb{R}} \subseteq \mathcal{I}^\alpha_{m,\text{sep-}\mathbb{R}}$ for all $\alpha \in (0, 0.5]$.

Next, it is also easy to check that election $E^I$ falls under the special case of $f_{\text{fail}}$. Hence, $f_{\text{fail}}(E^I) = a_1$. Note that $C(a_1) \geq (2n - 2m) \cdot |0 - \frac{1}{2}| > n - m$ because $a_1$ is at distance $\frac{1}{2}$ from all but $2m$ voters located at $\frac{1}{2}$. In contrast, $C(a_t) \leq 2m \cdot 1$ because $a_t$ is at zero distance from all but $2m$ voters (and its distance from those $2m$ voters is at most 1). Thus, $\phi(f_{\text{fail}}, I) \geq (n - m)/2m$. Since $n$ is unbounded, $\phi_{\mathcal{I}^\alpha_{m,\text{sep-}\mathbb{R}}}(f_{\text{fail}})$ is also unbounded.

Finally, we show that $\phi_{\mathcal{I}^\alpha_{m,\text{sep-}\mathbb{R}}}(\widehat{f}_{\text{fail}})$ is upper bounded by a constant. Fix an instance $I \in \mathcal{I}^\alpha_{m,\text{sep-}\mathbb{R}}$. For notational simplicity, we refer to the number of candidates in $I$ as $m$, though the proof below works if it is less than $m$. First, assume $|\rho(A)| \neq 2$ (we will handle the case $|\rho(A)| = 2$ later). Without loss of generality, assume that $\rho(a_1) \leq \ldots \leq$

$\rho(a_m)$, and that for a fixed $q \in [m]$, candidates $a_1, \ldots, a_q$ are affiliated with party $-1$ and the rest are affiliated with party 1.

Let $I_{-1}$ and $I_1$ be the primary instances corresponding to $I$. Let $a_{OPT}$ be an optimal candidate for $I$. Without loss of generality, suppose it is affiliated with party $-1$.

In the proof of Theorem 18, $\phi(\widehat{f}, I)$ depends only on the distortion of $f$ on the primary instance of the party that $a_{OPT}$ is affiliated with. Hence, if primary election $E^{I_{-1}}$ does not fall under the special case of $f_{\text{fail}}$, then $f_{\text{fail}}$ applies a Condorcet-consistent rule on $I_{-1}$, ensuring that $\phi(f, I_{-1})$ is at most 3. In this case, by Theorem 18, $\phi(\widehat{f}, I)$ is also upper bounded by a constant.

Suppose $E^{I_{-1}}$ falls under the special case of $f_{\text{fail}}$. Let $t = (q+1)/2$ and $d_{-1} = d(a_1, a_q)$. Then, by Lemma 24, $\min\{d(a_t, a_1), d(a_t, a_q)\} \geq d_{-1}/4$. From now on, we shall use asymptotic notation liberally for simplicity.

Recall that there is a set of voters $S \subset V_{-1}$ whose top candidate was $a_t$, and

$$|S| = |V_{-1}| - 2q = \Omega(|V_{-1}|) = \Omega(n),$$

where the second transition holds because in the special case, $|V_{-1}| \geq q^2$, and the final transition holds because $|V_{-1}| \geq \alpha n$.

Note that for every $i \in S$ and $j \in V_1$, $d(i,j) \geq d_{-1}/4$. And $|V_1| \geq \alpha n$. Hence, we have $\Omega(n)$ pairs of voters $(i,j)$ such that $d(i,j) \geq d_{-1}/4$. Further, $d(a_{OPT}, i) + d(a_{OPT}, j) \geq d(i,j)$. Hence, it follows that

$$C(a_{OPT}) = \Omega(n) \cdot d_{-1}. \tag{5.4}$$

Let $a^* = \widehat{f}(I)$. If $a^* = a^*_{-1} = a_1$, then we have

$$\begin{aligned} C(a^*) &\leq C(a_{OPT}) + n \cdot d(a^*, a_{OPT}) \\ &\leq C(a_{OPT}) + n \cdot d_{-1} = O(C(a_{OPT})), \end{aligned}$$

yielding a constant upper bound on $\phi(\widehat{f}_{\text{fail}}, I) = C(a^*)/C(a_{OPT})$.

On the other hand, if $a^* = a^*_1$, we have

$$\begin{aligned} C(a_{OPT}) &\geq C(a_1) - n \cdot d(a_1, a_{OPT}) \\ &\geq \frac{n}{2} \frac{d(a_1, a^*)}{2} - n \cdot d_{-1} \\ &\geq \frac{n}{2} \frac{d(a_{OPT}, a^*)}{2} - O(C(a_{OPT})). \end{aligned}$$

Here, the second transition follows because in the general election, at least $n/2$ voters vote for $a^*$ over $a_1$ and $d(a_1, a_{OPT}) \leq d_{-1}$, and the final transition follows from Equation (5.4). This implies

$$C(a_{OPT}) = \Omega(n \cdot d(a_{OPT}, a^*)). \tag{5.5}$$

On the other hand, we have

$$C(a^*) \leq C(a_{OPT}) + n \cdot d(a_{OPT}, a^*) = O(C(a_{OPT})),$$

where the last transition follows due to Equation (5.5). Hence, we again have the desired constant upper bound on $\phi(\widehat{f}_{\text{fail}}, I)$.

Finally, if $\rho(A) = \{x_1, x_2\}$ (w.l.o.g., $x_1 < x_2$), due to separability, we have two options:

1. $\rho(A_{-1}) = x_1$ and $\rho(A_1) = x_2$: Therefore, $a_{OPT} \in \{a*_{-1}, a*_1\}$, and since $a*$ is the majority winner, $a* = a_{OPT}$, and the distortion is 1.

2. $\rho(A_{-1}) = \{x_1, x_2\}$ and $\rho(A_1) = x_2$: If $a_{OPT} \in \{a*_{-1}, a*_1\}$, then it is similar to the previous case. Otherwise, this means $\rho(a_{OPT}) = x_1$ and $\rho(a*_{-1}) = \rho(a*_1) = x_2$. Separability means the voters of party 1 lie in $\{x|x \geq x_2\}$, so $C(a_{OPT}) \geq |V_1| \cdot d(x_1, x_2) \geq \alpha \cdot n \cdot d(x_1, x_2)$. While $C(a*) \leq C(a_{OPT}) + |V_{-1}| \cdot d(x_1, x_2) \leq C(a_{OPT}) + n(1-\alpha)d(x_1, x_2)$. Combining these two equations we get the distortion is $\frac{1+2\alpha}{\alpha}$.

□

## 5.6 Relaxing the assumptions I: voter participation

We have analyzed above the primary system under two restrictive assumptions: every voter participates in exactly one of the two primaries, and every voter participates in the general election. In reality, only a fraction of eligible voters participate in the primaries or in the general election. *What happens to the distortion under the primary system when not all eligible voters vote, but we still measure the social cost of the winning candidate with respect to all eligible voters?* We show that under such a relaxation, the distortion of any voting rule under the primary system changes by at most a constant factor. Hence, all of our results continue to hold.

To relax the assumptions, we extend our formal framework as follows. Recall that so far, we denoted an instance by $I = (V, A, M, \rho, \pi)$, where $V$ is a set of $n$ voters, $A$ is a set of candidates, and party affiliation $\pi : V \cup A \to \{-1, 1\}$ maps every voter and candidate to one of the two parties. We further assumed that each voter $v$ participates in the primary of party $\pi(v)$ as well as in the general election.

In the extended framework, we denote an instance by $I = (V, A, M, \rho, \pi, \tau)$. Here, the party affiliation function $\pi : V \cup A \to \{-1, 0, 1\}$ is allowed to map a voter to 0, which indicates that the voter does not participate in either party's primary election. This incorporates not only independent voters, but also voters affiliated with a party that do not participate in the party's primary election. We still require that $\pi(a) \in \{-1, 1\}$ for every $a \in A$, i.e., that every candidate is affiliated with a party and participates in that party's primary election. We also have the additional function $\tau : V \to \{0, 1\}$, which maps each voter to 1 if the voter participates in the general election, and to 0 otherwise. This relaxes the assumption that all voters participate in the general election.

Finally, recall that for $\alpha \in [0, 0.5]$, we defined a family of instances $\mathcal{I}^{\alpha}_{m,\mathcal{M}}$ such that in every instance $I \in \mathcal{I}^{\alpha}_{m,\mathcal{M}}$ with $n$ voters, at least $\alpha n$ voters are affiliated with each party (and participate in that party's primary election). In the relaxed framework, for $\alpha \in [0, 0.5]$, $\beta \in [0, 1]$, and $\gamma \in [0, 1]$, we define a family of instances $\mathcal{I}^{\alpha,\beta,\gamma}_{m,\mathcal{M}}$ such that in every instance $I \in \mathcal{I}^{\alpha,\beta,\gamma}_{m,\mathcal{M}}$ with $n$ voters, at least $\alpha n$ voters participate in each primary election, at least $\beta n$ of voters participated in the primaries (obviously, $\beta \geq 2\alpha$), and at least $\gamma n$ voters participate in the general election. Formally, we require $|\pi^{-1}(p)| \geq \alpha n$ for each $p \in \{-1, 1\}$, $|\pi^{-1}(-1) \cup \pi^{-1}(1)| \geq \beta n$ and $|\tau^{-1}(1)| \geq \gamma n$. Note that the relaxed family $\mathcal{I}^{\alpha,\beta,\gamma}_{m,\mathcal{M}}$ includes all instances from the restricted family $\mathcal{I}^{\alpha}_{m,\mathcal{M}}$, which are obtained when $\beta = \gamma = 1$.

While our restricted framework required that $\pi(v) \in \{-1, 1\}$ and $\tau(v) = 1$ for every voter $v$, the extended framework allows the following possibilities.

1. $\pi(v) = 0$ and $\tau(v) = 1$: the voter does not participate in either primary but participates in the general election.

2. $\pi(v) \in \{-1, 1\}$ and $\tau(v) = 0$: the voter participates in a primary election but does not participate in the general election.[5]

3. $\pi(v) \in 0$ and $\tau(v) = 0$: the voter does not participate in any election, primary or general.

Our goal is to show that for any voting rule $f$ and constants $\alpha, \beta, \gamma > 0$, the distortion of $f$ under the primary system in the relaxed framework ($\phi_{\mathcal{I}^{\alpha,\beta,\gamma}_{m,\mathcal{M}}}(\widehat{f})$) is no more than a constant times higher than in the distortion in the restricted framework ($\phi_{\mathcal{I}^{\alpha}_{m,\mathcal{M}}}(\widehat{f})$); the fact that $\phi_{\mathcal{I}^{\alpha,\beta,\gamma}_{m,\mathcal{M}}}(\widehat{f}) \geq \phi_{\mathcal{I}^{\alpha}_{m,\mathcal{M}}}(\widehat{f})$ follows trivially from the fact that the relaxed framework is strictly more general than the restricted framework.

We begin by proving a generalization of Lemma 21 in the relaxed framework, in which we establish that the social costs of the two primary winners must not differ by more than a constant factor.

**Lemma 26.** *Consider an instance $I \in \mathcal{I}^{\alpha,\beta,\gamma}_{m,\mathcal{M}}$ where $\alpha, \beta, \gamma > 0$. In the primary system, let $a^*_{-1}$ and $a^*_1$ be the two primary winners, $a^* \in \left\{a^*_{-1}, a^*_1\right\}$ be the winner of the general election, and $\widehat{a} \in \left\{a^*_{-1}, a^*_1\right\} \setminus \{a^*\}$. Then,*

$$C(a^*) \leq \left(\frac{4}{\gamma} - 1\right) \cdot C(\widehat{a}).$$

*Proof.* Because at least $\gamma n$ voters participate in the general election and $a^*$ wins by a majority vote, there must exist $X \subseteq V$ with $|X| \geq \frac{\gamma n}{2}$ such that $d(i, a^*) \leq d(i, \widehat{a})$ for every $i \in X$. Combining with the triangle inequality $d(a^*, \widehat{a}) \leq d(a^*, i) + d(i, \widehat{a})$, we get $d(i, \widehat{a}) \geq \frac{d(a^*, \widehat{a})}{2}$ for every $i \in X$. Hence,

$$C(\widehat{a}) \geq \frac{\gamma n}{2} \cdot \frac{d(a^*, \widehat{a})}{2} \Rightarrow d(a^*, \widehat{a}) \leq \frac{4}{\gamma n} \cdot C(\widehat{a}). \tag{5.6}$$

Now, we have

$$\begin{aligned} C(a^*) &= \sum_{i \in X} d(i, a^*) + \sum_{i \in V \setminus X} d(i, a^*) \\ &\leq \sum_{i \in X} d(i, \widehat{a}) + \sum_{i \in V \setminus X} (d(i, \widehat{a}) + d(a^*, \widehat{a})) \\ &\leq C(\widehat{a}) + \left(n - \frac{\gamma n}{2}\right) \cdot d(a^*, \widehat{a}), \end{aligned} \tag{5.7}$$

where the second transition follows because $d(i, a^*) \leq d(i, \widehat{a})$ for $i \in X$, and $d(i, a^*) \leq d(i, \widehat{a}) + d(a^*, \widehat{a})$ due to the triangle inequality. Substituting Equation (5.6) into Equation (5.7), we get the desired result. $\square$

We are now ready to prove the main result of this section.

**Theorem 27.** *For every choice of $\mathcal{M}$, $m \in \mathbb{N}$, affiliation-independent voting rule $f$, and constants $\alpha, \beta > 0$,*

$$\phi_{\mathcal{I}^{\alpha,\beta,\gamma}_{m,\mathcal{M}}}(\widehat{f}) \leq \left(\frac{4}{\gamma} - 1\right)\left(1 + \frac{4}{\beta}\right) \cdot \phi_{\mathcal{I}^{\alpha}_{m,\mathcal{M}}}(\widehat{f}).$$

[5] While data indicates that almost all voters who participate in the primary elections also participate in the general election [103], we consider this possibility for the sake of completeness.

*Proof.* Consider an instance $I = (V, A, M, \rho, \pi, \tau) \in \mathcal{I}^{\alpha,\beta,\gamma}_{m,\mathcal{M}}$ with $n$ voters in the relaxed framework. Recall that function $\pi : V \cup A \to \{-1, 0, 1\}$ dictates whether each voter participates in the primary of party $-1$, in the primary of party 1, or in neither primary. Similarly, function $\tau : V \to \{0, 1\}$ dictates whether each voter participates in the general election.

Let us now construct an instance $\overline{I} = (\overline{V}, A, M, \rho, \overline{\pi}) \in \mathcal{I}^{\alpha}_{m,\mathcal{M}}$ in the restricted framework as follows: We set $\overline{V} = \left(\pi^{-1}(-1) \cup \pi^{-1}(1)\right) \cap V$; that is, we keep the voters who participate in the primaries in $\overline{I}$, and delete the voters who do not. Note that $|\overline{V}| \geq \beta n$. Let $\overline{\pi} : \overline{V} \cup A \to \{-1, 1\}$ be the restriction of $\pi$ to $\overline{V} \cup A$ (note that $\pi$ does not map any voter in $\overline{V}$ to 0, and cannot map any candidate in $A$ to 0). Note that at least $\alpha n \geq \alpha|\overline{V}|$ voters participate in each primary election; hence, this is a valid instance of $\mathcal{I}^{\alpha}_{m,\mathcal{M}}$.

Crucially, note that instances $I$ and $\overline{I}$ match in the set of voters that participate in each primary election. They only differ in the set of voters that participate in the general election.[6] Hence, the primary winners in instances $I$ and $\overline{I}$ must be identical.

Let $a_{-1}$ and $a_1$ denote the primary winners of parties $-1$ and 1, respectively, in both $I$ and $\overline{I}$. Let $a_{OPT} \in \arg\min_{a \in A} C^I(a)$ and $\overline{a_{OPT}} \in \arg\min_{a \in A} C^{\overline{I}}(a)$ denote the socially optimal candidates in $I$ and $\overline{I}$, respectively, and let $a^*$ and $\overline{a^*}$ denote the winners of the general elections in $I$ and $\overline{I}$, respectively. Hence, $\phi(\widehat{f}, I) = C^I(a^*)/C^I(a_{OPT})$ and $\phi(\widehat{f}, \overline{I}) = C^{\overline{I}}(\overline{a^*})/C^{\overline{I}}(\overline{a_{OPT}})$.

For any candidate $a$, we have

$$\begin{aligned}\beta n \cdot d(\overline{a_{OPT}}, a) \leq |\overline{V}| \cdot d(\overline{a_{OPT}}, a) &\leq \sum_{i \in \overline{V}} d(i, \overline{a_{OPT}}) + d(i, a) \\ &= C^{\overline{I}}(\overline{a_{OPT}}) + C^{\overline{I}}(a) \\ &\leq 2C^{\overline{I}}(a),\end{aligned}$$

where the first transition follows because $|\overline{V}| \geq \beta n$, the second transition follows from the triangle inequality, and the final transition holds because $\overline{a_{OPT}}$ is the socially optimal candidate in $\overline{I}$. Hence, for every candidate $a$, we have

$$d(\overline{a_{OPT}}, a) \leq \frac{2}{\beta n} C^{\overline{I}}(a). \tag{5.8}$$

[6]In the general election of $\overline{I}$, we removed voters who participated in the general election but not in the primaries in $I$, and added voters who participated in the primaries but not in the general election in $I$.

Now, we have

$$
\begin{aligned}
C^I(a^*) &\leq \left(\frac{4}{\gamma}-1\right) C^I(\overline{a^*}) = \left(\frac{4}{\gamma}-1\right) \sum_{i\in V} d(i,\overline{a^*}) \\
&\leq \left(\frac{4}{\gamma}-1\right) \cdot \sum_{i\in V} d(i,a_{OPT}) + d(a_{OPT},\overline{a^*}) \\
&\leq \left(\frac{4}{\gamma}-1\right) \cdot \left(C^I(a_{OPT}) + n\cdot \left(d(a_{OPT},\overline{a_{OPT}}) + d(\overline{a_{OPT}},\overline{a^*})\right)\right) \\
&\leq \left(\frac{4}{\gamma}-1\right) \cdot \left(C^I(a_{OPT}) + n\cdot \left(\frac{2}{\beta n} C^{\overline{I}}(a_{OPT}) + \frac{2}{\beta n} C^{\overline{I}}(\overline{a^*})\right)\right) \\
&\leq \left(\frac{4}{\gamma}-1\right) \cdot \left(\left(1+\frac{2}{\beta}\right) C^I(a_{OPT}) + \frac{2}{\beta} C^{\overline{I}}(\overline{a^*})\right), \qquad (5.9)
\end{aligned}
$$

where the first transition follows from the fact that both $a^*$ and $\overline{a^*}$ are primary winners in $I$ and from Lemma 26, the third and the fourth transitions follow from the triangle inequality, the fifth transition uses Equation (5.8), and the final transition uses the fact that $C^{\overline{I}}(a) \leq C^I(a)$ for any candidate $a$ since the set of voters in $I$ is a superset of the set of voters in $\overline{I}$. Finally, we have

$$
\begin{aligned}
\phi(\widehat{f}, I) &= \frac{C^I(a^*)}{C^I(a_{OPT})} \\
&\leq \left(\frac{4}{\gamma}-1\right) \cdot \frac{1}{C^I(a_{OPT})} \left(\left(1+\frac{2}{\beta}\right) \cdot C^I(a_{OPT}) + \frac{2}{\beta} \cdot C^{\overline{I}}(\overline{a^*})\right) \\
&\leq \left(\frac{4}{\gamma}-1\right) \cdot \left(\left(1+\frac{2}{\beta}\right) + \frac{2}{\beta} \cdot \frac{C^{\overline{I}}(\overline{a^*})}{C^I(a_{OPT})}\right) \\
&\leq \left(\frac{4}{\gamma}-1\right) \cdot \left(\left(1+\frac{2}{\beta}\right) + \frac{2}{\beta} \cdot \frac{C^{\overline{I}}(\overline{a^*})}{C^{\overline{I}}(\overline{a_{OPT}})}\right) \\
&= \left(\frac{4}{\gamma}-1\right) \cdot \left(\left(1+\frac{2}{\beta}\right) + \frac{2}{\beta} \cdot \phi(\widehat{f}, \overline{I})\right) \\
&\leq \left(\frac{4}{\gamma}-1\right) \left(1+\frac{4}{\beta}\right) \phi(\widehat{f}, \overline{I}),
\end{aligned}
$$

where the first transition uses Equation (5.9), and the second transition uses the fact that $C^I(a_{OPT}) \geq C^{\overline{I}}(a_{OPT}) \geq C^{\overline{I}}(\overline{a_{OPT}})$. Thus, we have the desired result. □

Since having independent voters in the general election and having voters who only participate in the primaries cannot improve the worst-case distortion of a voting rule under the primary system, our earlier results which establish that primaries are no better than direct elections (Theorems 17 and 22) continue to hold in this extended framework. Additionally, we proved that the existence of such voters can make the distortion under the primary system worse by at most a constant factor (which depends on $\beta$ and $\gamma$). Hence, our earlier results which establish that primaries cannot be much worse than direct elections (Theorem 18) or that primaries can be significantly better than direct elections (Theorem 25) also continue to hold in this extended framework.

## 5.7 Relaxing the assumptions II: multiple parties and complex structure

We now examine another relaxation of our assumptions; namely, so far we have assumed a 2-stage, 2-party primary election. Having multiple parties means some voting rule needs to be used to decide between the different general election candidates (since majority, which we used in the 2-party setting, may not give a definite answer). Moreover, we wish to explore scenarios which involve more complex settings, in which the decision process extends beyond a 2-stage mechanism. For example, subsets of parties' electorate decide on the candidates for the party, which the party than further cuts down, presenting a single candidate for the general elections[7].

The extension to multiple parties is relatively straightforward. For $r$ parties we extend party membership such that the party affiliation function $\pi$ returns a value in $\{0, 1, \ldots, r\}$ (with $0$, as in Section 5.6, referring to voters who do no participate in any party's primary. We shall assume each party uses its own election rule for its internal process, so party $p \in \{1, \ldots, r\}$ uses voting rule $f_p$. The general election uses the voting rule $f$, with its electorate being $V' \subset V$. We shall mark the voting rule employing this structure as $\widetilde{f}$.

The lower bound on such a structure is unchanged from the 2-stage case; more specifically, simply using the same voting rule from Theorem 25's proof results in a similar improvement of the distortion for multi-party settings over direct settings. However, the upper bound in different.

We shall mark settings with $r$ parties and 2 electoral stages with $\mathcal{I}^{\alpha,\beta,\gamma}_{m,r,\mathcal{M}}$, with $\alpha, \beta$ and $\gamma$ serving the same bounds the did in Section 5.6: each party has at least $\alpha n$ voters, at least $\beta n$ voters participate in the primary stage (hence $\beta \geq r\alpha$), and at least $\gamma n$ voters participate in the final voting stage.

**Theorem 28.** *For more than 2 parties ($r \geq 3$), each of which using its own voting rule to select its candidate (i.e., party $p$ uses voting rule $f_p$), and any general election voting rule $f$, for $I \in \mathcal{I}^{\alpha,\beta,\gamma}_{m,r,\mathcal{M}}$ (with each party being $I_p$, and the general electorate being $I_{V'}$)*

$$\phi(I, \widetilde{f}) \leq \frac{1 - \gamma + \phi(I_{V'}, f)}{\gamma} \cdot \max_{p \in \{1,\ldots,r\}} \left(\frac{1 - \alpha + \phi(I_p, f_p)}{\alpha}\right)$$

*Proof.* Let $a_{OPT}$ be the optimal candidate (i.e., $a_{OPT} = \arg\min_{a \in A}(C(a))$), and $a^*$ the candidate that won the overall election. Let $p* = \pi(a_{OPT})$ be the party to which $a_{OPT}$ belongs, and let $a^*_{p*}$ be the candidate which won the elections in party $p*$, which used $f_{p*}$. Thanks to Lemma 20 we know $\frac{C(a^*_{p*})}{C(a_{OPT})} \leq \frac{1-\alpha+\phi(I_{p*}, f_{p*})}{\alpha}$. Showing the relation between $C(a^*)$ and $C(a^*_{p*})$ is done quite similarly to that lemma as well (in a sense, one can consider the agents not voting in the general election as the "non voting" party) :

$$\begin{aligned} C(a^*) = \sum_{i \in V} d(i, a^*) &= \sum_{i \in V'} d(i, a^*) + \sum_{i \in V \setminus V'} d(i, a^*) \leq C_{V'}(a^*) + \sum_{i \in V \setminus V'} d(i, a^*_{p*}) + d(a^*, a^*_{p*}) \\ &\leq \phi(I_{V'}, f) C_{V'}(a^*_{p*}) + C_{V \setminus V'}(a^*_{p*}) + (n - |V'|) d(a^*, a^*_{p*}) \\ &\leq \phi(I_{V'}, f) C(a^*_{p*}) + (n - |V'|) d(a^*, a^*_{p*}) \end{aligned} \tag{5.10}$$

with the second inequality stemming from the distortion definition – for any candidate $a$ in the general election,

[7] For example, in the British party system, after candidates for party leadership announce their candidacy, the parties' Members of Parliament cull down the number of candidates (in the Labour party to 6 candidates at most; in the Conservative party to 2 candidates). The "surviving" candidates are put forward to the party membership, and the winner in each party leads their respective parties in the general election.

$\frac{C_{V'}(a^*)}{C_{V'}(a)} \leq \phi(I_{V'}, f)$. For each voter $i \in V$, and in particular, for voters who participated in the general election ($i \in V'$), $d(a^*, a^*_{p*}) \leq d(i, a^*) + d(i, a^*_{p*})$. Therefore

$$\begin{aligned}
d(a^*, a^*_{p*}) &\leq \frac{1}{|V'|} \sum_{i \in V'} d(i, a^*) + d(i, a^*_{p*}) = \frac{1}{|V'|}(C_{V'}(a^*) + C_{V'}(a^*_{p*})) \\
&\leq \frac{1}{|V'|}(\phi(I_{V'}, f) C_{V'}(a^*_{p*}) + C_{V'}(a^*_{p*})) = \frac{C_{V'}(a^*_{p*})}{|V'|}(1 + \phi(I_{V'}, f)) \\
&\leq \frac{C(a^*_{p*})}{|V'|}(1 + \phi(I_{V'}, f)) \qquad (5.11)
\end{aligned}$$

the second inequality stems, as in the previous equation, from the distortion definition, and the third inequality is since for $V' \subseteq V$ and any candidate $a$, $C_V(a) \geq C_{V'}(a)$.

Combining equations 5.10 and 5.11,

$$\begin{aligned}
C(a^*) &= C(a^*_{p*})(\phi(I_{V'}, f) + \frac{n - |V'|}{|V'|}(1 + \phi(I_{V'}, f))) \\
&\leq C(a^*_{p*})(\phi(I_{V'}, f) + \frac{1-\gamma}{\gamma}(1 + \phi(I_{V'}, f))) = \\
&= C(a^*_{p*}) \frac{1 - \gamma + \phi(I_{V'}, f)}{\gamma}
\end{aligned}$$

Finally:

$$\phi(I, \widehat{f}) = \frac{C(a^*)}{C(a_{OPT})} = \frac{C(a^*)}{C(a^*_{p*})} \cdot \frac{C(a^*_{p*})}{C(a_{OPT})} \leq \frac{1 - \gamma + \phi(I_{V'}, f)}{\gamma} \cdot \frac{1 - \alpha + \phi(I_{p*}, f_{p*})}{\alpha}$$

□

Naturally, Theorem 28 can be applied multiple times for more complex multi-stage voting. For example, a mechanism in which $\gamma n$ voters participate in the general elections, and each of the $r$ parties has at least $\alpha n$ voters, and in each party, there are $r'$ committees, each consisting of $\alpha'$ portion of the party voters, which choose one candidate to represent them. In this case, the distortion is bounded by applying Theorem 28 twice: $\frac{1-\gamma+\phi(f)}{\gamma} \max_{p \in \{1,\ldots,r\}}(\frac{1-\alpha+\phi(f_{p*})}{\alpha} \cdot \max_{p' \in \{1,\ldots,r'\}}(\frac{1-\alpha'+\phi(f_{p'})}{\alpha'}))$.

In order to formalize this, we introduce the idea of a tree structure for the decision making process:

**Definition 13.** A *tree-structured* primary process is a rooted tree, in which each node $x$ is comprised of a set of voters and a voting rule $f_x$. Each leaf also contains a set of candidates $A_x$. Each leaf holds an election, choosing a single candidate ($a^*_x$) that will participate in the election in its' parent node.

The ultimate winner is the candidate which will be victorious in the election of the tree root.

Note that this tree structure allows for modeling of multi-staged decision making processes (as in the British party system, suggested above) and much more complex structures as well.

**Corollary 29.** *Let $I$ be a tree-structured primary process, such that each node $x$ contains at least $\alpha_x n$ voters. The distortion for this process is bounded from above by the maximum (over all branches $x_1, x_{2,}, \ldots, x_k$ from the tree root to a leaf) value of $\prod_{i=1}^{k} \frac{1-\alpha_{x_i}+\phi(f_{x_i})}{\alpha_{x_i}}$.*

*Proof.* This is a result of multiple applications of Theorem 28 using induction on the height of the tree. Theorem 28 is the base case for a tree of height 2, and for any tree of height $\ell > 2$, one applies Theorem 28 on all the leafs of the maximal height. Notice that each node has an affiliated distortion value, which depends on its voting rule. So for a leaf $x$, its distortion value is $\phi(f_x)$. However, when we apply Theorem 28 multiple times, we are simply assigning new distortion values to leaf nodes that incorporate the process that led to the candidates they can vote on. That is, if node $y$ of height $\ell - 1$ has children $y_1, \dots, y_r$, its new distortion will now be $\phi(f_y) \max_{1,\dots,r} \frac{1-\alpha_{y_i}+\phi(f_{y_i})}{\alpha_{y_i}}$. Now we have a tree of height $\ell - 1$, with updated distortion values for its leaf nodes, and we use the induction hypothesis to complete the proof. □

## 5.8 Using simulations to go beyond worst case

So far we compared the distortion of a voting rule under the direct and primary systems, taken in the *worst case* over a family of instances. In practice, such worst-case instances may not arise naturally. In this section, we investigate the distortion of a voting rule under the direct and primary systems, in the *average case* over simulated instances. Our simulations are limited to the two-stage, two-party primary process that we focus on in the bulk of this paper. We generate the simulated instances by varying a number of parameters; to keep the number of simulations reasonable, when varying one parameter, we use default values of the other parameters[8]:

**Total voters** The number of voters $n$: default = 500, range = 100 to 2100 in increments of 200.

**Total candidates** The number of candidates $m$: default = 50, range = 10 to 210 in increments of 20.

**Independent voters** The percentage of voters who are independent: default = $0\%$, range = $0\%$ to $90\%$ in increments of $10\%$.

**Independent candidates** The percentage of candidates who are independent: default = $0\%$, range = $0\%$ to $90\%$ in increments of $10\%$.

**Party voter balance** The percentage of voters who are affiliated with party $-1$: default = $50\%$,[9] range = $10\%$ to $90\%$ in increments of $10\%$.

**Party candidate balance** The percentage of candidates who are affiliated with party $-1$: default = $50\%$,[10] range = $10\%$ to $90\%$ in increments of $10\%$.

**Metric space dimension** The dimension $k$ of the Euclidean metric space $[0,1]^k$: default = 4, range = $\{1, 4, 7, 10\}$.

For a given combination of the parameter values, we generate random instances as follows. First, we place a set $V$ of $n$ voters at uniformly random locations in $[0,1]^k$. Next, if the ratio of the number of voters in the two parties is supposed to be $x : (1-x)$, then we find a hyperplane dividing voters into this $x : (1-x)$ ratio. Due to symmetry, we simply find a threshold $t$ on the $k^{\text{th}}$ coordinate such that the locations of $x$ fraction of the voters (call this set $V_{-1}$) have $k^{\text{th}}$ coordinate at most $t$, while the locations of the rest (call this set $V_1$) have $k^{\text{th}}$ coordinate at least $t$. We do not set

[8]As we observe later, the dimension of the metric space had perhaps the most significant effect on the distortion. To verify that this was not an artifact of the default values of the other parameters, we varied the dimension along with every other parameter. Our figures, however, are generated with the default value for the dimension.

[9]Given that there are no independent voters in these simulations, by default we split voters equally between the two parties.

[10]Given that there are no independent candidates in these simulations, by default we split the candidates equally between the two parties.

| | Primary is better | Direct is better | No detected difference |
|---|---|---|---|
| plurality split | 183 | 7 | 18 |
| plurality random | 203 | 0 | 5 |
| STV split | 8 | 178 | 22 |
| STV random | 152 | 13 | 48 |
| Borda split | 2 | 202 | 4 |
| Borda random | 8 | 179 | 21 |
| maximin split | 0 | 207 | 1 |
| maximin random | 0 | 208 | 0 |
| Copeland split | 0 | 207 | 1 |
| Copeland random | 0 | 208 | 0 |

Table 5.1: The table shows the number of settings (out of 208) in which each of primary and direct systems leads to a lower average distortion than the other. Statistical significance is measured using a paired t-test with $p = 0.05$.

voter affiliations yet. Next, if the ratio of the number of candidates in the two parties is supposed to be $x : (1-x)$, then we place $\lceil x \cdot m \rceil$ candidates (call this set $A_{-1}$) uniformly at random on one side of the hyperplane, and the remaining candidates (call this set $A_1$) uniformly at random on the other side. Finally, if $x\%$ of the voters (resp. candidates) are supposed to be independent, then we choose $x\%$ of the voters (resp. candidates) from $V_{-1}$ and $V_1$ each (resp. from $A_{-1}$ and $A_1$ each), remove them from the respective sets, and distribute them at uniformly random locations in $[0, 1]^k$.

Once the locations of the voters and candidates are fixed, we create two instances. In one instance (called "split"), we assign $V_{-1} \cup A_{-1}$ to party $-1$, and assign $V_1 \cup A_1$ to party 1. In this instance, we have party separability. In the other instance (called "random"), we assign $|V_{-1}|$ voter and $|A_{-1}|$ candidates chosen uniformly at random to party $-1$, and among the rest, assign $|V_1|$ voter and $|A_1|$ candidates chosen uniformly at random to party 1. In this instance, we do not have party separability. This allows us to compare the effect of party separability on the distortion. We run five voting rules — plurality, Borda, STV, maximin, and Copeland — on both instances under the direct and primary systems, and measure the distortion. Note that their distortion under the direct system would be identical for party separable and random instances because the two instances only differ in party affiliations. Thus, for each rule, we obtain three numbers: Direct, Primary-split, and Primary-random. For each combination of parameter values, we repeat this 1000 times and take the average distortion.

Borda, which is a positional scoring rule like plurality, surprisingly tracks the Condorcet-consistent rules Copeland and maximin. This provides some support to a long line of papers in the literature establishing that Borda is "close to being Condorcet consistent" [42, 99, 44, 46, 45].

### 5.8.1 Primary versus direct: summary

We now present our results comparing the primary and direct systems. Our experiments result in 208 settings (combination of parameter values). For each setting, we compare the average primary distortion and the average direct distortion of each voting rule under each party affiliation model (split or random), and evaluate which system results in a better average distortion. For statistical significance, we use the paired t-test with $p = 0.05$. The results are presented in Table 5.1.

Without party separability (i.e. in the random case), our theoretical results indicate that primary is no better than direct in the worst case (Theorem 22). While this is also true in our experimentally generated average cases for Borda, Copeland, and maximin, we see that for plurality and STV, primary almost always outperforms direct in the average case.

With party separability (i.e. in the split case), direct outperforms primary for all voting rules except plurality. While our theoretical result shows that for plurality direct also outperforms primary in the worst case (Proposition 23), we see that this is not true in the average case.

It is also interesting to see that for STV, party separability crucially affects which system works better.

### 5.8.2 Primary versus direct: margins

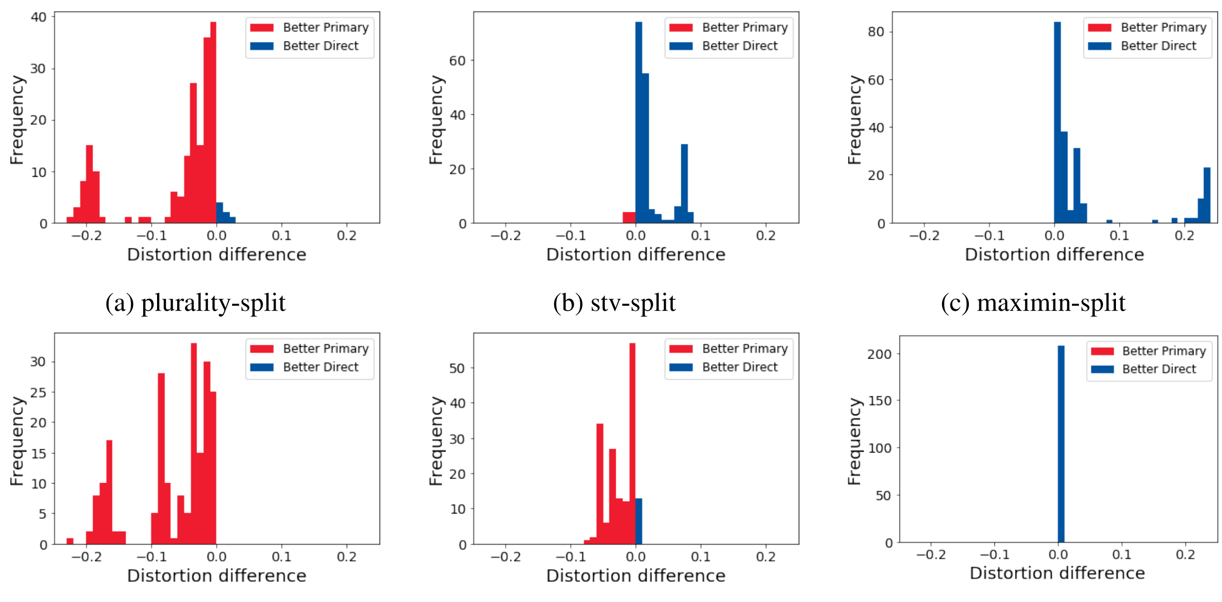


(a) plurality-split (b) stv-split (c) maximin-split

(d) plurality-random (e) stv-random (f) maximin-random

Figure 5.1: These histograms show the difference between the average distortion under the primary system and that under the direct system.

Table 5.1 shows the number of settings in which each system outperforms the other, but it does not tell us about the margin by which it outperforms. It could very well be the case that when one system outperforms the other, it does so by a large margin, whereas when the latter outperforms the former, it only does so by a small margin.

To investigate this, we look at the the difference between the average distortion under the primary system and the average distortion under the direct system. The results for plurality, STV, and maximin are given in Figure 5.1. The figures for Borda and Copeland are omitted because they were similar to the figures for maximin.

The results are quite varied. In some cases where primary mostly outperforms direct (e.g. in plurality-split or plurality-random), primary sometimes outperforms direct by a large margin whereas direct only outperforms primary by a small margin. But we see that the opposite is also true (e.g. in maximin-split). Hence, neither system seems to have a significant advantage over the other in terms of the difference in average distortions.

### 5.8.3 Effect of varying parameters on the distortion of primary

While the direct system and its distortion have long been studied in the computational social choice literature, the primary system has not received much attention. We now take a closer look at the distortion under the primary system, and how it is affected by different parameters.

#### The number of voters

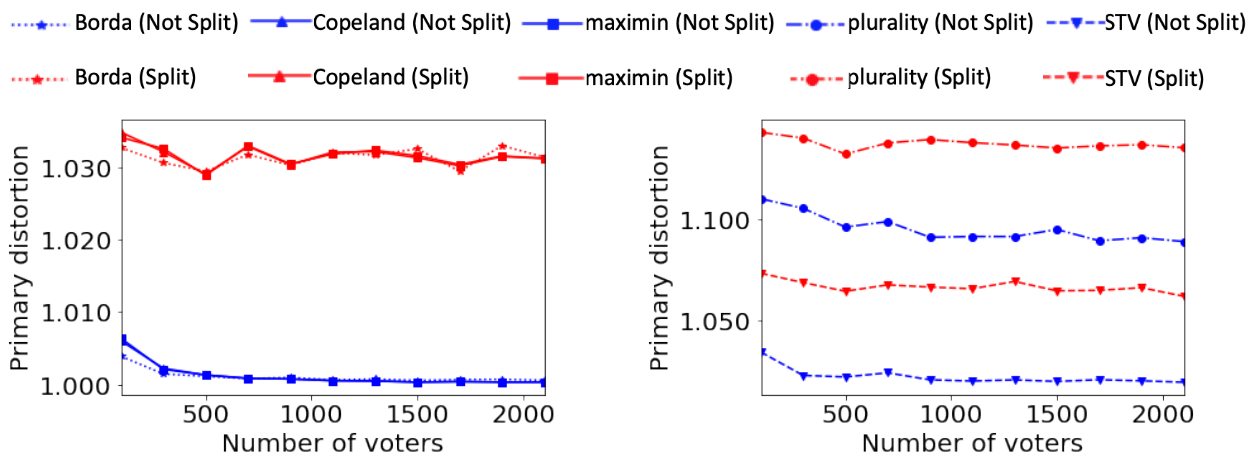


(a) Borda, Copeland and maximin

(b) Plurality and STV

Figure 5.2: The average distortion under the primary system as a function of the number of voters.

Figure 5.2 shows the effect of the number of voters. Initially, the average distortion decreases as the number of voters increases (presumably because the effect of outliers is reduced), but it quickly flatlines as the number of voters grows further. This is not surprising: with an infinite number of voters, it should converge to the distortion with respect to the underlying distribution.

#### The number of candidates

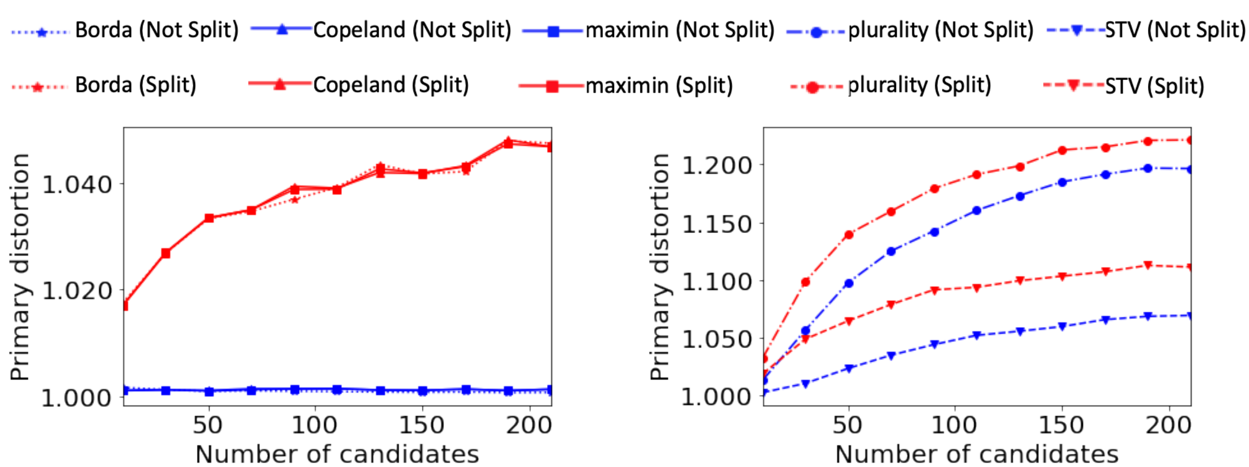


(a) Borda, Copeland and maximin

(b) Plurality and STV

Figure 5.3: The average distortion under the primary system as a function of the number of candidates.

The number of candidates has a more interesting effect on the average distortion, depicted in Figure 5.3. For plurality and STV, the average distortion grows with the number of candidates, which mimics the worst-case behavior. For Borda, Copeland, and maximin, the average distortion grows in the split case, but shows minimal change in the random case. In the random case, the behavior of Borda, Copeland and mamimin is so similar that their distortions overlap in the figure.

### The relative number of voters in each party

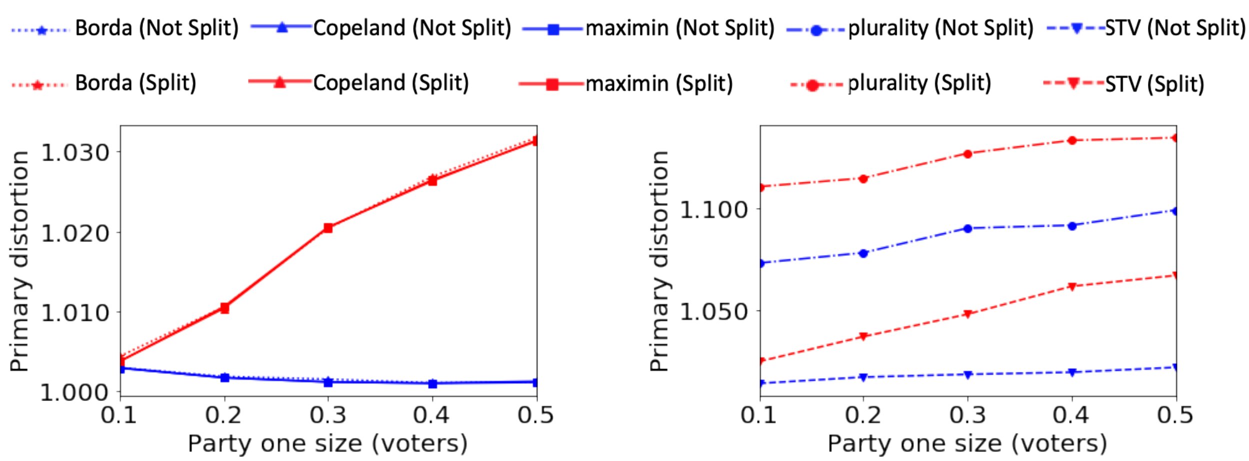


(a) Borda, Copeland and maximin

(b) Plurality and STV

Figure 5.4: The average distortion under the primary system as a function of the fraction of the 500 voters in the first party.

The relative size of the parties has an interesting effect on the average distortion. As both sub-figures of Figure 5.4 there was a striking difference between cases where parties were separated and not separated. In general when parties were not separated, the relative size of parties, in terms of voters, did not have a large impact on distortion[11]. On the other hand when parties are separated, distortion tends to increase as the number of voters in each party becomes more balanced, this is especially evident when using the Condorcet consistent rules (and Borda).

[11] In the non-separable case, the distortion for STV does increase slightly as the number of voters in each party becomes balanced.

**The number of candidates in each party**

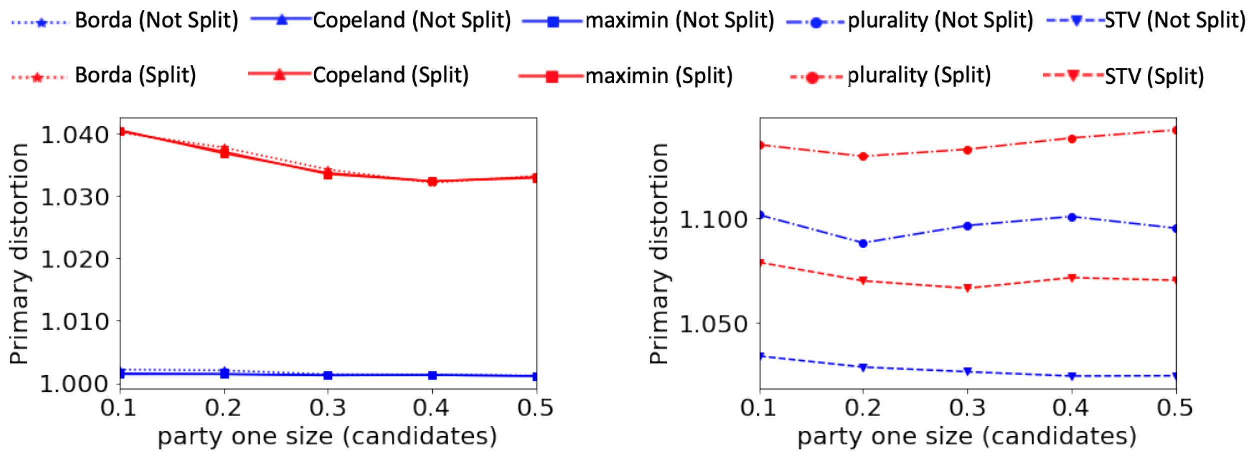


(a) Borda, Copeland and maximin

(b) Plurality and STV

Figure 5.5: The average distortion under the primary system as a function of the fraction of the 50 candidates in the first party.

As Figure 5.5 shows, in most cases the average distortion of the primary election was not impacted by the relative size of parties, in terms of the number of candidates. The only exception to this was the Condorcet consistent rules (and Borda) which shows a slight decrease in distortion as the parties became closer in terms of number of candidates.

**The percentage of independent voters**

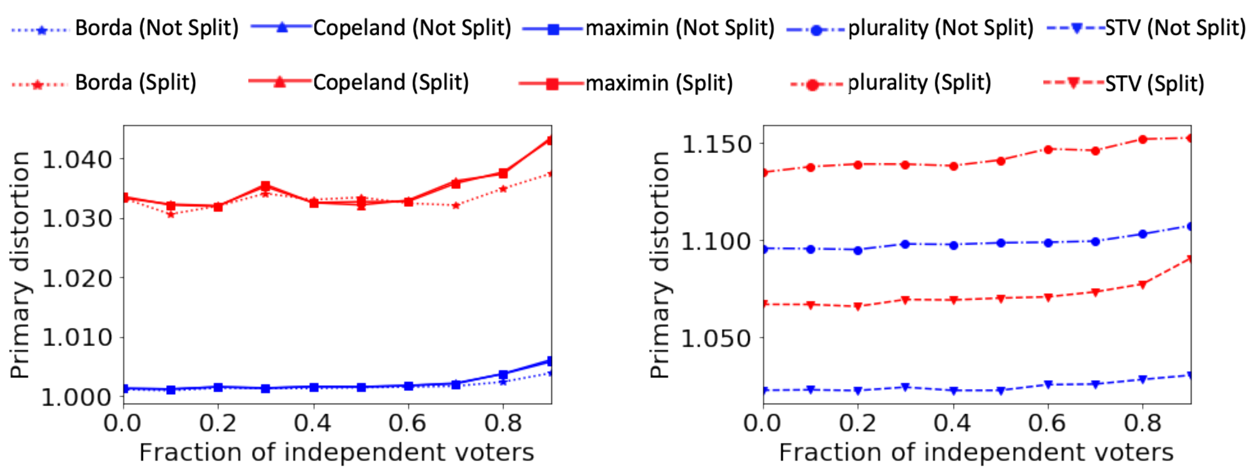


(a) Borda, Copeland and maximin

(b) Plurality and STV

Figure 5.6: The average distortion under the primary system as a function of the percentage of independent voters.

The effect of independent voters is also interesting. With more independent voters, we would expect the average distortion to increase because the independent voters become more important in the general election, but do not get a voice in selecting the primary winners over which they are asked to vote. In Figure 5.6, we see that the average distortion increases in each case; however, the effect is mild except in the extreme region where more than 80% voters are independent.

### The percentage of independent candidates

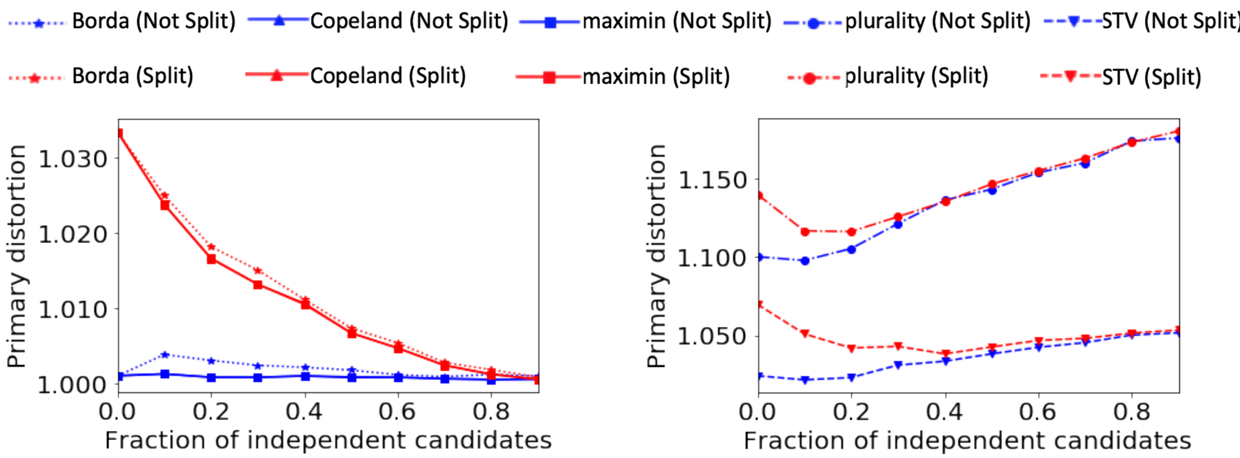


(a) Borda, Copeland and maximin

(b) Plurality and STV

Figure 5.7: The average distortion under the primary system as a function of the percentage of independent candidates.

There are two reasonable interpretations of how independent candidates may impact the average distortion. On the one hand, if the primary winners are desirable candidates, having too many independent candidates in the general election can hurt by overshadowing them. On the other hand, if the primary winners are not desirable, then independent candidates can serve as viable alternatives. In Figure 5.7, we see that this non-trivial effect shows up in the case of plurality and STV. For Borda, Copeland, and maximin, we once again see a dramatic difference between the split and random cases.

### The dimension of the metric space

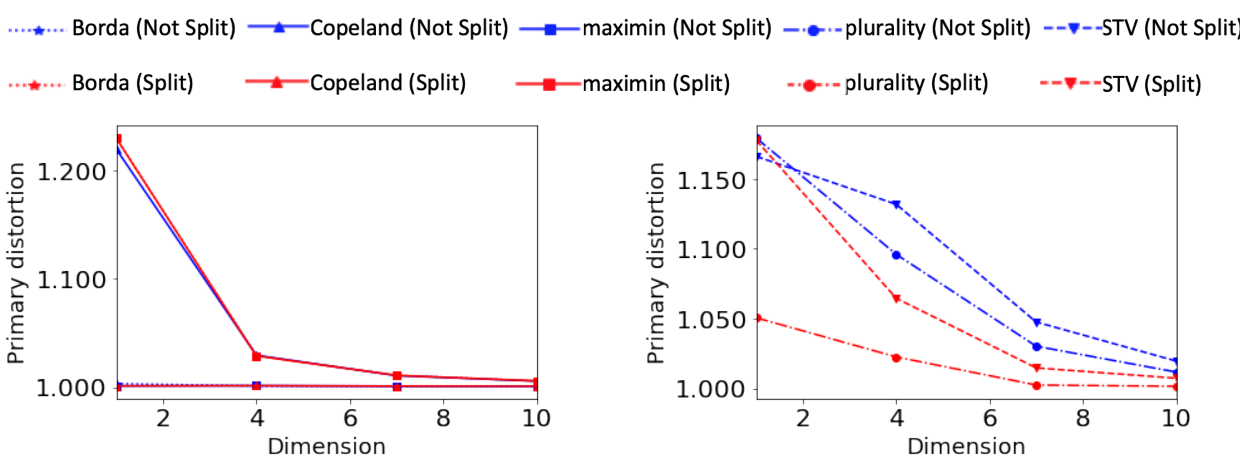


(a) Borda, Copeland and maximin

(b) Plurality and STV

Figure 5.8: The average distortion under the primary system as a function of the dimension of the party embedding.

Perhaps the most consistent pattern across all our simulations was the impact of the dimension the parties were embedded in had on the average primary election distortion. As Figure 5.8 shows as the dimension of the embedding increases the distortion goes down in all. While Figure 5.8 only shows the distortion curve with the default values for each variable we encountered similar patterns in all settings. Increasing the dimension while holding all other variables constant consistently led to lower distortions.

## 5.9 Conclusion

In this chapter we looked at the novel quantitative study of multi-stage elections, also known as primary elections and compared them to single-stage (or direct) elections.

Our contribution was twofold. First, we formulated a model which allows a quantitative comparison of the two voting systems. Our model is a spatial model of voting in which voters and candidates are located in a underlying multi-dimensional space. For our first results we analyzed a two-party setting, in which all voters are affiliated with a party, each party selecting a single candidate, and both candidates are presented to the general voting public. We then extended this setting to include various voters – independent voters (i.e., not party affiliated) and voters which do not participate in the general election – as well as multiple parties and multiple decision stages. Secondly, we use this model to present a comparison of the direct and primary voting systems not in worse-case bounds but in simulations. We showed how different voting systems behave quite differently under primary and direct elections. In particular, we showed plurality generally benefits strongly from using primaries, while Condorcet consistent rules are mainly better off under the direct system. We explore the effects of various parameters (independent voters, party sizes, etc.) on the distortion, and show which settings reduce distortion and which increase it. Furthermore, we saw similarities with various voting rules, often plurality and STV behaved similarly in terms of distortion while Condorcet consistent rules and Borda also tracked each other.

# Chapter 6

# Conclusion

Everyone leaves unfinished business.

Amos Burton
The Expanse

Broadly our work focused on three directions and their impact on the outcome of elections. We can alternatively think of these directions as an examination of different agents who participate in the election. In this sense we can view our work from the perspective of game theory and artificial intelligence.

In chapter 3 we examine the voters themselves. We examine how their preferences and manipulations of the voting mechanism lead to a dynamic game known as iterative voting. We show the settings we study (restricted voter dynamics and district-based elections) lack the convergence properties of simpler settings seen previously. Despite this, through empirical simulation we found most games are well behaved and converge quickly to equilibrium. In addition, we found that when voters engage in this myopic behaviour they often see an increase in various measures of winner quality.

In chapter 4 we examined the role of the political entities which draw election boundaries. These entities when acting in a partisan fashion are known as gerrymanders. For a restricted setting we study the impact of the ever increasing urban-rural divide. To better evaluate the ability to gerrymander we introduced the *gerrymandering power* metric which measures how many seats beyond the proportional vote share a party can take. As the strength of the urban-rural divide increases we found their gerrymandering power quickly tapers off. At extreme levels of division only a few regions prove fruitful for exploitation. We next developed a simulated annealing based method which was able to effectively gerrymander real data. This near optimal method, when compared against 538 and their gerrymandering project, is easily generalizable to a variety of settings. We hope to apply it to a variety of questions and make it available to other researchers.

In chapter 5 we examined the role of the political party and the mechanism they use to pick their candidates. This process is known as the primary process in the United States. We were mainly concerned with the utility of the (non-strategic) voter who interacts with this mechanism. Our main contribution was applying the distortion metric, a type of worst case measure for voter utility under some voting rule, to voters and candidates embedded in a metric space. Comparing the distortion of a primary system to its non-primary analog we were able to provide strong bounds on the difference in distortion for various types of elections. Furthermore, we used simulated data to show the effect on distortion of changing various parameters of the election, such as the number of voters, candidates or what space the political parties occupied. We found some surprising results in terms of distortion; namely certain voting rules behaved very similarly to each other in terms of distortion, plurality and STV were highly correlated while Condorcet consistent rules followed the same pattern as Borda.

## 6.1 Next steps

Here we will briefly examine what we think are potential next steps for research. Firstly, we will talk about extensions contained entirely within each chapter. Afterwards we will look at potential research directions which mix ideas from different chapters.

### 6.1.1 Strategic voting extensions

Here we will look at some potential extensions to the work found in chapter 3. An obvious extension would be more experiments that could be performed with various voting rules and voter dynamic combinations. In section 3.3, where we examined district-based elections, we only focused on the best response dynamic. Adding other voter dynamics, like those we saw in section 3.2, could be interesting as they seemed to have some unintuitive effects. On the topic of dynamics we're curious if there is a natural dynamic that does converge for any voting rule. In subsection 2.3.4 we briefly discussed (what we consider) unnatural dynamics that do converge, but defining what is natural is a subjective question.

A persistent theme we noticed was the influence a Condorcet winner had on the willingness, or more aptly put lack of willingness, of voters to engage in strategic behaviour. There are many more Condorcet consistent rules we could examine so a line of work focusing on this seems interesting.

Finally, the influence of the distribution over preference types on the outcomes is something we only scratched the surface of. In all of chapter 3 we only looked at uniformly random and single-peaked preference distributions. These are a good starting point, as uniformly random ensures all preference types are equally likely and single-peaked ensures a Condorcet winner exists. But there are many more commonly used distributions. One of particular interest is the Mallows-model, where there is a ground truth preference and individual agents randomly sample a nearby preference[1]. We hypothesize that iterative voting would lead to the preferred candidate in the ground truth ranking reemerging as the iterated winner more often than not, but this is only a conjecture.

### 6.1.2 Gerrymandering extensions

Here we will look at some potential extensions to the work found in chapter 4. Our theoretical results in section 4.2 only examine the extreme settings of two districts or an infinite graph. Naturally we would like to find some theoretical results for finite graphs with more than two districts. More general results would be directly transferable to real world settings providing a sort of upper (or lower) bound for optimization techniques. Furthermore, in Theorem 14 we were able to incorporate convex districts into our proof. However, in Theorem 16 our results require non-convex districts. We would be interested in seeing if we could incorporate a formal notion of compactness into our theoretical results. While there are many potential definitions of compactness (often conflicting ones) there is one we are particularly

[1]Nearness is quantified using the Kendal-Tau (or bubble sort) distance.

interested in. Briefly, we define the compactness of a connected component of a graph to be the ratio of boundary nodes (those adjacent to another district) to total nodes of that component. This is a particularly appealing definition of compactness since it is entirely graph theoretical (unlike those which are defined in terms of surface area and perimeter) and can thus be inserted into both our theoretical and empirical work.

Extending on our work in section 4.3 we would like to improve our technique so that in North Carolina we achieve the optimal gerrymandering for the Democrats. Currently we are investigating other optimization techniques and using a larger number of cores. Additionally we'd like to integrate other objectives and constraints into our program. For example we could integrate compactness as both an objective or constraint. To integrate it as an objective we would just take some definition of compactness, such as the one suggested above, and add it to the energy function (Equation 4.5). Now the program would optimize for some mixture, possibly a linear combination, of compact and partisan districts. To integrate it as a constraint one could say any plan where the level of compactness is insufficient will be rejected when proposed (the rejection would occur in Algorithm 2). In some initial experiments we've found integrating compactness in this fashion is a bit restrictive, as this greatly reduces the ability of the Markov chain to move from state to state since many neighbours of a potential state now violate the constraint. Another potential way to integrate compactness would be to allow initially very non-compact states but as Algorithm 2 makes progress the required level of compactness increases until it reaches the desired amount[2]. We are currently working on expanding the code base so that researchers may integrate their own objectives and constraints in the ways we've just described.

One extension in chapter 4 that involves both lines of research is applying Algorithm 2 to the urban-rural model found in section 4.2. The simulated annealing method is clearly well suited for working on large and complex problems, this was an unfortunate limitation of our integer linear programming approach. As a first we could study the urban-rural divide problem on larger grid graphs and multiple urban centres. Another direction would be to modify our voter dispersion model for the urban-rural divide on real data. Perhaps the geographic oddities of real data would lead to different results than those found on grid graphs, this seems like an obvious question worth answering.

### 6.1.3 Primary extensions

Here we will look at some potential extensions to the work found in chapter 5.

Some directions are fairly straightforward extensions of our results. The most straightforward question is to tighten our bounds. Beyond that, it is important to consider whether or not our result on the benefit of primaries over direct elections that hold with respect to a synthetic voting rule could also hold for more common voting rules (e.g., STV) or not. Moreover, our results contrasting separable and non-separable metric spaces might possibly be extended to spaces which are "nearly-separable", or more generally we might consider a suitable parameterized definition of party separation and study how results change as the parameter is changed to force more or less separation. There is also the question of explaining the trends we observe in the average case, which sometimes differ from our worst-case results. In particular, we would like to have a good understanding as to why plurality and STV differ so dramatically (as in the experimental results summarized in Table 5.1) with regard to the party separation model. A next step would be to study realistic distributional models of voter preferences and candidate positions in the political spectrum, and analyze their effect on distortion. Other extensions are seemingly more involved. The multi-stage process offers various directions of exploration. With regard to the multi-party tree model (Definition 13) in section 5.7, what tree-graph structures

[2] The reader may notice this is similar to a simulated annealing method where moves towards non-compact states become less permissible over time.

produce tighter results? What are voting rule combinations that work well together? Examining the use of multiple and different voting rules as [79] does for two-step voting (though without candidate elimination between stages) is an enticing direction. For example, plurality might be used in the general election, while STV might be used by the parties in their primaries. It would be interesting as well to examine manipulations by parties, by candidates, and by voters in primary systems. In particular, it is reasonable to believe that candidates may strategically shift, to some extent, their location following the primaries, to make themselves more appealing to the general electorate. Another topic where more research is needed is investigating multi-winner elections in party elections (e.g., in party lists, in countries where this is common). We believe that the study of multi-stage elections and party mechanisms can not only contribute novel theoretical challenges to tackle, but can also bring research on computational social choice closer to reality and increase its impact.

### 6.1.4 Combining gerrymandering and strategic voting

Here we will look at a potential extension using ideas found in both chapter 4 and chapter 3. In chapter 4 we assumed that there were only two parties and thus strategic voting was not a concern[3]. If we were to increase to multiple parties it would be interesting to see how voter behaviour changes when forced into gerrymandered districts. Would the voter's strategic behaviour adversely effect the gerrymanderer's goals? Could a two-level game be conceived where the gerrymanderer builds districts which are resistant to strategic behaviour? There are many possible settings for this problem. For example, we can use the iterative voting model, or one where many voters may move at once. It is also unclear what level of information the gerrymanderers should have access to. For example, if the voting rule is plurality the gerrymanderer can only learn the top of the submitted preferences.

Another, perhaps less realistic, model could involve voters intentionally moving from district to district in an effort to increase the impact of their vote. While we personally believe very few people would make their voting power the main deciding factor of where they live it is nevertheless an interesting idea proposed in some works [68][4]. Again, to pursue this idea there are several questions one would need to settle. For example, we would need to model how voters move within the region. We could model voters who have a preference for improving their vote impact while not moving too far from their current home. Another option would be a similar approach to a Schelling segregation model [101] where voters prefer to stay near other voters with similar preferences.

### 6.1.5 Combining primaries and strategic voting

Here we will look at a potential extension using ideas found in both chapter 3 and chapter 5.

The primary mechanism proposed in chapter 5 at its core is just a voting rule, thus all of the iterative voting and dynamic questions we examined in section 3.2 can be directly applied to it. There is an interesting balance primary voters may want to consider. Should they behave strategically and vote for a candidate they prefer to the ones in the other party, but not necessarily the one they prefer the most in their own party, or should they pick their most preferred candidate even if they would fare poorly against competition from the opposite party? This balance is not unlike the globally minded and locally minded voter type we introduced in section 3.3, but now it is further complicated by the electability chances when faced with candidates from another party. Strategizing need not be

[3]For two candidates the optimal strategy is always to vote for your most preferred of the two in any voting rules we've considered.

[4]The authors of this article refer to moving to "battleground" states where the young liberal voter would have more power in the electoral college, but nevertheless the idea is the same.

limited to misrepresenting preferences, a crafty voter may partake in the primary for the opposite party in an effort to pick a candidate who would fail in the general election. Partaking in another primary need not even be a malicious activity, perhaps a voter would be fine with any member of their party winning and thus would like to ensure the member of the opposite party is above their acceptable threshold (alternatively they may help an unelectable candidate win the opposing primiary[5]). We've only described these ideas at a high level since we would need to decide how to model these considerations.

[5]This is known as crossover voting. While its effectiveness is questionable [78] it is an intriguing problem since it is allowed in several states.

# Bibliography


[1] AEC. Above the line and below the line voting. `https://www.aec.gov.au/about_aec/research/files/sbps-atl-and-btl-voting.pdf`, 2016. Accessed: 2019-01-01.

[2] M. Altman. The computational complexity of automated redistricting: Is automation the answer. *Rutgers Computer & Technology Law Journal*, 23:81–142, 1997.

[3] E. Anshelevich, O. Bhardwaj, and J. Postl. Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 777–783, Austin, Texas, January 2015.

[4] E. Anshelevich and J. Postl. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research*, 58(1):797–827, January 2017.

[5] N. P. G. at The Smithsonian. Elbridge gerry. `https://npg.si.edu/object/npg_NPG.77.297?fbclid=IwAR1gskWUivhL6OsiQZlZ2JBCQBWm7GwtnGZmFr4FcKHGL17PMz5UMdr-7cE`.

[6] Y. Bachrach, O. Lev, Y. Lewenberg, and Y. 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.

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

[8] 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.

[9] J. J. Bartholdi III, C. A. Tovey, and M. A. Trick. How hard is it to control an election? *Mathematical and Computer Modelling*, 16(8-9):27–40, August-September 1992.

[10] S. Bervoets and V. Merlin. Gerrymander-proof representative democracies. *International Journal of Game Theory*, 41:1–16, 08 2012.

[11] S. Bervoets and V. Merlin. On avoiding vote swapping. *Social Choice and Welfare*, 46:495–509, 2016.

[12] B. Bishop. *The Big Sort: Why the Clustering of Like-Minded America is Tearing Us Apart*. Mariner Books, 2009.

[13] D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56:23–34, 1948.

[14] A. Borodin, O. Lev, N. Shah, and T. 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.

[15] A. Borodin, O. Lev, N. Shah, and T. Strangway. Primarily about primaries. In *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI)*, Honolulu, USA, January-February 2019.

[16] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.

[17] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, March 2016.

[18] 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.

[19] M. Brill and V. Conitzer. Strategic voting and strategic candidacy. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 819–826, Austin, Texas, January 2015.

[20] G. Brown and R. Dell. Formulating integer linear programs: A rogues' gallery. *Informs Transactions on Education*, 7:153–159, 01 2007.

[21] E. Canada. History of representation in the house of commons of canada. `https://www.elections.ca/content.aspx?section=res&dir=cir/red/book&document=rep2&lang=e`, 2012. Accessed: 2020-01-01.

[22] J. Chen and J. Rodden. Tobler's law, urbanization, and electoral bias: Why compact, contiguous districts are bad for the democrats, 2009.

[23] J. Chen and J. Rodden. Unintentional gerrymandering: Political geography and electoral bias in legislatures. *Quarterly Journal of Political Science*, 8:239–269, 2013.

[24] M. Cohen, D. Karol, H. Noel, and J. Zaller. *The Party Decides: Presidential Nominations Before and After Reform*. University of Chicago Press, 2008.

[25] A. Cohen-Zemach, Y. Lewenberg, and J. S. Rosenschein. Gerrymandering over graphs. In *Proceedings of The 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (IJCAI)*, Stockholm, Sweden, July 2018.

[26] G. Cohensius, S. Mannor, R. Meir, E. Meirom, and A. Orda. Proxy voting for better outcomes. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 858–866, São-Paulo, Brazil, May 2017.

[27] V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. In *Proceedings of the 25th National Conference on Artificial Intelligence (AAAI)*, pages 638–643, San Francisco, California, August 2011.

[28] W. P. Cross and A. Blais. Who selects the party leader? *Party Politics*, 18(2):127–150, January 2011.

[29] W. P. Cross and R. S. Katz, editors. *The Challenges of Intra-Party Democracy*. Oxford University Press, 2013.

[30] P. Dunleavy and R. Diwakar. Analysing multiparty competition in plurality rule elections. *Party politics*, 19:1–32, 10 2011.

[31] B. Dutta, M. O. Jackson, and M. Le Breton. Strategic candidacy and voting procedures. *Econometrica*, 69(4):1013–1037, July 2001.

[32] B. Dutta, M. O. Jackson, and M. Le Breton. Voting by successive elimination and strategic candidacy. *Journal of Economic Theory*, 103(1):190–218, March 2002.

[33] M. Dyer and A. Frieze. On the complexity of partitioning graphs into connected subgraphs. *Discrete Applied Mathematics*, 10:139–153, 1985.

[34] 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.

[35] H. Enten. Ending gerrymandering won't fix what ails america. *FiveThirtyEight*, 26 January 2018.

[36] G. Erdélyi, E. Hemaspaandra, and L. A. Hemaspaandra. More natural models of electoral control by partition. In T. Walsh, editor, *Algorithmic Decision Theory*, pages 396–413, Cham, 2015. Springer International Publishing.

[37] FEC. Federal election commission of the united states. `https://www.fec.gov/about/`, 2020. Accessed: 2020-01-01.

[38] M. Feldman, A. Fiat, and I. Golomb. On voting and facility location. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 269–286, Maastricht, The Netherlands, July 2016.

[39] J. A. Ferejohn and M. P. Fiorina. The paradox of not voting: A decistion theoretic analysis. *American Political Science Review*, 68:525–536, June 1974.

[40] B. Fifield, M. Higgins, K. Imai, and A. Tarr. A new automated redistricting simulator using markov chain monte carlo, January 2018.

[41] P. C. Fishburn. Paradoxes of voting. *The American Political Science Review*, pages 537–546, 1974.

[42] P. C. Fishburn and W. V. Gehrlein. Borda's rule, positional voting, and condorcet's simple majority principle. *Public Choice*, 28(1):79–88, 1976.

[43] S. D. Fisher. Definition and measurement of tactical voting: the role of rational choice. *British Journal of Political Science*, 34:152?166, 2004.

[44] W. V. Gehrlein. The impact of social homogeneity on the condorcet efficiency of weighted scoring rules. *Social Science Research*, 16(4):361–369, 1987.

[45] W. V. Gehrlein and F. Plassmann. A comparison of theoretical and empirical evaluations of the borda compromise. *Social Choice and Welfare*, 43(3):747–772, 2014.

[46] W. V. Gehrlein and F. Valognes. Condorcet efficiency: A preference for indifference. *Social Choice and Welfare*, 18(1):193–205, 2001.

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

[48] 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.

[49] B. Grofman and G. King. The future of partisan symmetry as a judicial test for partisan gerrymandering after lulac v. perry. *Election Law Journal*, 6(1):2–35, 2007.

[50] A. Hamilton, J. Madison, and J. Jay. *The Federalist Papers*. The Independent Journal, 1787.

[51] D. Hartvigsen. Vote trading in public elections. *Mathematical Social Sciences*, 52(1):31 – 48, 2006.

[52] R. Y. Hazan. The 1996 intra-party elections in Israel: Adopting party primaries. *Electoral Studies*, 16(1):95–103, March 1997.

[53] R. Jobson and M. Wickham-Jones. Gripped by the past: Nostalgia and the 2010 labour party leadership contest. *British Politics*, 5(4):525–548, December 2010.

[54] J. M. Jones. On social ideology, the left catches up to the right. Gallup, MAY 22 2015.

[55] A. Kahng, S. Mackenzie, and A. D. Procaccia. Liquid democracy: An algorithmic perspective. In *Proceedings of the 32nd Conference on Artificial Intelligence (AAAI)*, pages 1095–1102, New Orleans, Louisiana, February 2018.

[56] R. N. Karasev. Equipartition of several measures. arXiv:1011.4762, 2010.

[57] O. Kenig. Classifying party leaders' selection methods in parliamentary democracies. *Journal of Elections, Public Opinion and Parties*, 19(4):433–447, October 2009.

[58] M. J. Kim. *Optimization Approaches to Political Redistricting Problems*. PhD thesis, Ohio State University, 2011.

[59] B. Klaas. Gerrymandering is the biggest obstacle to genuine democracy in the united states. so why is no one protesting? *Washington Post*, 10 February 2017.

[60] A. Koolyk, T. Strangway, O. Lev, and J. S. Rosenschein. Convergence and quality of iterative voting under non-scoring rules. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 273–279, Melbourne, Australia, August 2017.

[61] J. Lang, N. Maudet, and M. Polukarov. New results on equilibria in strategic candidacy. In B. Vöcking, editor, *Algorithmic Game Theory*, pages 13–25, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[62] 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.

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

[64] O. Lev and T. Strangway. Iterative voting in district-based elections. In *Games and Information Workshop (GAIW)*, Montreal, Canada, May 2019.

[65] J. Levitt. Where are the lines drawn? `https://redistricting.lls.edu/where.php/`.

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

[67] 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`.

[68] A. MacGillis. Go midwest, young hipster. New York Times, October 23 2016.

[69] K. C. Martis. The original gerrymander. *Political Geography*, 27(8):833 – 839, 2008.

[70] N. Mattei and T. Walsh. Empirical evaluation of real world tournaments. *CoRR*, abs/1608.01039, 2016.

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

[72] R. Meir. Iterative voting. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 4, pages 69–86. AI Access, 2017.

[73] 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.

[74] 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.

[75] R. Meir, M. Polukarov, J. S. Rosenschein, and N. 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.

[76] MGGG. Metric geometry and gerrymandering group (mggg). `https://mggg.org`, 2020. Accessed: 2020-01-01.

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

[78] J. Nagler and R. Michael Alvarez. *Analysis of Crossover and Strategic Voting (expert witness report)*, volume Bysiewicz. California Party v. Bill Jones, 6 1997.

[79] N. Narodytska and T. Walsh. Manipulating two stage voting rules. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 423–430, St. Paul, Minnesota, May 2013.

[80] J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, September 1951.

[81] New York Times Editorial Board. A shot at fixing american politics. *New York Times*, page A22, 30 September 2017.

[82] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[83] R. Noack. The urban-rural divide that bolstered trump isn't just an american thing; it's prevalent in europe, too. *Washington Post*, 27 November 2016.

[84] H. Norpoth. From primary to general election: A forecast of the presidential vote. *PS: Political Science & Politics*, 37(4):737–740, October 2004.

[85] Y. Nourani and B. Andresen. A comparison of simulated annealing cooling strategies. *Journal of Physics A: Mathematical and General*, 31(41):8373–8385, oct 1998.

[86] H. Nurmi. *Voting Paradoxes and How to Deal with Them*. Springer-Verlag, 1999.

[87] 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.

[88] W. Pegden, A. D. Procaccia, and D. Yu. A partisan districting protocol with provably nonpartisan outcomes. ArXiv:1710.08781, October 2017.

[89] M. Polukarov, S. Obraztsova, Z. Rabinovich, A. Kruglyi, and N. R. Jennings. Convergence to equilibria in strategic candidacy. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI)*, pages 624–630, Buenos Aires, Argentina, July 2015.

[90] A. D. Procaccia and J. S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, pages 317–331, Edinburgh, Great Britain, September 2006.

[91] C. Puppe and A. Tasnádi. A computational approach to unbiased districting. *Mathematical and Computer Modelling*, 48(9–10):1455–1460, 2008. Mathematical Modeling of Voting Systems and Elections: Theory and Applications.

[92] 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.

[93] 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.

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

[95] T. C. P. Report. The cook political report. `https://cookpolitical.com/`, 2018. Accessed: 2020-01-01.

[96] 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.

[97] J. Rodden. *Why Cities Lose: The Deep Roots of the Urban-Rural Political Divide*. New York, Basic Books, 2019.

[98] R. Rumi. 2016 election explainer in 4 words: The urban-rural divide. *HuffPost*, 6 December 2017.

[99] D. G. Saari. The optimal ranking method is the borda count. Discussion Paper 638, Northwestern University, Center for Mathematical Studies in Economics and Management Science, February 1985.

[100] 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, April 1975.

[101] T. C. Schelling. Models of segregation. *The American Economic Review*, 59(2):488–493, 1969.

[102] N. Schofield. *The Spatial Model of Politics*. Number 95 in Routledge Frontiers of Political Economy. Routledge, 2008.

[103] J. Sides, C. Tausanovitch, L. Vavreck, and C. Warshaw. On the representativeness of primary electorates. *British Journal of Political Science*, pages 1–9, March 2018.

[104] N. Silver. 538: The gerrymandering project. `https://fivethirtyeight.com/tag/the-gerrymandering-project/`, 2020. Accessed: 2020-01-01.

[105] P. Skowron and E. Elkind. Social choice under metric preferences: Scoring rules and STV. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 706–712, San Francisco, California, February 2017.

[106] P. Soberón. Balanced convex partitions of measures in rd. *Mathematika*, 58(1):71–76, 2012.

[107] 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.

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

[109] D. Wasserman. Hating gerrymandering is easy. fixing it is harder. *FiveThirtyEight*, 25 January 2018.

[110] B. Wauters. Explaining participation in intra-party elections: Evidence from belgian political parties. *Party Politics*, 16(2):237–259, September 2010.

[111] S. Wilentz. *The Rise of American Democracy*. Norton, 2005.

[112] 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.

[113] J. Yang. Some np-complete edge packing and partitioning problems in planar graphs. ArXiv:1409.2426, September 2014.

# Chapter 7

# Appendix

> “How do you know I’m mad?” said Alice.
> “You must be,” said the Cat, “or you wouldn’t have come here.”
>
> Lewis Carroll’s Alice in Wonderland

In this chapter we will include supplementary material that was excluded from the main body.

## 7.1 Missing proofs for strategic voting

In chapter 3 we did not include all of the proofs of non-convergence in subsection 3.2.1 since many are similar and do not add to the understanding of the material. Here we include them for completeness (this includes ones already presented in the main body). We again provide (the same) note on reading the examples of non convergence: 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 (after ties are broken). The i-th row 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.

### 7.1.1 Maximin

**Theorem 30.** *Maximin with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT and SWAP.*

*Proof.* For BR, the example is:

$$\begin{tikzcd}
BAC & BAC \arrow[r] & CBA & CBA \arrow[lll, bend right] \\
CAB \arrow[r] & ABC & ABC \arrow[r] & CAB \\
CAB & CAB & CAB & CAB \\
BCA & BCA & BCA & BCA \\
\mathbf{B} & \mathbf{A} & \mathbf{B} & \mathbf{C}
\end{tikzcd}$$

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:

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

For KT, the example is:

$$\begin{array}{ccccccc}
DBACE & \longrightarrow & BDCAE & & BDCAE & & BDCAE \\
EDCAB & & EDCAB & \to & DECAB & & DECAB \\
CEADB & & CEADB & & CEADB & \to & EACBD \\
\mathbf{A} & & \mathbf{B} & & \mathbf{D} & & \mathbf{A} \\
 & & & & & & \\
 & & DBCAE & & DBCAE & & DBCAE \\
 & & EDCAB & & EDCAB & \leftarrow & DECAB \\
 & & CEADB & \leftarrow & EACBD & & EACBD \\
 & & \mathbf{C} & & \mathbf{E} & & \mathbf{D}
\end{array}$$

And, finally, for SWAP, the example is:

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

□

### 7.1.2 Copeland

**Theorem 31.** *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 $\alpha$.*

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

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:

$$\begin{tikzcd}[row sep=tiny]
DABC & DABC \arrow[r] & ACBD & ACBD \arrow[lll, bend right] \\
BDAC \arrow[r] & BACD & BACD \arrow[r] & BDAC \\
CDBA & CDBA & CDBA & CDBA \\
\mathbf{D} & \mathbf{B} & \mathbf{A} & \mathbf{B}
\end{tikzcd}$$

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:

$$\begin{tikzcd}[row sep=tiny]
ACB & ACB \arrow[r] & ABC & ABC \arrow[lll, bend right] \\
BCA \arrow[r] & CBA & CBA \arrow[r] & BCA \\
\mathbf{A} & \mathbf{C} & \mathbf{A} & \mathbf{B}
\end{tikzcd}$$

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:

$$\begin{tikzcd}[row sep=tiny]
BCA & BCA \arrow[r] & CBA & CBA \arrow[lll, bend right] \\
ABC \arrow[r] & ACB & ACB \arrow[r] & ABC \\
\mathbf{B} & \mathbf{A} & \mathbf{C} & \mathbf{A}
\end{tikzcd}$$

□

### 7.1.3 Bucklin

**Theorem 32.** *Bucklin with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT and SWAP.*

*Proof.* The example for BR:

$$\begin{tikzcd}[row sep=tiny]
ACB & ACB \arrow[r] & CBA & CBA \arrow[lll, bend right] \\
BAC \arrow[r] & CBA & CBA \arrow[r] & BAC \\
ABC & ABC & ABC & ABC \\
CBA & CBA & CBA & CBA \\
A & \mathbf{B} & \mathbf{C} & B
\end{tikzcd}$$

The example for TOP:

$$\begin{tikzcd}[row sep=tiny, column sep=small]
BACD & BACD \arrow[r] & CBAD & CBAD \arrow[r] & ADBC & ADBC \arrow[lllll, bend right] \\
DBAC \arrow[r] & DCAB & DCAB \arrow[r] & DABC & DABC \arrow[r] & DBAC \\
BDAC & BDAC & BDAC & BDAC & BDAC & BDAC \\
CDBA & CDBA & CDBA & CDBA & CDBA & CDBA \\
ADCB & ADCB & ADCB & ADCB & ADCB & ADCB \\
\mathbf{B} & \mathbf{D} & \mathbf{C} & \mathbf{D} & \mathbf{A} & \mathbf{D}
\end{tikzcd}$$

The example for the TB dynamic:

$$
\begin{tikzcd}
ADBC & ADBC \arrow[r] & ACBD & ACBD \arrow[lll, bend right] \\
CBDA \arrow[r] & DBCA & DBCA \arrow[r] & CBDA \\
CADB & CADB & CADB & CADB \\
DBAC & DBAC & DBAC & DBAC \\
\mathbf{A} & \mathbf{D} & \mathbf{A} & \mathbf{C}
\end{tikzcd}
$$

The example for the KT dynamic:

$$
\begin{tikzcd}
ABCD & ABCD \arrow[r] & ACBD & ACBD \arrow[lll, bend right] \\
DCBA \arrow[r] & DBCA & DBCA \arrow[r] & DCBA \\
CBAD & CBAD & CBAD & CBAD \\
DACB & DACB & DACB & DACB \\
\mathbf{A} & \mathbf{B} & \mathbf{A} & \mathbf{C}
\end{tikzcd}
$$

The example for the SWAP dynamic:

$$
\begin{tikzcd}
DCBA & DCBA \arrow[r] & DBCA & DBCA \arrow[lll, bend right] \\
ABCD \arrow[r] & ACBD & ACBD \arrow[r] & ABCD \\
CDAB & CDAB & CDAB & CDAB \\
BDAC & BDAC & BDAC & BDAC \\
\mathbf{D} & \mathbf{C} & \mathbf{D} & \mathbf{B}
\end{tikzcd}
$$

$\square$

### 7.1.4 STV

**Theorem 33.** *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:

$$
\begin{tikzcd}
BCDA \arrow[r] & DCAB & DCAB \arrow[r] & BADC & BADC \arrow[lll, bend right] \\
ADBC & ADBC \arrow[r] & DBCA & DBCA \arrow[r] & ADBC \\
CDAB & CDAB & CDAB & CDAB & CDAB \\
\mathbf{A} & \mathbf{C} & \mathbf{D} & \mathbf{B} & \mathbf{A}
\end{tikzcd}
$$

The example for the TOP dynamic:

$$
\begin{tikzcd}
CDAB \arrow[r] & DABC & DABC \arrow[r] & CABD & CABD \arrow[lll, bend right] \\
ABCD & ABCD \arrow[r] & BACD & BACD \arrow[r] & ABCD \\
DACB & DACB & DACB & DACB & DACB \\
CBDA & CBDA & CBDA & CBDA & CBDA \\
\mathbf{A} & \mathbf{D} & \mathbf{B} & \mathbf{C} & \mathbf{A}
\end{tikzcd}
$$

The example for the TB dynamic:

$$
\begin{tikzcd}
CABD & CABD \arrow[r] & ACDB & ACDB \arrow[lll, bend right] \\
DBCA \arrow[r] & BADC & BADC \arrow[r] & DBCA \\
DBCA & DBCA & DBCA & DBCA \\
CDAB & CDAB & CDAB & CDAB \\
\mathbf{C} & \mathbf{B} & \mathbf{A} & D
\end{tikzcd}
$$

The example for the KT dynamic:

$$
\begin{tikzcd}[row sep=tiny]
DBAC & DBAC \arrow[r] & BDAC & BDAC \arrow[lll, bend right] \\
ACBD \arrow[r] & CABD & CABD \arrow[r] & ACBD \\
CDBA & CDBA & CDBA & CDBA \\
DABC & DABC & DABC & DABC \\
\mathbf{D} & \mathbf{C} & \mathbf{B} & \mathbf{A}
\end{tikzcd}
$$

The example for the SWAP dynamic:

$$
\begin{tikzcd}[row sep=tiny]
ADCB & ADCB \arrow[r] & DACB & DACB \arrow[lll, bend right] \\
DBAC \arrow[r] & BDAC & BDAC \arrow[r] & DBAC \\
CBDA & CBDA & CBDA & CBDA \\
\mathbf{A} & \mathbf{B} & \mathbf{C} & \mathbf{D}
\end{tikzcd}
$$

□

### 7.1.5 Second Order Copeland

**Theorem 34.** *SOC with linear order tie-breaking does not converge for the dynamics BR, TOP, TB, KT and SWAP.*

*Proof.* The example for BR is:

$$
\begin{tikzcd}[row sep=tiny]
BCDA & BCDA \arrow[r] & CBDA & CBDA \arrow[lll, bend right] \\
DABC \arrow[r] & DACB & DACB \arrow[r] & DABC \\
\mathbf{B} & \mathbf{D} & \mathbf{C} & \mathbf{A}
\end{tikzcd}
$$

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

### 7.1.6 Ranked Pairs

**Theorem 35.** *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:

$$
\begin{tikzcd}[row sep=tiny]
BDCA & BDCA \arrow[r] & DACB & DACB \arrow[r] & CBAD & CBAD \arrow[lllll, bend right] \\
CBAD \arrow[r] & CDAB & CDAB \arrow[r] & BACD & BACD \arrow[r] & CDAB \\
\mathbf{B} & \mathbf{C} & \mathbf{D} & \mathbf{A} & \mathbf{B} & \mathbf{C}
\end{tikzcd}
$$

The example for the TOP dynamic:

$$
\begin{tikzcd}[row sep=tiny]
ADBC & ADBC \arrow[r] & ACBD & ACBD \arrow[lll, bend right] \\
CDAB \arrow[r] & DABC & DABC \arrow[r] & CDAB \\
BCDA & BCDA & BCDA & BCDA \\
\mathbf{A} & \mathbf{D} & \mathbf{A} & \mathbf{C}
\end{tikzcd}
$$

The example for the TB dynamic:

$$
\begin{tikzcd}[row sep=tiny]
DABEC & DABEC \arrow[r] & BADEC & BADEC \arrow[lll, bend right] \\
CADEB \arrow[r] & CABED & CABED \arrow[r] & CADEB \\
ECADB & ECADB & ECADB & ECADB \\
BDECA & BDECA & BDECA & BDECA \\
\mathbf{D} & \mathbf{C} & \mathbf{B} & \mathbf{C}
\end{tikzcd}
$$

The example for the KT dynamic:

$$
\begin{tikzcd}[row sep=tiny]
DABC & DABC & DABC & DABC \arrow[r] & ADBC & ADBC \arrow[lllll, bend right] \\
BDCA & BDCA \arrow[r] & BDAC & BDAC & BDAC \arrow[r] & BDCA \\
CABD \arrow[r] & CADB & CADB \arrow[r] & CABD & CABD & CABD \\
CBDA & CBDA & CBDA & CBDA & CBDA & CBDA \\
\mathbf{B} & \mathbf{C} & \mathbf{D} & \mathbf{B} & \mathbf{A} & \mathbf{C}
\end{tikzcd}
$$

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

$\square$

## 7.2 Gerrymandering

In this section we will go over the missing components from chapter 4.

### 7.2.1 The origin of gerrymandering

As we have mentioned in the thesis, the word gerrymandering is a portmanteau. It is named after Elbridge Gerry and the salamander shaped district he approved. As the governor of Massachusetts, Gerry signed into law a redistricting plan which was designed to protect the Democratic-Republican party in the state senate. A portrait of Gerry and a widely shared political cartoon of the districting can be seen in Figure 7.1. There has never been conclusive evidence on who first coined the term gerrymandering, but it was widely used in 1812 [69].

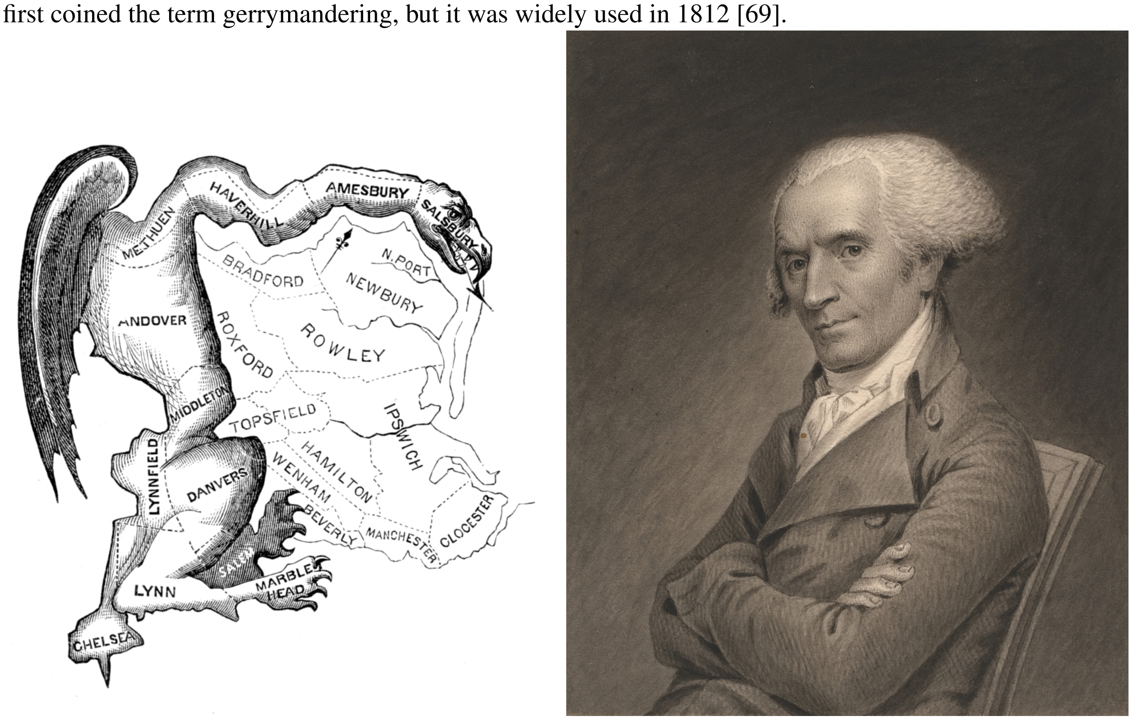


Figure 7.1: On the left a political cartoon of the first gerrymander. The cartoon was likely create by Elkanah Tisdale. His original wood carvings are stored in the Library of Congress [69]. On the right a portrait of Elbridge Gerry. Gerry served as the governor of Massachusetts from 1810 to 1812 and later as the Vice President of the United States, serving from 1813 until his death in 1814. Portrait by James Barton Longacre, based on a work by John Vanderlyn (for further information on this art work see [5]).

### 7.2.2 Optimal integer linear program for gerrymandering

Here we will provide the optimal integer linear programming that is used to recombine an arbitrary number of districts in section 4.2.2. Some quick notation reminders. Our nodes are embedded in a graph $G = (V, E)$, and we want to partition them into $K$ districts.

To ensure we have connected districts we will be utilizing network flows. Recall, in a network flow problem each node in $V$ is connected to a special source node $s$ and a special terminal node $t$. From $s$ we send out fractional values, the flow, which is sent between the nodes $V$ along the edges $E$ until it all reaches the terminal node $t$. The main constraint in flow problems is the *conservation of flow*, it states if a node has some units of flow incoming it must send the exact same amount of flow out. This constraints ensures all the flow sent from $s$ will arrive and terminate at $t$. We will show how to model a flow problem using constraints in our optimal integer linear program. To gerrymander for the urban party we use the following program:

$$
\begin{aligned}
\max \quad & \sum_{k\in[K]} w_k && \\
\text{s.t.} \quad & \sum_{v\in nbhd(u)\cup\{s\}} f_{(v,u)} = \sum_{v\in nbhd(u)\cup\{t\}} f_{(u,v)}, && u \in V \quad \text{(Conservation of flow)}, \\
& I_{(u,v)} \iff (f_{(u,v)} > 0), && (u,v) \in E \quad \text{(Flow indicator)}, \\
& \sum_{k\in[K]} x_{(u,k)} = 1, && u \in V \quad \text{(Each precinct node is in exactly one district)}, \\
& \sum_{v\in V} I_{(s,v)} = K, && \quad \text{(Source sends flow to } K \text{ nodes)}, \\
& \sum_{u\in V} I_{(s,u)} \wedge x_{(u,k)} = 0, && k \in [K] \quad \text{(Source sends flow to exactly one node per district)}, \\
& I_{(u,v)} \implies (\sum_{k\in[K]} x_{(u,k)} \wedge x_{(v,k)} = 1), && (u,v) \in E \quad \text{(If u sends flow to v then they are in the same district)}, \\
& 1-\delta \le \sum_{u\in V} n_u \cdot x_{(u,k)}/NK \le 1+\delta, && k \in [K] \quad \text{(Population constraint)}, \\
& (\sum_{u\in V} (\alpha_u^U - \alpha_u^R) \cdot x_{(u,k)} > 0) \iff w_k, && k \in [K] \quad (w_k \text{ is 1 if and only if the urban party has a majority in district } k)
\end{aligned}
$$

Note, we are using shorthand for the boolean operations, and, or, and implication (see [20] for how to convert these operation to constraints in the program). Each of these operations can be expressed as several integer and linear constraints. Recall $\alpha_u^P$ is the number of votes for party $P$ in node $u$, $\delta$ is our population slack (for our experiments this is taken to be 0), $K$ is the number of districts, $n_u$ is the population of node $u$ and $N$ is the overall population. The meaning of the variables are as follows:

1. The $w_k$ variable is a binary variable which indicates if the urban party won district $k$.
2. The $I_{(u,v)}$ variable is a binary variable which indicates if there is flow along the edge $(u,v)$.
3. The $x_{(u,k)}$ variable is a binary variable which indicates district $k$ contains node $u$.
4. The $f_{(u,v)}$ variable is a fractional variable in $[0,1]$ indicating the amount of flow on edge $(u,v)$.

In addition to enforcing conservation of flow our constraints ensure:

1. Two nodes being in the same district *iff* there is a path of flow from $s$ to $t$ which passes through both nodes.
2. The source activates exactly $k$ "root" nodes for each district. Each of these root nodes in turn forward that flow through the nodes which form its district on the way to the sink.
3. There is no flow between nodes in different districts.

These additional constraints ensure that tracing the flow from source to sink is equivalent to a partition of the graph into contiguous and non-overlapping districts. The program can also be adapted to gerrymander for the rural party. We can either change it from a maximization to minimization problem or change the meaning of the variable $w_k$ (along with the final constraint) to mean the rural party wins district $k$.

### 7.2.3 Comparing starting points

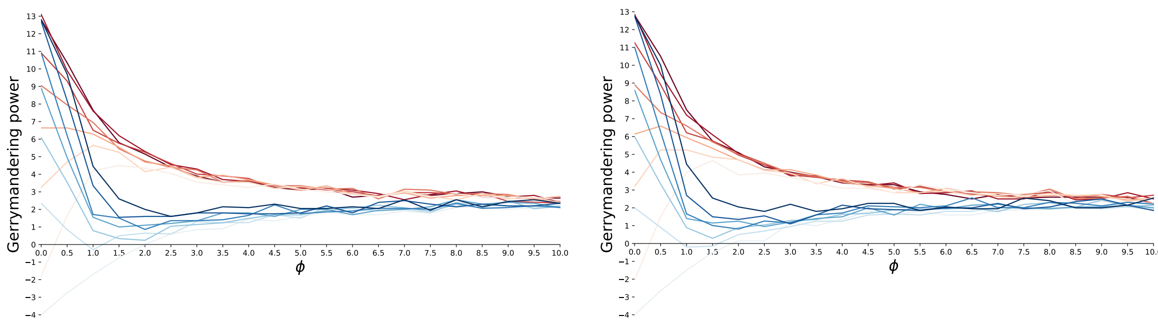


Figure 7.2: A comparison of starting points from the hand drawn rows based distrcits (the left image) and starting points from the tree-based algorithm (the right image). Each figure shows the average gerrymandering power of the parties versus the strength of the urban-rural divide ($\phi$). The urban/rural party is in blue/red, and a darker colour represents a higher vote share of the gerrymandering party.

In section 4.2.2 we mentioned for our simulations of the urban-rural divide we used two types of initial districtings for our maps (each map was a $16 \times 16$ grid graph). The first involved dividing each row into two districts, that is an 8 node line graph. The second was a random tree-based starting point generated from Algorithm 1. In either case for each combination of urban-rural divide strength and urban vote advantage before passing our starting point to our optimization procedure we scrambled it further by using an iterative process which swaps a pair of nodes between adjacent districts $100,000$ times (for full details of this process see section 4.2.2).

For each combination of starting point type (tree-based or divided along rows), strength of urban-rural divide and urban vote advantage we generate twenty maps to optimize over. Thus for each combination of urban-rural divide strength and urban vote advantage we are optimizing over 40 different starting maps. The results shown in Figure 4.3 are the averages over these 40 maps. If we instead plot based on starting point, that is now we're averaging over 20 maps, we see there is effectively no discernible difference (Figure 7.2). We still have the pattern of increasing the urban-rural divide ($\phi$) leads to a convergence of gerrymandering power for each party. We also see the Rural party still tends to have an advantage in gerrymandering power for most values of $\phi$.

From a quantitative standpoint there is effectively no difference between the two groups in Figure 7.2. For the rural party on average the gerrymandering power was $3.8$ and $3.9$ districts for the row and tree-based starting points respectively. The largest deviation between the averages in equal settings for the rural gerrymandering power (that is over the $20$ maps for a fixed level of $\phi$ and $\alpha^U$) was $0.55$ districts. For the urban party on average the gerrymandering power was $2.1$ districts for both starting points. The largest deviation between the averages for urban gerrymandering power in equal settings (that is over the 20 maps for a fixed level of $\phi$ and $\alpha^U$) was $0.65$ districts.

### 7.2.4 Figures for change in gerrymandering power

In this section we will provide the missing figures for section 4.2. In the main body we provided two figures showing the change in gerrymandering power for various levels of the urban rural divide. The previous two figures were chosen since they show how gerrymandering power is influenced by the urban rural divide for competitive elections and for unbalanced elections. For completeness we will present all of the figures (including the previously shown two). For formatting reasons we present our plots in two groups. The first groups is when the urban party does not have a vote advantage (Figure 7.3). The second group is for when the rural party does not have a vote advantage (Figure 7.4).

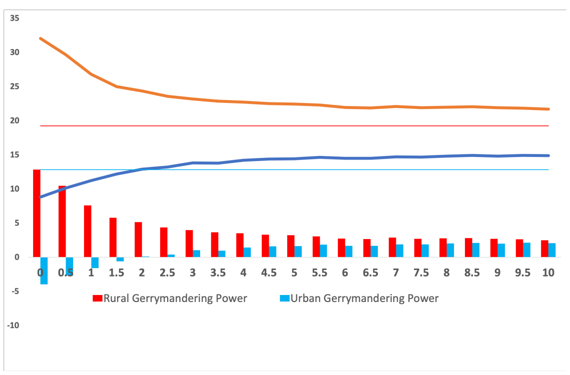


(a) $\alpha^U = 0.40$

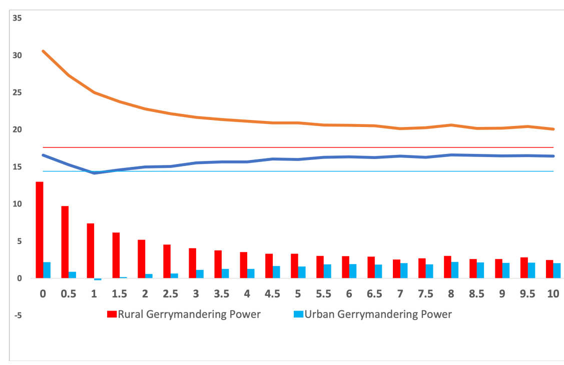


(b) $\alpha^U = 0.45$

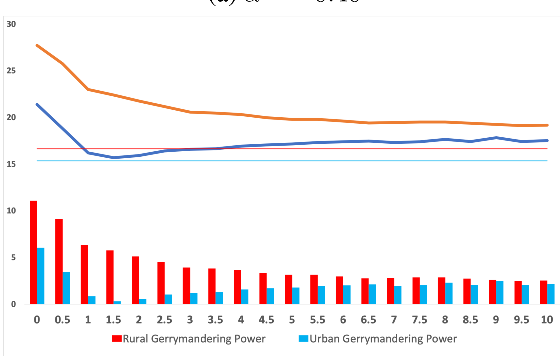


(c) $\alpha^U = 0.48$

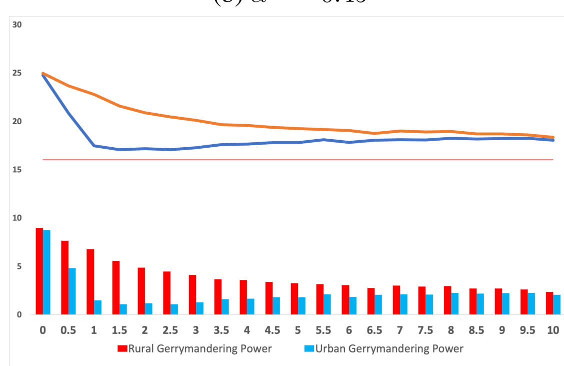


(d) $\alpha^U = 0.50$

Figure 7.3: The number of districts won (top) and gerrymandering power (bottom) compared to $\phi$ for the cases where the urban party does not have a vote advantage (including no vote advantage $\alpha^U = 0.50$). The horizontal blue/red (in the final figure only one is visible since they overlap) line represents how many seats the urban/rural party would win in an entirely proportional system. The blue/brown line represents how many seats the urban/rural party won when they were gerrymandering. The blue/red bar graph at the bottom represents the gerrymandering power of the urban/rural party.

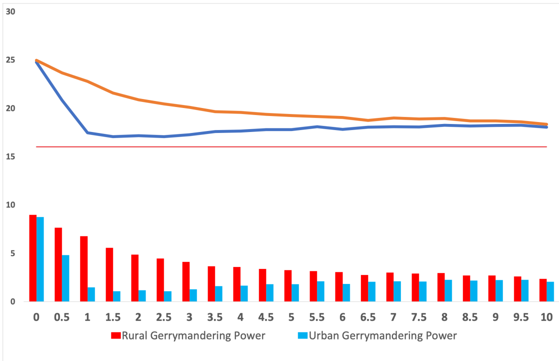


(a) $\alpha^U = 0.50$

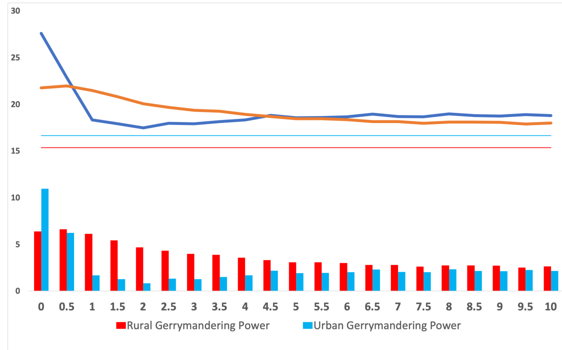


(b) $\alpha^U = 0.52$

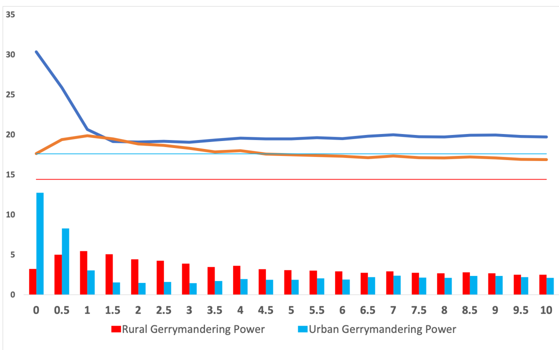


(c) $\alpha^U = 0.55$

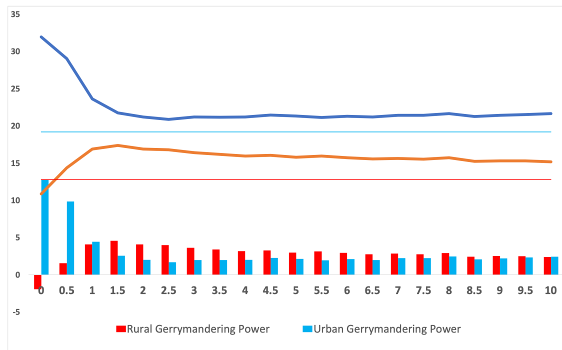


(d) $\alpha^U = 0.60$

Figure 7.4: The number of districts won (top) and gerrymandering power (bottom) compared to $\phi$ for the cases where the rural party does not have a vote advantage (including no vote advantage $\alpha^U = 0.50$). The horizontal blue/red (in the first figure only one is visible since they overlap) line represents how many seats the urban/rural party would win in an entirely proportional system. The blue/brown line represents how many seats the urban/rural party won when they were gerrymandering. The blue/red bar graph at the bottom represents the gerrymandering power of the urban/rural party.

### 7.2.5 538 gerrymandering model

In this section we will cover the missing details of the 538 model for gerrymandering found in subsection 4.3.1. Recall, for any group of voters the 538 model aggregated their historic behaviour and outputted the probability those voters would elect a democrat (equivalently the probability they would not elect a republican). To compare our algorithm against the 538 optimal results we needed to understand what model 538 used. Unfortunately, from the information they published this was not immediately obvious. While we were able to rebuild what 538 did we left some of the details out of the main body. Here we will cover those missing elements.

#### The Cook PVI

Recall, the Cook Partisan Voting Index (Cook PVI), referenced in section 4.3.1, is a measure of how partisan a group of voters is. It is given by the form:

$$100 \cdot \left(\frac{\frac{N_1^D}{N_1^D+N_1^R} + \frac{N_2^D}{N_2^D+N_2^R}}{2} - \beta_D\right) \tag{7.1}$$

Where $N_i^D$ and $N_i^R$ are the number of votes this group of voters gave to the Democratic and Republican party respectively in the $i^{th}$ last presidential election. And $\beta_D$ is the overall fraction of the votes the Democrats received for the last two presidential elections across the entire country. To determine this value we sourced data from the United States Federal Election Commission (FEC) the official government body which was "created to promote confidence and participation in the democratic process" [37] .The exact number of votes received by each party in the 2012 election were:

- For the Barack Obama and Joe Biden of the Democratic party $65,915,795$ votes.
- For Mitt Romney and Paul Ryan of the Republican party $60,933,504$ votes.

For the 2016 election the exact results were:

- For the Hillary Clinton and Tim Kaine of the Democratic party $65,853,514$ votes.
- For Donald Trump and Mike Pence of the Republican party $62,984,828$ votes.

Using the above information we get the value of $\beta_D$ would be:

$$\frac{\frac{65{,}915{,}795}{65{,}915{,}795+60{,}933{,}504}+\frac{65{,}853{,}514}{65{,}853{,}514+62{,}984{,}828}}{2} \tag{7.2}$$

Which is exactly $51.53857559136132\%$. Note because the United States uses the electoral college system Donald Trump and Mike Pence won the 2016 election despite receiving fewer votes than Hillary Clinton and Tim Kaine. It is also important to note that while the United States is effectively a two party system there are other candidates who run for various offices including president. As an example Gary Johnson and Joe Weld of the Libertarian party received $4,489,341$ votes nationally in 2016 (over $3\%$ of the total vote). Since the Cook PVI is meant to be a direct comparison between the Democratic and Republican party it does not factor in third-party votes.

We should note that there is another interpretation of how one could calculate $\beta_D$. We also consider it reasonable to take the overall Democratic average of the total votes in the last two presidential elections. That is one could calculate it as:

$$\frac{65,915,795+65,853,514}{65,915,795+60,933,504+65,853,514+62,984,828} \tag{7.3}$$

If this were the case, then to get a "apples to apples" comparison then when calculating the PVI for a district we would use the following formula:

$$100\cdot\left(\frac{N_1^D+N_2^D}{N_1^D+N_1^R+N_2^D+N_2^R}-\beta_D\right) \tag{7.4}$$

instead of Equation 4.2. Unfortunately we were unable to find the exact details of how the PVI is calculated. Neither 538 or the Cook political report indicated the exact formula, instead the only information we could find was in this quote from 538:

"PVI measures how much more Democratic or Republican a district voted relative to the national result in an average of the last two presidential elections"

Luckily, in their data release 538 did calculate the PVI for single district states[1]. Given this and access to the vote totals in the 2012 and 2016 presidential elections in Vermont (a single district state) we were able to confirm Equation 7.2 and Equation 4.2 are the correct formulas for calculating the Cook PVI.

### The 538 sigmoid

The next step in the 538 model (which we examined in section 4.3.1) is going from the Cook PVI for a hypothetical district to how probable it is that district elects a Democrat. As explained in the subsection 4.3.1 we suspect 538 went with the sigmoid function (in that section we also explained the beneficial properties of the sigmoid). Recall, the sigmoid function takes the form:

$$\sigma(x) = \frac{1}{1 + e^{-w \cdot x}} \tag{7.5}$$

This function is fitted to data by adjusting the weight parameter $w$, if we look at a plot of the function (with $w = 1$) we see the characteristic $s$ shape (first subfigure of Figure 7.5). Luckily, 538 did publicly report the Cook PVI and their derived probability of a Democratic win for all of district plans in their catalogue for each state[2]. The Cook PVI, plotted against the probability of a Democratic win (second subfigure of Figure 7.5), clearly shows a sigmoid shape. From here we just need to derive what weight parameter $w$ they use is. To figure this out we first invert the sigmoid function using the log-odds (or logit) function:

$$logit(x) = \log_e \frac{1}{1 - x} \tag{7.6}$$

Inverting the sigmoid function with the logit function would produce a line given by $y = w \cdot x$, thus we simply need to invert any two data point in the second subfigure of Figure 7.5 and take the slope of the resulting line as $w$ (since this is a linear function of one variable any two distinct points are sufficient to determine it). Briefly, we mention two important points. Firstly, the sigmoid, and hence the line from the logit, may have a bias term associated with them. We found 538 did not include one since their their sigmoid passes through $(0, 50)$[3] and the resulting logit line passes through $(0, 0)$. Secondly, the points 538 published do not perfectly follow a sigmoid, instead there is a small amount of "jitter" on some of the points in the second subfigure of Figure 7.5. This deviation could simply be a rounding issue or minor transcription errors, in either case the points still very closely follow the sigmoid pattern. Because of the small amount of noise the resulting inverted plot found with the logit function will not be perfectly linear. Thus our choice of the two points for the inference of $w$ would (very slightly) influence the outcome. To mitigate this issue we take the ordinary least squares (OLS) regression line, also known as the line of best fit, for all of the points (after inverting them with the logit). We found the slope of the OLS line was $0.3047121945377743$ which we ended up using for the $w$ parameter in our sigmoid model. Our resulting model can be seen in the third subfigure of Figure 7.5. This is a near perfect fit for the 538 model since they form the line $y = x$ when plotted against each other (the final subfigure of Figure 7.5)

[1] We require a single district state since the 538 data release does not include the exact vote totals or precinct breakdown. So for states with multiple districts it is impossible to tell which PVI formula is correct.

[2] In total there are 3424 plans. These plans include all of their created plans such as the partisan plans, competitive plans and plans that emphasize compactness. They also include the current congressional plans.

[3] There is exactly one data point with a PVI of 0 and a Democratic probability of winning of $50\%$.

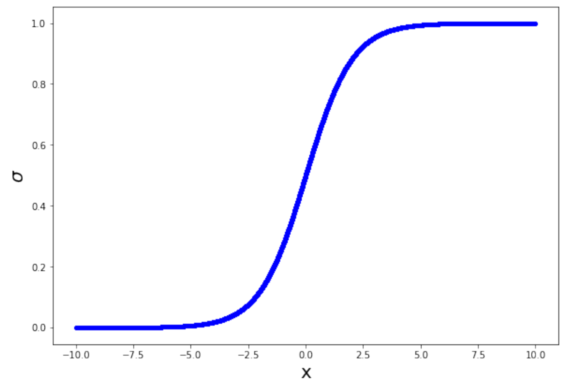


(a) Standard sigmoid curve

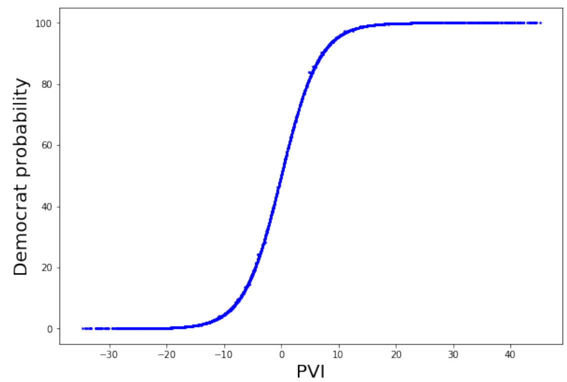


(b) 538 model

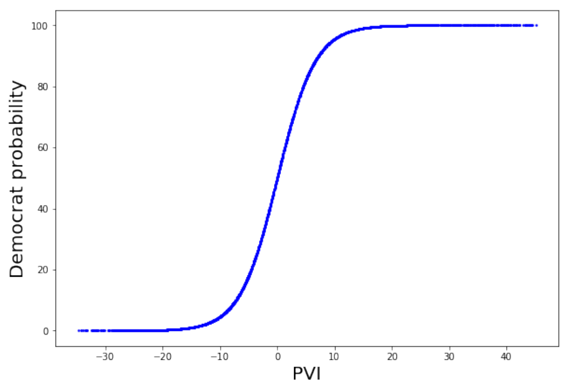


(c) Our model

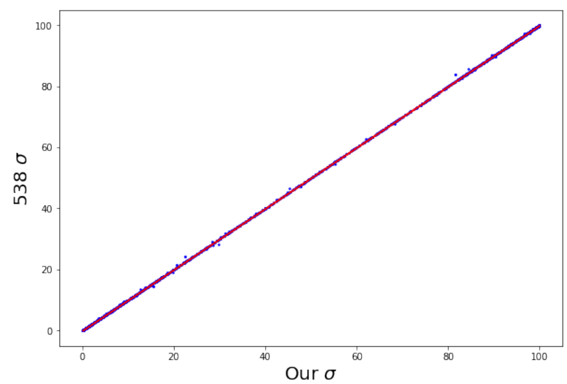


(d) 538 vs our model

Figure 7.5: Subfigure (a) is a standard sigmoid curve. Subfigure (b) is the reported 538 model (note 538 reported their probability of a Democrat winning on a 0 to 100 scale) for every districting created by 538. Subfigure (c) is our derived model (where the response is also plotted on a 0 to 100 scale) applied to every districting created by 538. Subfigure (d) is the output of our model and the 538 model plotted against each other for every districting created by 538 (in blue) the line $y = x$ is plotted in red. For subfigures (c) and (d) as inputs we used the PVI as provided by 538, this is because they do not provide a vote or precinct breakdown for their districts.

### 7.2.6 Gerrymandering for real data

In this section we will include all of the figures for the optimal gerrymandering in subsection 4.3.4. That is for all five states found in Table 4.3 we include four figures, the optimal outcomes 538 found when gerrymandering for the Republicans and Democrats and the outcomes our algorithm found when gerrymandering for the Republicans and Democrats.

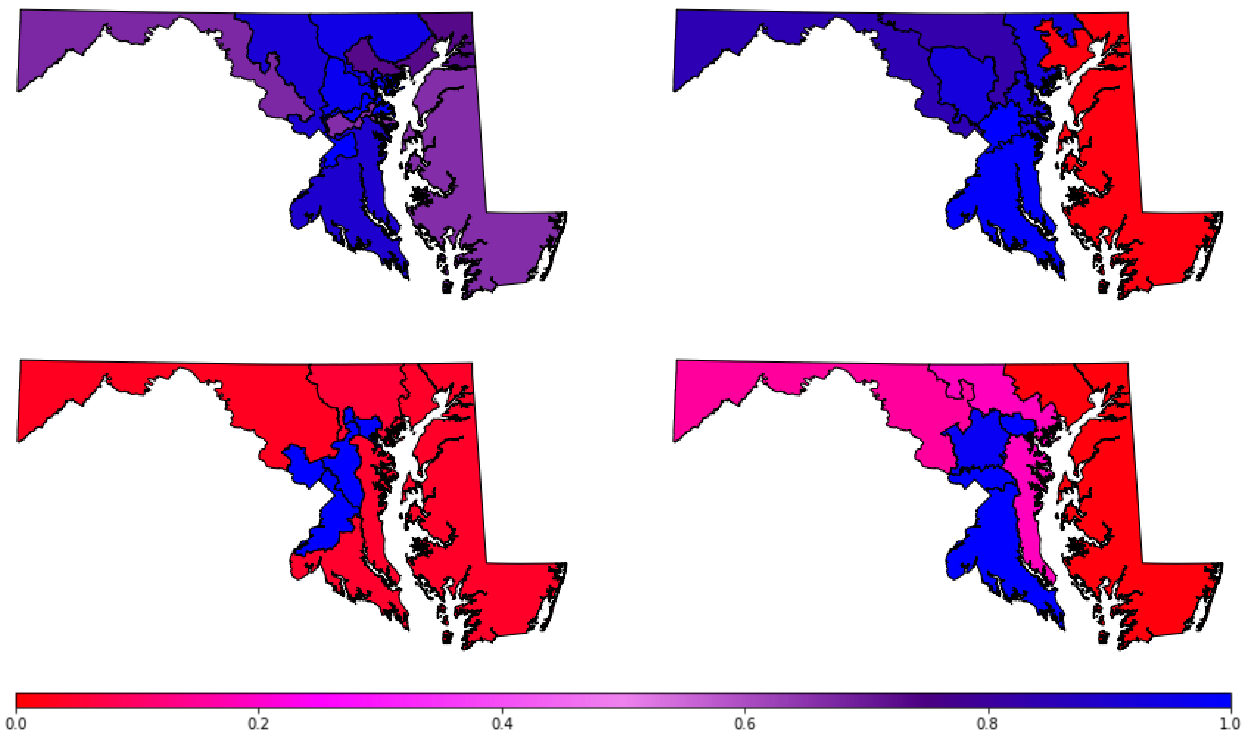


Figure 7.6: Gerrymandering outcomes for Maryland (MD). The top-left map is how 538 optimally gerrymandered for the Democrats, the top-right map figure is how our algorithm gerrymandered for the democrats. The bottom-left map is how 538 optimally gerrymandered for the Republicans, the bottom-right map is how our algorithm gerrymandered for the Republicans. The bottom bar shows the scale for the probability the Democrats win the district.

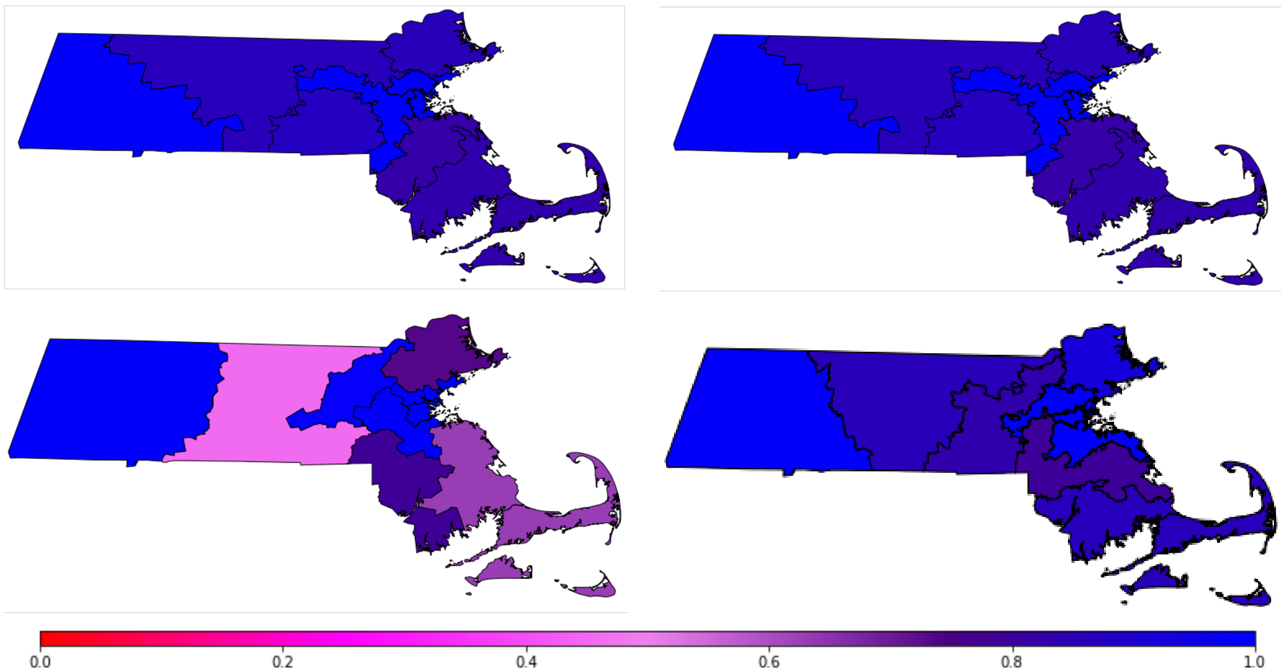


Figure 7.7: Gerrymandering outcomes for Massachusetts (MA). The top-left map is how 538 optimally gerrymandered for the Democrats, the top-right map figure is how our algorithm gerrymandered for the democrats. The bottom-left map is how 538 optimally gerrymandered for the Republicans, the bottom-right map is how our algorithm gerrymandered for the Republicans. The bottom bar shows the scale for the probability the Democrats win the district.

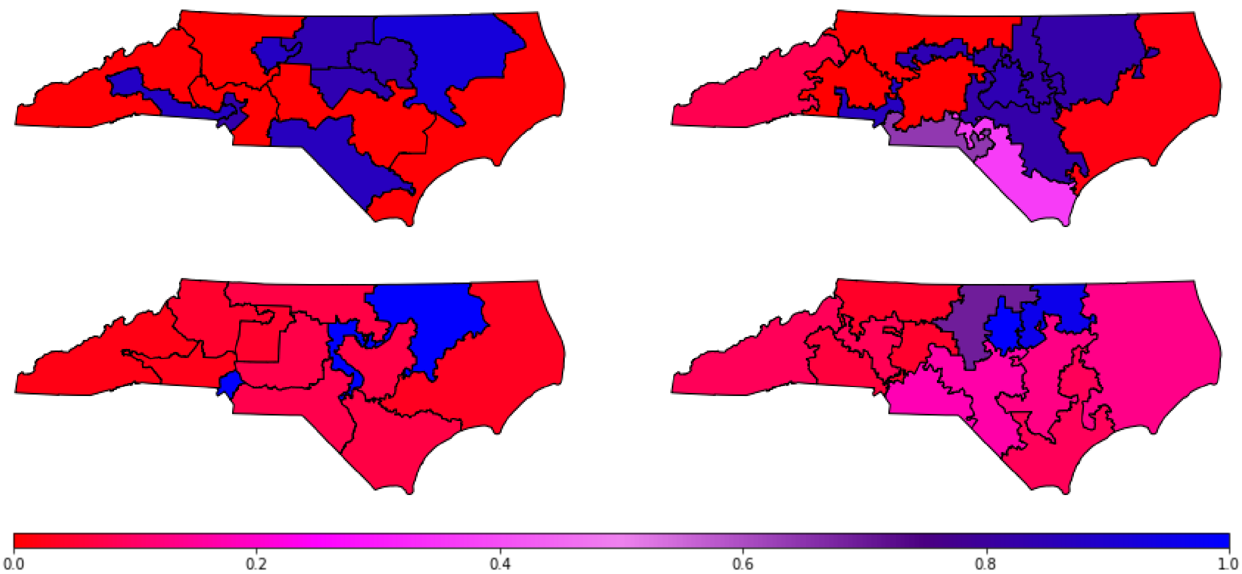


Figure 7.8: Gerrymandering outcomes for North Carolina (NC). The top-left map is how 538 optimally gerrymandered for the Democrats, the top-right map figure is how our algorithm gerrymandered for the democrats. The bottom-left map is how 538 optimally gerrymandered for the Republicans, the bottom-right map is how our algorithm gerrymandered for the Republicans. The bottom bar shows the scale for the probability the Democrats win the district.

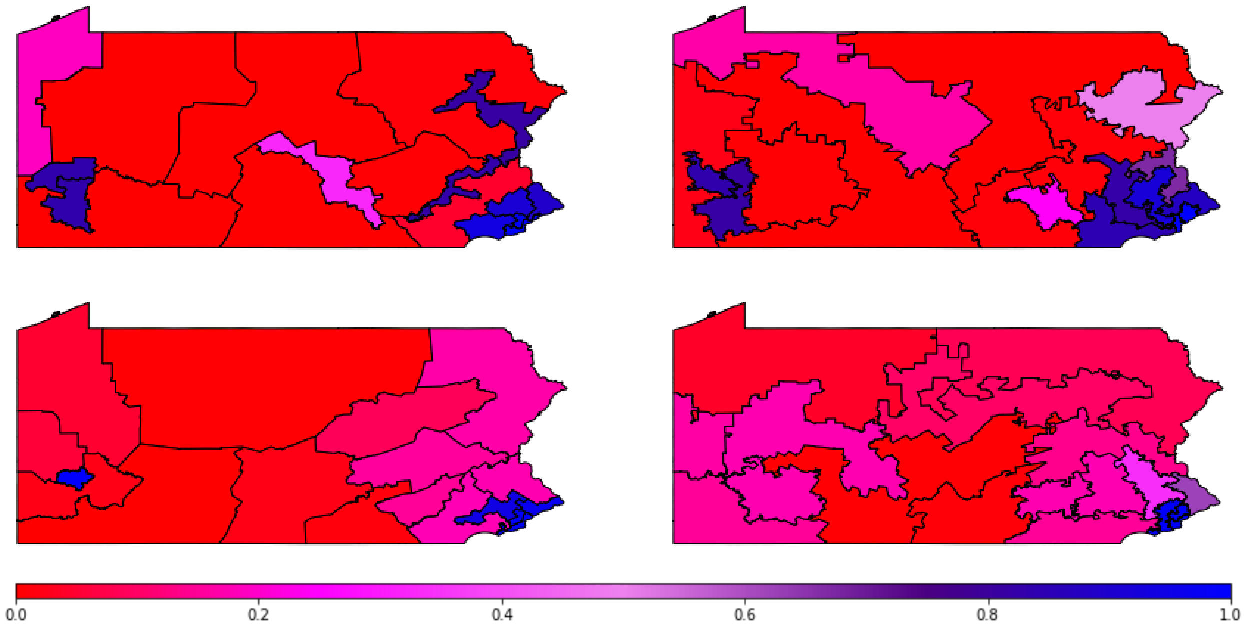


Figure 7.9: Gerrymandering outcomes for Pennsylvania (PA). The top-left map is how 538 optimally gerrymandered for the Democrats, the top-right map figure is how our algorithm gerrymandered for the democrats. The bottom-left map is how 538 optimally gerrymandered for the Republicans, the bottom-right map is how our algorithm gerrymandered for the Republicans. The bottom bar shows the scale for the probability the Democrats win the district.

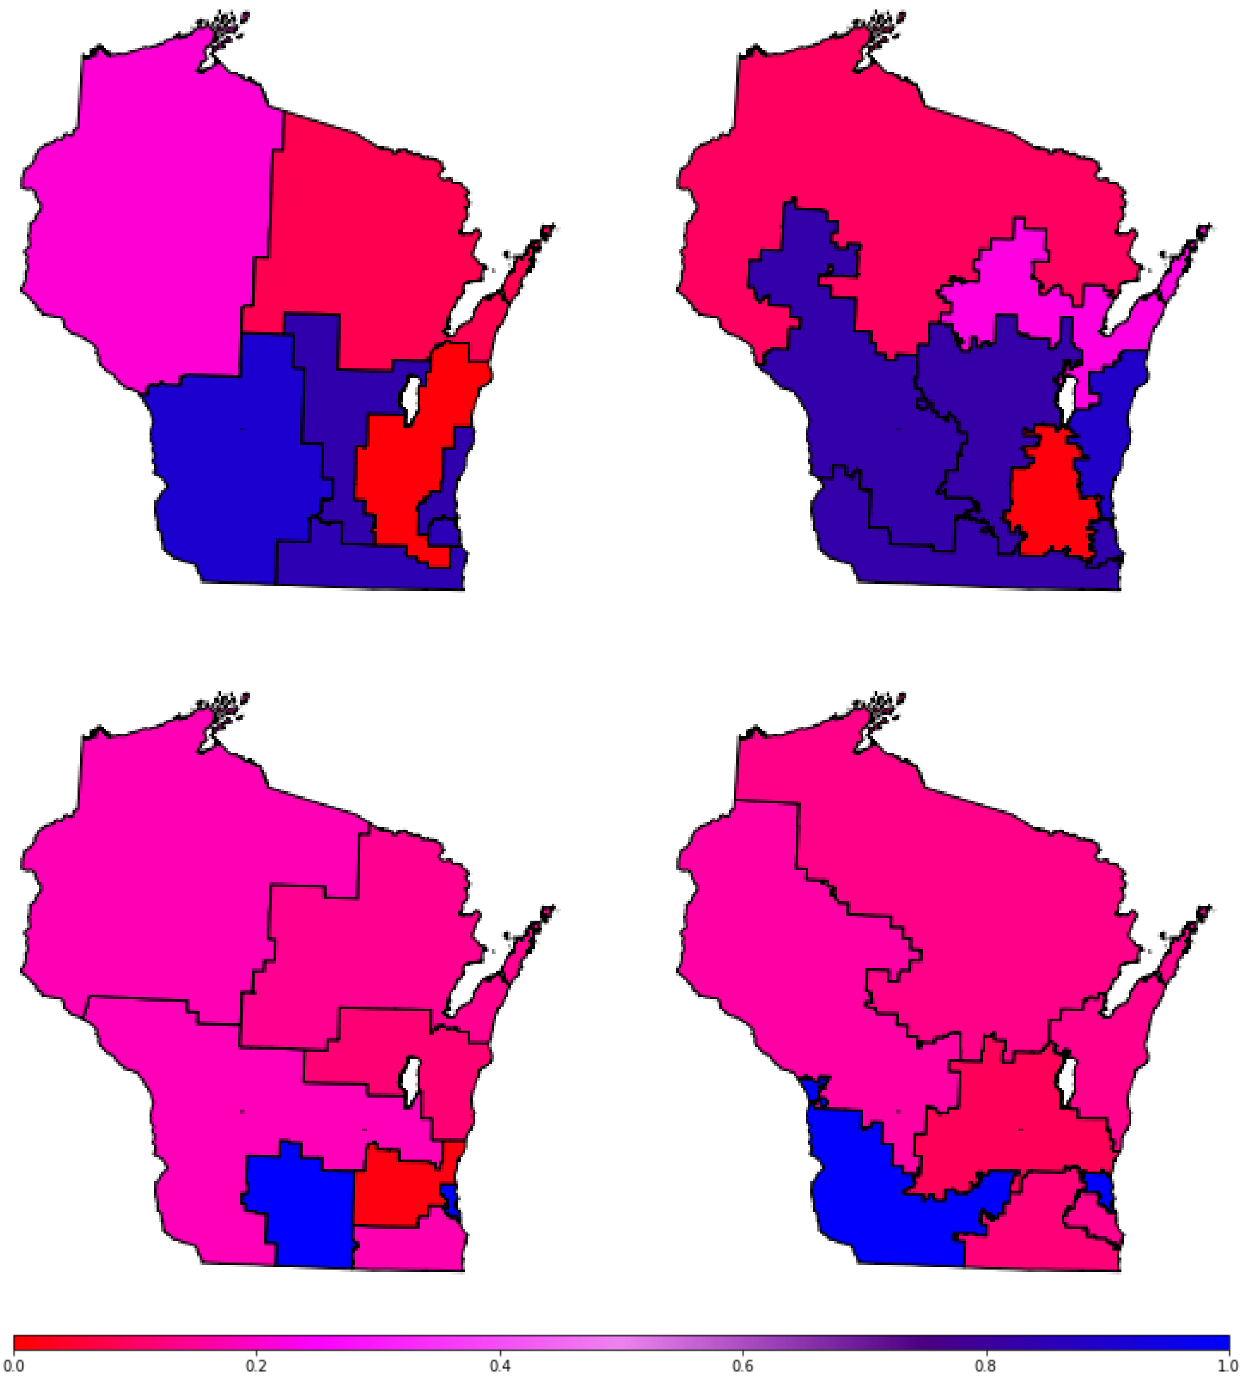


Figure 7.10: Gerrymandering outcomes for Wisconsion (WI). The top-left map is how 538 optimally gerrymandered for the Democrats, the top-right map figure is how our algorithm gerrymandered for the democrats. The bottom-left map is how 538 optimally gerrymandered for the Republicans, the bottom-right map is how our algorithm gerrymandered for the Republicans. The bottom bar shows the scale for the probability the Democrats win the district.