Primarily About Primaries

Allan Borodin[a], Omer Lev[b], Nisarg Shah[a], Tyrone Strangway[b]

[a] *University of Toronto, Ontario, Canada*
[b] *Ben-Gurion University of the Negev, Beersheba, Israel*

Abstract

Much of the social choice literature examines *direct* voting systems, in which voters submit their ranked preferences over candidates and a voting rule picks a winner. Real-world elections and decision-making processes are often more complex and involve multiple stages. For instance, one popular voting system filters candidates through *primaries*: first, voters affiliated with each political party vote over candidates of their own party and the voting rule picks a set of candidates, one from each party, who then compete in a general election.

We present a model to analyze such multi-stage elections, and conduct what is, to the best of our knowledge, the first quantitative comparison of the direct and primary voting systems in terms of the quality of the elected candidate, using the metric of *distortion*, which attempts to quantify how far from the optimal winner is the actual winner of an election. Our main theoretical result is that voting rules (which are independent of party affiliations, of course) are guaranteed to perform in the primary system within a constant factor of the direct, single stage setting. Surprisingly, the converse does not hold: we show settings in which there exist voting rules that perform significantly better under the primary system than under the direct system. Using simulations, we see that plurality benefits significantly from using a primary system over a direct one, while Condorcet-consistent rules do not.

1. Introduction

> *If I could not go to heaven but with a party, I would not go there at all.*
>
> – Thomas Jefferson, 1789

Thomas Jefferson, like many of the US constitution's authors, believed that political parties and factions are a bad thing (see also Hamilton et al. [1]). This view stemmed from a long history of British and English political history, in which prison sentences and executions were possible outcomes in the battle between factions for supremacy at the Royal court [2]. However, both in Britain and in the Unites States, once their respective legislative assemblies gained political force, parties turned out to be quite unavoidable. Even Jefferson eventually had to start his own party (the *Republican* party[1]), which ended up quite successful, and was able to vanquish the opposing party (the *Federalist* party) from political existence [3].

Moving on to current times, political parties have become the bedrock of parliamentary politics throughout the world. In particular, one of political parties' main roles – if not the most important (especially in presidential systems) – is to select the candidates which are voted on by the general public. The mechanisms by which parties make this selection are varied, and they have evolved significantly throughout the

Email addresses: `bor@cs.toronto.edu` (Allan Borodin), `omerlev@bgu.ac.il` (Omer Lev), `nisarg@cs.toronto.edu` (Nisarg Shah), `strangwa@post.bgu.ac.il` (Tyrone Strangway)

[1] Commonly referred in research as the *Democratic-Republican* party, as it is not a precursor of the modern Republican party.

past 150 years. But in the past few decades there has been a marked shift by parties throughout the world towards increasing the ability of individual party members to influence the outcome, and in some cases, to be the only element to determine party candidates [4]. In particular, US parties have changed their election methods since the 1970s to focus the selection of presidential, congressional and state-wide candidates on popular support by party members via primaries [5].

Despite this long and established role of parties in whittling down the candidate field in elections, the treatment of a party's role in elections within the multiagent systems community has been quite limited. While various candidate manipulation attacks have been investigated (e.g., Sybil attacks [6]), and there is recent research into parties as a collection of similar minded candidates (e.g., in gerrymandering, across different districts), the role of parties in removing candidates has not been analyzed.

The focus of this paper is the *primary* voting system, in which each party's electorate selects a winner from among the party's candidates, and among these primary winners, an ultimate election winner is selected by the general public. We compare this system to the *direct* voting system, in which all voters directly vote over all candidates. Moreover, we expand this voting mechanism to encompass various multi-stage processes, by which parties may have multiple stages before the general election (for example, in Britain, before the party membership votes, each party's parliamentary members can decide which candidates will be voted on).

A simple example of such a mechanism is illustrated in Figure 1, showing how a rather extreme candidate (candidate d) can get elected in a regular direct election. However, a primary lets one party select a more widely-accepted candidate, and the removal of various candidates (thanks to the primary process) lets this candidate become the ultimate winner. However, Figure 2 shows how primaries can also cause extreme candidates to win. Details about these two possible effects of a primary system are provided in Example 1.

Though ostensibly about parties, voters and elections, the multi-stage model considered here applies beyond the realm of political affairs, to a variety of decision-making processes used by agents. This can be used to manage the number of candidates – when the potential number of options is huge, it is common to use subdivisions to cull the options and present only a few of them for discussion and vote. For example, in many universities academic departments bring forward only a subset of the candidates vying for a position to the faculty forum which decides on hiring. Or a city may ask its regional subdivisions to assess which roads require urgent fixing, and then the city council decides from these options where to invest its efforts. The multi-stage model can also be used for *gatekeeping* purposes, i.e., to ensure only "good-enough" options are presented (for example, by having a forum of experts vet options before being brought forward for general consideration).

1.1. Our Results

Our contribution is twofold. First, we formulate a model which allows a quantitative comparison of direct elections vs. a primary system. Our model is a spatial model of voting in which voters and candidates are located in an underlying metric space, and voter preferences are single-peaked (i.e., voters prefer candidates which are closer to them according to the metric of the space). This allows us to compare each candidate's social utility in terms of its total distance to the voters. We formalize the evaluation objective by using the notion of *distortion* as advocated in a recent line of research [7, 8].

Our results begin by analyzing 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 show that while arbitrarily sized parties can increase an election's distortion to the maximal value possible, when each political party has some constant proportion of the electorate, the primary system will increase distortion by at most a constant factor. Moreover, it can also improve distortion in some cases – when parties are separable (i.e., the ideological metric space can be divided between the parties), distortion can *decrease* by a factor linear in the number of voters for some voting rules. We then extend this setting to include additional voters – independent voters (i.e., not party affiliated) and voters who do not participate in the general election – as well as multiple parties and multiple decision stages, showing that the potential increase of distortion by primaries remains at multiplicative constant factor (the constant factor depending, as before, on the share of population voting in each primary, and the share of independent voters).

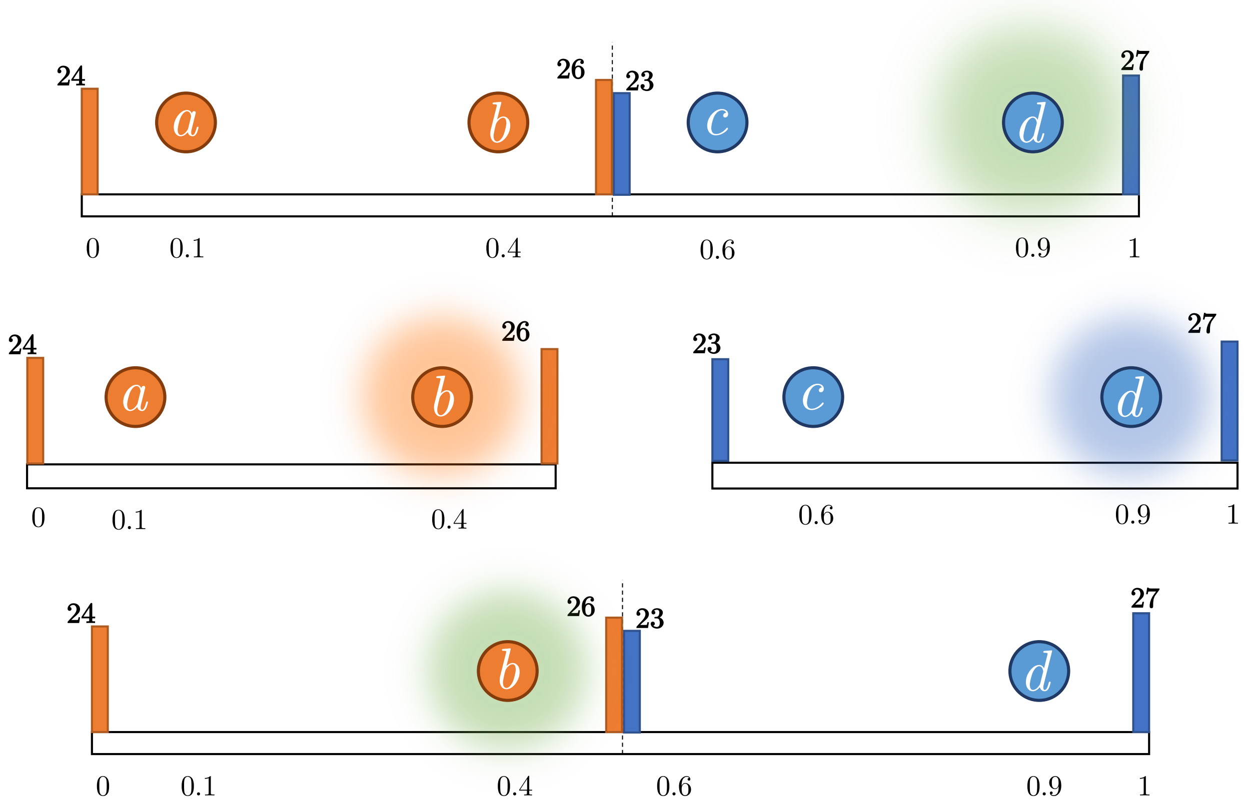

Figure 1: A direct plurality vs. a primary system (as described in Example 1). Party -1 is in orange, while party 1 is in blue. The top line is the direct election (the winner highlighted in green). The second and third lines show the primary election – the second line shows the party winners, and the bottom line depicts the general election in which only primary winners participate.

Second, we use this model to present a comparison of the direct and primary voting systems not in worse-case bounds but in simulations. We show how different voting systems behave quite differently under primary and direct elections. In particular, we show 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.

This paper is an extension of a AAAI 2019 paper [9]. Sections 8 and 9 are new additions in this version. Several proofs missing in the proceedings version of the AAAI 2019 paper have been added. Also, many proofs have been rewritten, simplified, or reorganized for clarity.

2. Related Work

The analysis of regular, direct elections is long and varied, both in the social sciences and in AI [10]. In our particular setting, the voters are located in a metric space, with their preferences related to their distance from candidates. Such settings have been widely investigated in the social science literature since the work of Downs [11], recently summarized by Schofield [12]. In particular, we focus on the concept of distortion, introduced by Procaccia and Rosenschein [7]. While they analyze it in the context of a direct election in which voters have arbitrary utilities for candidates, we consider distortion in a primary setting in which the costs (negative utilities) of voters for candidates are determined by their locations in an underlying metric space, which represents their opinions. Voter distortion in metric spaces in regular elections (which we term *direct*) was investigated in a series of papers [13, 14, 15, 16, 17, 18] for most common voting rules. Feldman et al. [19] explored such a setting for strategyproof mechanisms. See also the survey [20].

Discussing changes to the set of candidates has mainly focused on two paths of research. Strategic candidates, investigation of which began with the work of Dutta et al. [21] – followed by Dutta et al. [22] discussing strategic candidacy in tournaments, and more recently further explored by Brill and Conitzer [23], Polukarov et al. [24] and others – deals mainly with finding equilibria. The other is the addition and removal of candidates, as a form of control manipulation, which was studied by Bartholdi III et al. [25]; see the summaries by Brandt et al. [10] and Rothe [26].

Investigating parties' selection methods and their effect on the election has mostly been done in the social sciences. Kenig [27] details the range of selection methods parties use, and there has been significant focus on more democratic methods for leader selection [28], which seems to be a general trend in many Western countries [4]. There is also significant literature on particular party elections in various countries, such as Britain [29], Belgium [30], Israel [31], and many others. The most widely examined country is the US, in which political parties have been a fixture of political life since its early days [3]. The most recent extensive summary of research on it is Cohen et al. [5], who try to explain how party power-brokers influence the party membership vote. Norpoth [32] uses primary data to predict election results, and notably Sides et al. [33] show that primary voters are very similar to "regular" voters. In computational fields, recent interest in proxy voting [34, 35], in which voters give other agents the ability to vote for them, may be related to how modern parties are viewed and analyzed. Concretely, party analysis has mostly been about having the same set of candidates in different voting domains, as in district elections [36], and their related manipulation problem, gerrymandering [37, 38, 39, 40, 41, 42]. Note that we do not deal with how parties, as strategic agents, may try to divide the ideological space between them, as done (implicitly) in Voronoi games [43, 44, 45, 46, 47], though a recent paper [48, 49] began addressing the question of candidates manipulation of this sort in Hotelling-Downs settings.

3. 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*, which are the alternatives voters wish to select from. We assume that voters and candidates lie in an underlying metric space $M = (S, d)$, where S is a set of points and d is a distance function satisfying the triangle inequality, symmetry, and the identity of indiscernibles. 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 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.

Less formally, this metric space serves as an ideological space, so proximity between points means similar views. The number of dimensions reflects the number of policy areas, or, more generally, whatever salient different topics there are that are the constituent parts of reaching a decision ranking the different candidates.

We shall assume in our model that voters and candidates also 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 8 (and Appendix A) 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 $n_p, m_p \geqslant 1$ for each $p \in \{-1, 1\}$.

Overall, an *instance* is a tuple $I = (V, A, M, \rho, \pi)$. Given I, we can use its information to find a winning candidate $a \in A$. 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 $C^I_p(a) = \sum_{i \in V_p} d(i, a)$. Hence $C^I(a) = C^I_{-1}(a) + C^I_1(a)$. For $X \subseteq V$, we also use $C^I_X(a) = \sum_{i \in X} d(i, a)$. 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, and drop the subscript when referring to all existing voters (V).

However, our assumptions (as in most voting literature, see Brandt et al. [10]), are such that they do not allow us to observe the full instance and know which candidate is a_{OPT}. Specifically, we do not know the underlying metric M or the embedding function ρ. 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 $a, b \in A$, $a \succ_i b \Rightarrow d(i,a) \leqslant d(i,b)$. 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 for common metric spaces – in a Euclidean space, for example, any preference order, for any number of voters and candidates can be expressed in a Euclidean metric space with a large enough dimension. While we do not know the utility functions of all the voters (the assumption made by Procaccia and Rosenschein [7]), we shall use the metric distance as a proxy to that utility function (an approach used in most distortion papers of the last decade [13, 14, 15, 16]).

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. We shall use the notation $\mathcal{M}$ to denote the space we operate in. Usually it will be the metric space, but sometimes it will also specify particular properties of the ρ and π functions. For some $\mathcal{M}$, let $\mathcal{I}^{\alpha}_{m,\mathcal{M}}$ be the family of instances satisfying the following conditions:

- Each party has at least an α fraction of the voters affiliated with it, i.e., $n_p \geqslant \alpha \cdot n$ 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 α in $\mathcal{I}^{\alpha}_{m,\mathcal{M}}$.
- The number of candidates is 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 ρ must be such that $\rho(V_{-1} \cup A_{-1})$ and $\rho(V_1 \cup A_1)$ are linearly separable.[2] In this case we also take the metric to be $M = (\mathbb{R}^k, d)$ with d as the standard Euclidean distance.

The *separable* space $\mathcal{M} = \text{sep-}\mathbb{R}^k$ captures the case of a highly polarized society. The reason to examine separable spaces is that, as mentioned above, each dimension of the metric space $\mathbb{R}^k$ may represent the range of opinions on an issue, and voters and candidates may be embedded in this space based on their opinions on the different issues. In general, we would expect voters and candidates affiliated with any party to be mapped to nearby points since they may share somewhat similar opinions on major issues. Thus, a separable space means that each party has a particular ideology (i.e., occupies a particular area of the ideological space) that does not overlap with any other party. In the single dimension, this means there would exist 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. In reality, this alignment may not be perfect, but when there is significant polarization, voters and candidates affiliated with one party may still be sufficiently far from those affiliated with the other party, giving rise to a hyperplane separating them. Note that this choice of $\mathcal{M}$ restricts the embedding ρ based on the party affiliation π.

These families of instances are related by the following relation. For all k, the relations between the families can be seen here:

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

[2] Two sets of points are linearly separable if the *interiors* of their convex hulls are disjoint, i.e., if there exists a hyperplane that contains each set in a different closed halfspace.

3.1. Voting Rules and Distortion

A *voting rule* f takes an election as input, and returns a winning candidate from A. This means the voting rule does not have access to the information contained in the ideological metric space, which means it will not necessarily choose the candidate whose distance from all voters in the metric space is minimal. *Distortion* is a metric that aims to quantify, for each voting rule, how far can it stray from the optimal candidate, and by how much it will increase the social cost of the voters vs. an oracle who has access to the ideological metric space. This is conceptually similar to metrics such as price of anarchy [50, 51], which tell us how far from the optimum will the cost be of limiting ourselves to Nash equilibria only. We say that the cost-approximation of f on instance I is

$$\phi(f, I) = \frac{C^I(f(E^I))}{\min_{a \in A} C^I(a)},$$

that is, the ratio of the cost of the elected and the cost of the optimal candidate. Given a family of instances $\mathcal{I}$, the distortion of f with respect to $\mathcal{I}$ is

$$\phi_{\mathcal{I}}(f) = \sup_{I \in \mathcal{I}} \phi(f, I).$$

Just like price of anarchy, distortion is a worst-case notion, so if $\mathcal{I} \subseteq \mathcal{I}'$, $\phi_{\mathcal{I}}(f) \leqslant \phi_{\mathcal{I}'}(f)$ for every voting rule f. This means, as we will see below on plurality, there are cases where the worst-case distortion is quite different than what we see in common settings. Moreover, as a metric, distortion, at least in some metric spaces, has a downside of lack of normalization. If all voters are condensed in a small area of the space (e.g., in the interval $[0, 1]$), a single voter who can be located far, far away, can significantly influence the distortion value. However, in this paper – as in much of the literature on distortion [16] and in the game theoretical fields [50] – we shall assume that a smaller distortion, denoting higher sum of voter utilities, is a desirable goal.

Standard voting rules choose the winning candidate independently of party affiliations. These include rules such as plurality, Borda, and STV. We refer readers to Brandt et al. [10] for a detailed list of various voting rules, but we shall just define several which we shall mention along the way:

Plurality Each voter gives a single point to its top-ranked candidate. The candidates with the highest scores are the winners (and a tie-breaking rule is used to determine a single one).

Borda Each voter gives $m-1$ points to their top-ranked candidate, $m-2$ to their second ranked candidate, and so on: $m-i$ points are given to the candidate ranked i. The candidates with the highest scores are the winners (and a tie-breaking rule is used to determine a single one).

STV (Single Transferable Vote) Each voter gives a single point to its top-ranked candidate. If no candidate has more than $\frac{n}{2}$ points, the candidate with the lowest score (if there are several – a tie-breaking rule is used) is eliminated, and all voters who gave their point to that candidate give their point to the highest-ranked candidate according to their preference which has not been eliminated. The process continues until there is a candidate with more than $\frac{n}{2}$ points.

Maximin Each candidate $a \in A$ is assigned a score – $\min_{b \in A} |\{i \in V : a \succ_i b\}|$ – which represents its core support in any pairwise election. The winners are the candidates with the highest such score (and a tie-breaking rule is used to determine a single one).

Copeland For each candidate $a \in A$ we define a set $A_+^a = \{b \in A : \frac{n}{2} > |\{i \in V : b \succ_i a\}|\}$, the candidates it strictly beats in a pairwise election, and $A_-^a = \{b \in A : \frac{n}{2} > |\{i \in V : a \succ_i b\}|\}$, the candidates it is strictly beaten by in a pairwise election. Each candidate's score is $|A_+^a| - |A_-^a|$.[3] The winners are the candidates with the highest score (and a tie-breaking rule is used to determine a single one).

[3] This is sometimes known as the *symmetric* Copeland score; see [10, Section 2.2]. As Brandt et al. [10] point out (Footnote 12 in Chapter 2), there are other definitions of the Copeland score used in the literature that are equivalent, i.e., which lead to the same set of winners.

The final two voting rules are *Condorcet consistent*: when there is a candidate that can beat any other one in a pairwise election (i.e., a candidate $a \in A$ such that $|\{i \in V : a \succ_i b\}| > \frac{n}{2}$ for any $b \in A$), that candidate (the *Condorcet winner*) will be the election winner. We note that such a candidate is not guaranteed to exist.

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.

3.2. 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 the obvious choice (see May's Theorem [52]), and what most voting rules become when dealing with only two candidates.

This setting resembles systems employed by the main parties in US, Canada and other countries, in which a party's members vote on their party's candidates to select a winner of their primary.[4] In other systems, the selection by a party could be a multi-stage process. While we assume for now that every voter votes in the general election, in Section 8 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 paper 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).

Example 1. Suppose $M = ([0, 1], d)$ (d being the Euclidean metric), with 100 voters, of which 50 are of party -1, and 50 of party 1, with the point 0.5 being the dividing line between the party (i.e., this is a separable space). As Figure 1 shows, 24 voters are at point 0, 26 at point $0.5 - \varepsilon$ (for some $\varepsilon < 0.1$), 23 voters at $0.5 + \varepsilon$, and 27 voters at point 1.

There are 4 candidates: candidate a at 0.1, with a social cost of $46.3 - 3\varepsilon$; candidate b at 0.4 with social cost $30.7 - 3\varepsilon$; candidate c at 0.6 with social cost $30.1 + 3\varepsilon$; and candidate d at 0.9 with social cost $43.9 + 3\varepsilon$. Thus, the optimal candidate is c.

A regular plurality election will result in d winning with 27 votes (vs. 24 for a, 26 for b, and 23 for c), resulting in a distortion of $\frac{43.9+3\varepsilon}{30.1+3\varepsilon} \sim 1.46$. However, a primary system will result in party -1 electing candidate b (26 votes vs. 24 for a) and party 1 electing candidate d (27 votes vs. 23 for c). Thus, the general election is run between candidates b and d, resulting in b winning with 73 votes (vs. 27 for candidate d), with a distortion of $\frac{30.7-3\varepsilon}{30.1+3\varepsilon} \sim 1.02$. Note that this improves the winner from one that is 47% worse than the optimal to one that is only 2% worse than the optimal.

[4] Sometimes each party uses a different voting rule, but in this paper, we are mainly interested in comparing the direct vs. primary cases. However, as will be seen below, the distortion bounds we prove still hold when parties' voting rules are different.

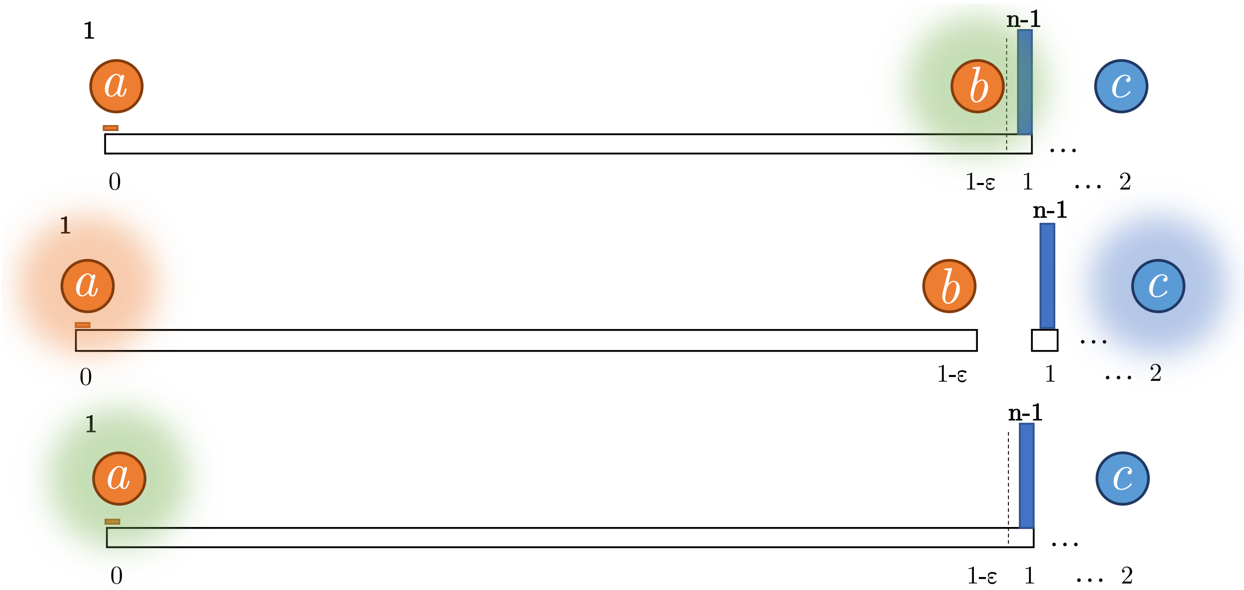

Figure 2: The construction from the proof of Theorem 1 for the case of $m = 3$ candidates, where the primary system results in an unbounded distortion. Party -1 is in orange, while party 1 is in blue. The top figure shows the result of a direct plurality election, in which the winner is also the optimal candidate (highlighted in green). The middle and bottom figures show the result under a primary system based on the plurality rule – the middle figure highlights the primary winners and the bottom figure shows the general election in which only the primary winners participate.

4. Very Small Primaries are Terrible

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

We first show that when there are no restrictions on α, and in particular, 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 1. *For $m \geqslant 3$, $\phi_{\mathcal{I}_{m,\text{sep-}\mathbb{R}}}(\widehat{f}) = \infty$ for every affiliation-independent unanimous voting rule f.*

Proof. Consider an instance $I \in \mathcal{I}_{m,\text{sep-}\mathbb{R}}$ in which one voter is located at 0 and affiliated with party -1, while the remaining $n-1$ voters are located at 1 and affiliated with party 1. The candidates are located as follows: one is at 0 (affiliated with -1), $m-2$ are at $1-\epsilon$ for some $\frac{1}{n^2} > \epsilon > 0$ (also affiliated with party -1), and one at 2 (affiliated with party 1).

The candidate at 2 becomes the primary winner of party 1, while the candidate a^* at 0 becomes the primary winner of party -1 as well as the general election winner. Its social cost is $C(a^*) = n-1$. In contrast, an optimal candidate a_{OPT} at $1-\epsilon$ has social cost $C(a_{\text{OPT}}) = (1-\epsilon)+\epsilon(n-1) = 1+\epsilon(n-2) \leqslant 2$. Hence, $\frac{C(a^*)}{C(a_{\text{OPT}})} \geqslant \frac{n-1}{2}$. Since the number of voters n is unbounded, we have $\phi_{\mathcal{I}_{m,\text{sep-}\mathbb{R}}}(\widehat{f}) = \infty$. □

Theorem 1 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 (or even all but one) of the candidates located at $1-\epsilon$ to 2 and assign them to party 1, and the proof would still hold (see Figure 2, which shows a 3 candidate variant of this description). Moreover, the results will still hold even if $\alpha = o(1)$, since the cost of the candidate at 0 will still be larger than $0.9n$, while the cost of the candidate at $1-\epsilon$ will be bounded by $n\epsilon + o(n)$, and the ratio between those is still unbounded.

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. Large Primaries are Never Much Worse than Direct Elections

In this section, we show that if at least a constant fraction of all voters are affiliated with each party, then for every affiliation-independent voting rule f, we can bound the distortion of $\widehat{f}$ in terms of the distortion of f and share of population that votes in each party *for every instance*. Note that this is stronger than comparing the worst-case distortions of f and $\widehat{f}$ over a family of instances (as we will later see in Theorem 7).

The key challenge in bounding the distortion of $\widehat{f}$ is that in each primary, only the voters of the party are discarding candidates, thus the distortion of f, $\phi(f)$, only bounds the damage to the social cost of these voters. But we want to understand the damage to the social cost of *all* voters, not only those of one particular party. The following lemma shows that bounding all voters can still be achieved if the primary is large (i.e., at least a constant fraction of the voters are affiliated with each party). More generally, the lemma bounds the distortion of an election in terms of the distortion of the voting rule and the fraction that voted.

Lemma 2. *Let $I = (V, A, M, \rho, \pi)$ be an instance. Let $I' = (V', A', M, \rho, \pi)$ be a sub-instance in which a subset of voters $V' \subseteq V$ vote over a subset of candidates $A' \subseteq A$. Let $|V'| \geqslant \alpha \cdot |V|$ for $\alpha > 0$. Let a^* be the winner in I' under voting rule f. If C is the social cost function under instance I, then*

$$C(a^*) \leqslant \frac{1 - \alpha + \phi(f, I')}{\alpha} \cdot \min_{a \in A'} C(a).$$

Proof. Denote $\theta = \phi(f, I')$ (hence $\theta \geqslant 1$). Recall that C_X denotes the social cost with respect to a set of voters X (thus, $C = C_V$). Let $n = |V|$ and $n' = |V'|$.

Fix $a \in A'$. Due to the definition of distortion, we have $C_{V'}(a^*) \leqslant \theta \cdot C_{V'}(a)$. We want to bound $C(a^*)$ in terms of $C(a)$. Note that

$$\begin{aligned} C(a^*) &= C_{V'}(a^*) + C_{V \setminus V'}(a^*) \\ &\leqslant \theta \cdot C_{V'}(a) + C_{V \setminus V'}(a) + (n - n') \cdot d(a, a^*) \\ &\leqslant \theta \cdot C(a) + (n - n') \cdot d(a^*, a), \end{aligned} \tag{1}$$

where the first inequality holds due to the triangle inequality and because we have derived $C_{V'}(a^*) \leqslant \theta \cdot C_{V'}(a)$, and the final transition holds because $\theta \geqslant 1$ and d is symmetric. We also have $d(a^*, a) \leqslant d(a^*, i) + d(i, a)$ for any $i \in V'$. Summing over all $i \in V'$ and averaging, we get

$$d(a^*, a) \leqslant \frac{C_{V'}(a^*) + C_{V'}(a)}{n'} \leqslant \frac{1 + \theta}{n'} \cdot C_{V'}(a) \leqslant \frac{1 + \theta}{n'} \cdot C(a). \tag{2}$$

Substituting Equation (2) into Equation (1), and using $\frac{n-n'}{n'} \leqslant \frac{n(1-\alpha)}{\alpha n} = \frac{1-\alpha}{\alpha}$, we get

$$\begin{aligned} C(a^*) &\leqslant \theta \cdot C(a) + (n - n') \cdot \frac{1 + \theta}{n'} \cdot C(a) \leqslant \left(\theta + \frac{1 - \alpha}{\alpha}(1 + \theta)\right) \cdot C(a) \\ &= \frac{1 - \alpha + \theta}{\alpha} \cdot C(a), \end{aligned}$$

as needed. ◻

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 ρ_p and π_p are restrictions of ρ and π to $V_p \cup A_p$. The primary election E_p of party p is precisely the election corresponding to instance I_p. Let a_p^* be the primary winner of party p under voting rule f, and $n_p = |V_p| \geqslant \alpha n$. Recall that C denotes the overall social cost under instance I. We can immediately bound $C(a_p^*)$ in terms of $\min_{a \in A_p} C(a)$ by applying Lemma 2 using I_p as the sub-instance.

Corollary 3. *Let a_p^* denote the primary winner of party p. Then*

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

It now remains to compare the social cost of the general election winner a^* to the social costs of the primary winners. Note that the voting rule used in the general election is the majority rule, which can be thought of as plurality for two candidates. Anshelevich et al. [13] show that the distortion of plurality for m candidates is $2m-1$. Substituting $m=2$, we get that the distortion of the majority rule is 3, which yields the desired comparison.

Corollary 4. *Let a^*_{-1} and a^*_1 be the two primary winners and $a^* \in \left\{a^*_{-1}, a^*_1\right\}$ be the winner of the general election. Then,*

$$C(a^*) \leqslant 3 \cdot \min\left\{C(a^*_{-1}), C(a^*_1)\right\}.$$

By combining Corollaries 3 and 4, we immediately get the following result.

Theorem 5. *Let $I = (V, A, M, \rho, \pi)$ be an instance. For $p \in \{-1, 1\}$ and $\alpha > 0$, let I_p be the primary instance of party p and $n_p = |V_p| \geqslant \alpha n$. Let $a_{\mathrm{OPT}} \in \arg\min_{a \in A} C(a)$ be a socially optimal candidate. Then,*

$$\phi(\widehat{f}, I) \leqslant 3 \cdot \frac{1-\alpha+\phi(f, I_{\pi(a_{\mathrm{OPT}})})}{\alpha} \leqslant 3 \cdot \frac{1-\alpha+\max_{p\in\{-1,1\}}\phi(f, I_p)}{\alpha}.$$

Proof. Let $p = \pi(a_{\mathrm{OPT}})$ denote the party with which the socially optimal candidate a_{OPT} is affiliated. Then, Corollary 4 implies that $C(a^*) \leqslant 3 \cdot C(a^*_p)$, and Corollary 3 implies that $C(a^*_p) \leqslant \frac{1-\alpha+\phi(f,I_p)}{\alpha} \cdot C(a_{\mathrm{OPT}})$. Combining the two inequalities, we get that

$$\phi(\widehat{f}, I) = \frac{C(a^*)}{C(a_{\mathrm{OPT}})} \leqslant 3 \cdot \frac{1-\alpha+\phi(f, I_p)}{\alpha},$$

as needed. □

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 6. *For $\alpha \in (0, 0.5]$, $k \in \mathbb{N}$, family of instances $\mathcal{I} \in \left\{\mathcal{I}^\alpha_{m,\star}, \mathcal{I}^\alpha_{m,\mathbb{R}^k}, \mathcal{I}^\alpha_{m,sep\text{-}\mathbb{R}^k}\right\}$, and affiliation-independent voting rule f, we have*

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

Since $\phi_{\mathcal{I}}(f) \geqslant 1$ by definition, we can write $\phi_{\mathcal{I}}(\widehat{f}) \leqslant \frac{6}{\alpha} \cdot \phi_{\mathcal{I}}(f)$. In other words, for a fixed party system (i.e., the parties contain some share of the population), every affiliation-independent voting rule f's 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.

Notice that we do not use the assumption that both parties use the same voting rule f in their primaries: if each party p uses a different voting rule f_p, then $\max_{p\in\{-1,1\}}\phi(f, I_p)$ in Theorem 5 can simply be replaced by $\max_{p\in\{-1,1\}}\phi(f_p, I_p)$. Thus, the distortion under the primary system can still be bounded in terms of the largest distortion of the two primary election systems. Additionally, we show in Section 8 (and Appendix A) 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. The proof of this generalized version is almost as straightforward as the proof of Theorem 5 courtesy of Lemma 2.

6. Without Party Separability, Large Primaries are not Better

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 will show that it cannot really reduce the distortion either.

We will show that a distortion reached with m candidates without primaries will not be higher than can be reached by primaries with $m+1$ candidates (though, if there is a candidate in the original m that can be removed without changing the outcome, our result holds for m candidates as well). This holds even if we require each party to have at least a constant fraction of the voters.

Note that this result is different in character than Theorem 5 because it is a worst-case comparison instead of an instance-wise comparison. However, it still applies to all voting rules f, when what we know of the metric space is limited (either it is $\star$, i.e., we do not know anything about its metric, ρ, or π, or it is $\mathbb{R}^k$).

Theorem 7. *For $\alpha \in [0, 0.5]$, $k \in \mathbb{N}$, $\mathcal{M} \in \{\star, \mathbb{R}^k\}$, and affiliation-independent voting rule f, we have*

$$\phi_{\mathcal{I}^{\alpha}_{m+1,\mathcal{M}}}(\widehat{f}) \geqslant \phi_{\mathcal{I}^{0.5}_{m+1,\mathcal{M}}}(\widehat{f}) \geqslant \phi_{\mathcal{I}^{\alpha}_{m,\mathcal{M}}}(f).$$

This result holds even when the underlying space is not metric (i.e., the triangle inequality is not satisfied). For metric spaces, the result can be improved to

$$\phi_{\mathcal{I}^{\alpha}_{m+1,\mathcal{M}}}(\widehat{f}) \geqslant \phi_{\mathcal{I}^{0.5}_{m+1,\mathcal{M}}}(\widehat{f}) \geqslant 2\phi_{\mathcal{I}^{\alpha}_{m,\mathcal{M}}}(f) - 1 \geqslant \phi_{\mathcal{I}^{\alpha}_{m,\mathcal{M}}}(f),$$

where the final inequality holds because the distortion $\phi_{\mathcal{I}^{\alpha}_{m,\mathcal{M}}}(f)$ is at least 1.

Proof. We want to show that for every instance $I \in \mathcal{I}^{\alpha}_{m,\mathcal{M}}$, there exists an instance $I' \in \mathcal{I}^{0.5}_{m+1,\mathcal{M}}$ such that $\phi(\widehat{f}, I') \geqslant 2\phi(f, I) - 1$.

Fix an instance $I = (V, A, M, \rho, \pi) \in \mathcal{I}^{\alpha}_{m,\mathcal{M}}$. Let $a_{\text{OPT}} \in A$ denote an optimal candidate in I, and $a^* = f(E^I)$. Note that $\phi(f, I) = \frac{C^I(a^*)}{C^I(a_{\text{OPT}})}$. Construct instance $I' = (V', A', M, \rho', \pi') \in \mathcal{I}^{0.5}_{m+1,\mathcal{M}}$ as follows:

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

We have $|V'_{-1}| = |V'_1| = |V'|/2$, which satisfies the constraint corresponding to every $\alpha \in [0, 0.5]$.

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 $f(E^{I'_{-1}}) = f(E^I) = a^*$. For I'_1, because $\widetilde{a}^*$ is the only candidate in party 1, it is the primary winner of that party. Since $\rho'(\widetilde{a}^*) = \rho(a^*)$, we are indifferent as to which of them is the general election winner (both have identical social cost). For simplicity, we shall assume that a^* is the general election winner.

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

$$\phi(\widehat{f}, I') = \frac{C^{I'}(a^*)}{C^{I'}(a_{\text{OPT}})} \geqslant \frac{C^I(a^*)}{C^I(a_{\text{OPT}})} = \phi(f, I). \tag{3}$$

Note that so far, the triangle inequality has not been used. Hence, this result holds for arbitrary spaces (not just metric spaces).

For metric spaces, this can be improved further.[5] Thanks to the triangle inequality on I, $|V|d(a_{\mathrm{OPT}}, a^*) \geqslant C^I(a^*) - C^I(a_{\mathrm{OPT}})$. Together, they yield

$$\begin{aligned}\phi(\widehat{f}, I') =& \frac{C^{I'}(a^*)}{C^{I'}(a_{\mathrm{OPT}})} = \frac{C^I(a^*) + |V| \cdot d(a_{\mathrm{OPT}}, a^*)}{C^I(a_{\mathrm{OPT}})} \\ \geqslant& \frac{C^I(a^*) + C^I(a^*) - C^I(a_{\mathrm{OPT}})}{C^I(a_{\mathrm{OPT}})} \\ =& \frac{2C^I(a^*) - C^I(a_{\mathrm{OPT}})}{C^I(a_{\mathrm{OPT}})} = 2\phi(f, I) - 1. \end{aligned} \tag{4}$$

Note that in cases in which it is possible to remove a candidate from A while keeping the winner the same (for example, when all voters rank the same candidate in the last place), removing that candidate will maintain the distortion result with m overall candidates, i.e., we will have $\phi_{\mathcal{I}^\alpha_{m,\mathcal{M}}}(\widehat{f}) \geqslant \phi_{\mathcal{I}^\alpha_{m,\mathcal{M}}}(f)$. □

7. The Advantages of Party Separability

We now turn to analysing the separable case, as analysing $\mathcal{I}_{m,\text{sep-}\mathbb{R}^k}$ is not as straightforward as in the previous section. Previously, in the proof of Theorem 7, we co-located the new voters affiliated with party 1 and a_{OPT} affiliated with party -1. This was allowed because for $\mathcal{M} \in \{\star, \mathbb{R}^k\}$, we have no constraints on the embedding.

In the separable case we need the voters and candidates affiliated with one party to be separated from those affiliated with the other. Hence, this "trick" 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 $\rho(V \cup A)$. That way, we could locate the voters as close as we would like to a_{OPT}, i.e., set $0 < \epsilon < \frac{1}{n^2}$, and locate a_{OPT} within ϵ of the convex hull of $\rho(V \cup A)$, and all voters of party 1 are located at a point b such that $\min_{a \in \rho(V \cup A)} |a - b| < \epsilon$ and $|a_{\mathrm{OPT}} - b| < 2\epsilon$.

Notice that because distortion is a worst-case metric, we would only need to show we could do this in an instance in which the direct distortion for f was a *worst-case instance*. That is, we are finding one $I \in \mathcal{I}^\alpha_{m,\text{sep-}\mathbb{R}^k}$ such that $\phi(f, I) = \phi_{\mathcal{I}^\alpha_{m,\text{sep-}\mathbb{R}^k}}(f)$. Equation (4) would then yield the desired result.

More generally, it is sufficient if, given any $\epsilon > 0$, we can find an instance I such that $\phi(f, I) \geqslant \phi_{\mathcal{I}^\alpha_{m,\text{sep-}\mathbb{R}^k}}(f) - \epsilon$ and a_{OPT} is at distance at most ϵ from the boundary of the convex hull of $\rho(V \cup A)$, as we could then construct a case as in Theorem 7 while maintaining party separability.

Interestingly, Anshelevich et al. [13] 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 8. *Let f be plurality or Borda. Then, $\phi_{\mathcal{I}^{0.5}_{m+1,sep\text{-}\mathbb{R}}}(\widehat{f}) \geqslant \phi_{\mathcal{I}^\alpha_{m,\star}}(f)$.*

However, known worst cases for the Copeland rule [13] and STV [15] 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 8 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 $\alpha \in (0, 0.5]$. 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. Not only are they single-peaked, but for instance, it is known that when voters and candidates are embedded on a line, there always exists a weak Condorcet

[5] We thank an anonymous reviewer for this improvement.

winner [53], and selecting such a candidate results in a distortion of at most 3 [13]. Hence, we have $\phi_{\mathcal{I}_{m,\text{sep-}\mathbb{R}}}(f) = 3$ for every Condorcet-consistent, affiliation-independent voting rule f.[6]

Our aim in this section is to construct an affiliation-independent voting rule f_{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 1. Let f_{fail} be an affiliation-independent voting rule that operates on election $E = (V, A, \overrightarrow{\succ})$ as follows. Let $A = \{a_1, \ldots, a_m\}$, and $t = (m+1)/2$.

- *Special Case*: If $m \geqslant 9$, m is odd, $n \geqslant m^2$, and $\overrightarrow{\succ}$ 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 $a_3 \succ_4 a_m \succ_4 a_2 \succ_4 a_1$.
  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 $m \geqslant 9$ 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 $n \geqslant m^2$ will be clear later.

We will now establish that a worst-case instance of f_{fail} falls under the special case; for this instance, we need to show that a_t is socially optimal; that f_{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 9. *Let $I = (V, A, M, \rho, \pi) \in \mathcal{I}_{m,sep\text{-}\mathbb{R}}$ be an instance for which the corresponding election E^I falls under the special case of f_{fail}. Then the following holds.*

1. *Either $\rho(a_1) \leqslant \ldots \leqslant \rho(a_m)$, or $\rho(a_1) \geqslant \ldots \geqslant \rho(a_m)$, or $|\rho(A)| = 2$.*
2. *If $|\rho(A)| \neq 2$, $\min\{d(a_t, a_1), d(a_t, a_m)\} \geqslant \frac{d(a_1, a_m)}{4}$.*

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 an 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 at an equal distance from all candidates (which means all candidates are located in the same location, or some are at the same 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., $\rho(a_1) \leqslant \ldots \leqslant \rho(a_m)$ or $\rho(a_1) \geqslant \ldots \geqslant \rho(a_m)$. If voter 1 is further away from all candidates (i.e, if $\rho(a_1) = \max_{a \in A} \rho(a)$, then $\rho(v_1) > \rho(a_1)$; if $\rho(a_1) = \min_{a \in A} \rho(a)$, then $\rho(v_1) < \rho(a_1)$), the ordering of the candidates is as voter 1 orders them. Otherwise, let ℓ 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 $\ell' > \ell$,

[6] Anshelevich et al. [13] also proved that no affiliation-independent (deterministic) voting rule can have distortion better than 3, even with respect to $\mathcal{I}_{m,\text{sep-}\mathbb{R}}$.

$\rho(a_{\ell'}) \neq \rho(a_1)$, 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 $d(a_t, a_1) \geqslant \frac{d(a_1,a_m)}{4}$. By symmetry, we also obtain $d(a_t, a_m) \geqslant \frac{d(a_1,a_m)}{4}$. Assume $|\rho(A)| \neq 2$, and without loss of generality, let $\rho(a_1) \leqslant \ldots \leqslant \rho(a_m)$ 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

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

This concludes the proof. □

We are now able to show that f_{fail} is a voting rule that has an unbounded distortion ($\mathcal{O}(n)$), but for which the primary system can significantly improve its distortion to $\mathcal{O}(1)$ (constant depending on the share of parties' electorate).

Theorem 10. *For $m \geqslant 9$ and constant $\alpha \in (0, 0.5]$, $\phi_{\mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}}}(\widehat{f}_{fail})$ is upper bounded by a constant depending on α, whereas $\phi_{\mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}}}(f_{fail})$ 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 $n \geqslant m^2$), $A = \{a_1, \ldots, a_m\}$ (m being odd), and $M = (\mathbb{R}, d)$ with d being the Euclidean distance on the line.

The embedding function ρ 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 $\epsilon < \frac{1}{n}$. The voters are embedded as follows.

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

Finally, we are interested in the distortion of an affiliation-independent voting rule f. Hence, the party affiliation is not important (the outcome of f is independent of π). Thus, we can use any π in which half of the voters are affiliated with each party. This construction is all we need since $I \in \mathcal{I}^{0.5}_{m,\text{sep-}\mathbb{R}} \subseteq \mathcal{I}^{\alpha}_{m,\text{sep-}\mathbb{R}}$ 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_{fail}. Hence, $f_{\text{fail}}(E^I) = a_1$. Note that $C(a_1) \geqslant (2n - 2m) \cdot |0 - \frac{1}{2} + \frac{1}{n}| > n - m - 2$ because a_1 is at distance at least $\frac{1}{2} - \frac{1}{n}$ from all but $2m$ voters located close to $\frac{1}{2}$. In contrast, $C(a_t) \leqslant 2m \cdot 1 + (2n - 2m) \cdot \frac{1}{n} \leqslant 2m + 2$ because a_t is at distance at most $\epsilon < \frac{1}{n}$ from all but $2m$ voters (and at distance at most 1 from those $2m$ voters). Thus, $\phi(f_{\text{fail}}, I) \geqslant (n - m - 2)/(2m + 2)$. 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 dependent on α. 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, too. First, assume $|\rho(A)| \neq 2$ (we will handle the case $|\rho(A)| = 2$ later). Without loss of generality, assume that $\rho(a_1) \leqslant \ldots \leqslant \rho(a_m)$, 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 5, $\phi(\widehat{f}, I)$ depends only α and on the distortion of f on the primary instance of the party with which a_{OPT} is affiliated. Hence, if primary election $E^{I_{-1}}$ does not fall under the special case of f_{fail}, then f_{fail} applies a Condorcet-consistent rule on I_{-1}, ensuring that $\phi(f, I_{-1})$ is at most 3. In this case, by Theorem 5, $\phi(\widehat{f}, I)$ is also upper bounded by a constant.

Suppose $E^{I_{-1}}$ falls under the special case of f_{fail}. Let $t = (q + 1)/2$ and $d_{-1} = d(a_1, a_q)$. Then, by Lemma 9, $\min\{d(a_t, a_1), d(a_t, a_q)\} \geqslant d_{-1}/4$. 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, $|V_{-1}| \geqslant q^2$, and the final transition holds because $|V_{-1}| \geqslant \alpha n$.

Note that for every $i \in S$ and $j \in V_1$, $d(i, j) \geqslant d_{-1}/8$, and $|V_1| \geqslant \alpha n$. Hence, we have $\Omega(n)$ pairs of voters (i, j) such that $d(i, j) \geqslant d_{-1}/8$. Further, $d(a_{\text{OPT}}, i) + d(a_{\text{OPT}}, j) \geqslant d(i, j)$. Hence, it follows that

$$C(a_{\text{OPT}}) = \Omega(n) \cdot d_{-1}. \tag{5}$$

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

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

The final transition follows from Equation (5). This yields a constant upper bound on $\phi(\widehat{f}_{\text{fail}}, I) = C(a^*)/C(a_{\text{OPT}})$.

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

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

Here, the second transition follows because in the general election, at least $\frac{n}{2}$ voters vote for a^* over a_1 and $d(a_1, a_{\text{OPT}}) \leqslant d_{-1}$, and the final transition follows from Equation (5). This implies

$$C(a_{\text{OPT}}) = \Omega(n \cdot d(a_{\text{OPT}}, a^*)). \tag{6}$$

On the other hand, we have

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

where the last transition follows due to Equation (6). 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, $a_{\text{OPT}} \in \{a^*_{-1}, a^*_1\}$, and since a^* is the majority winner, $a^* = a_{\text{OPT}}$, and the distortion is 1.
2. $\rho(A_{-1}) = \{x_1, x_2\}$ and $\rho(A_1) = \{x_2\}$ (the case where $\rho(A_{-1}) = \{x_1\}$ and $\rho(A_1) = \{x_1, x_2\}$ is symmetric): If $a_{\text{OPT}} \in \{a^*_{-1}, a^*_1\}$, then it is similar to the previous case. Otherwise, this means $\rho(a_{\text{OPT}}) = x_1$ and $\rho(a^*_{-1}) = \rho(a^*_1) = x_2$. Separability means the voters of party 1 lie in $\{x : x \geqslant x_2\}$, so $C(a_{\text{OPT}}) \geqslant |V_1| \cdot d(x_1, x_2) \geqslant \alpha \cdot n \cdot d(x_1, x_2)$. While $C(a^*) \leqslant C(a_{\text{OPT}}) + |V_{-1}| \cdot d(x_1, x_2) \leqslant C(a_{\text{OPT}}) + n(1-\alpha)d(x_1, x_2)$. Combining these two equations we get the distortion is $\frac{1+2\alpha}{\alpha}$.

This completes the proof. □

8. Extension to Multiple Parties and Multi-Stage Systems

So far, we analyzed the primary system as an alternative to the direct system, but in a restricted model that had the following assumptions:

- There are two parties.
- Every candidate belongs to *exactly one* of the two parties.
- Each party holds a primary election, and every voter votes in *exactly one* of the two primaries.
- The two primary winners compete in a general election, where *every* voter votes and the majority rule is used to decide the winner.

In practice, it is common for all of these assumptions to be violated. Often, there are multiple parties and independent voters or candidates (who do not participate in any primary); also, in the general election, only a fraction of the voters may show up to vote. Further, we wish to explore scenarios which involve more complex decision processes that extend beyond a two-stage system. For example, subsets of parties' electorate decide on the candidates for the party, which the party then further cuts down in an iterative process, presenting a single candidate for the general elections.[7]

In this section, we propose a general model that relaxes all of the aforementioned restrictions, and show that even multi-stage elections with multiple parties are not much worse than the direct system. This extends the main result of our paper, Theorem 5, which shows this for the restricted primaries model. In particular, we establish this through a straightforward application of Lemma 2, which was used to establish the special case of Theorem 5, thus demonstrating the power of this lemma.

Notice that even in the two-stage primaries model, when we allow the number of parties k to be larger than two, we need to decide which voting rule will be used in the general election to decide amongst the k primary winners. This is because the majority rule, which we used in the two-party setting, may not give a definite answer in an election between k candidates – it is possible that no candidate will have a majority. We can imagine this process as a rooted tree with $k+1$ nodes and height 2, with the root node representing the general election and its k children representing the k primaries, as shown in Figure 3a. The following definition generalizes this model to allow arbitrary rooted trees. Further, in the election at each node of the tree, it allows an arbitrary subset of the voters to participate, and an arbitrary subset of candidates to be added on top of the winners of the elections at the children nodes.

Definition 2 (Tree-Structured Primary Process)**.** A tree-structured primary process is a rooted tree, in which each node t has an associated tuple (V_t, A_t, f_t) such that $\cup_t A_t = A$. The elections are conducted leaves-to-root: in the election at each node t, voters in $V_t \subseteq V$ vote over candidates in $A_t \subseteq A$ as well as the winners of the elections at the children of node t, and voting rule f_t is used to select a winner a^*_t.

[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.

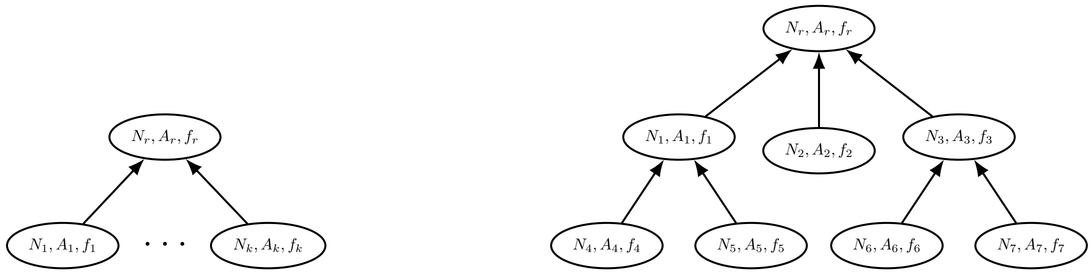

(a) 2-stage with k parties.

(b) Multi-stage.

Figure 3: Examples of tree-structured primary processes.

Formally, if $N(t)$ is the set of children of node t, then the election at node t is over the set of candidates $A_t \cup \{a^*_{t'} : t' \in N(t)\}$. The winner of the election at the root node is declared the overall winner under this process.

An example of a more general tree-structured process is shown in Figure 3b.

Crucially, note that we do not impose *any* restriction on the subsets of voters and candidates at the different nodes of the tree; they can overlap arbitrarily. In other words, each voter may participate in the elections at an arbitrary subset of the nodes, and the same holds for the candidates. The only restriction we impose is that every candidate must participate in at least one election ($\cup_t A_t = A$). This is because if a candidate does not participate in any election, it becomes impossible to bound the social cost of the winner in terms of the (entirely unknown) social cost of this candidate.

Therefore, in the two-stage primary process, this model allows independent voters and candidates. But moreover, it also allows more complex, multi-stage decision-making processes (such as in the British party system, suggested above).

We now show that an analogue of Theorem 5 holds for this more general model.

Theorem 11. *Let I be a tree-structured primary process. Let $\alpha_t = |V_t|/|V|$ denote the fraction of voters participating in the election at node t. Let a^* and a_{OPT} denote the winner and a socially optimal candidate, respectively, and let $(t_1, \ldots, t_k)$ denote a path from the root to a node t_k such that $a_{\mathrm{OPT}} \in A_{t_k}$. Then,*

$$C(a^*) \leqslant \left(\prod_{i=1}^{k} \frac{1 - \alpha_{t_i} + \phi(f_{t_i})}{\alpha_{t_i}} \right) \cdot C(a_{\mathrm{OPT}}).$$

Thus, if $\mathcal{P}$ is the set of all root-to-leaf paths in I, then the distortion under this process is at most $\max_{(t_1,\ldots,t_k)\in\mathcal{P}} \prod_{i=1}^{k} \frac{1-\alpha_{t_i}+\phi(f_{t_i})}{\alpha_{t_i}}$.

Proof. First, we observe that there always exists a node t_k with $a_{\mathrm{OPT}} \in A_{t_k}$ because we have assumed that $\cup_t A_t = A$. As in the statement of the theorem, let $(t_1, \ldots, t_k)$ be a path from the root to such a node. Let $N(t)$ denote the set of children of node t.

For each $i \in [k-1]$, applying Lemma 2 to the election at node t_i, we get

$$C(a^*_{t_i}) \leqslant \frac{1 - \alpha_{t_i} + \phi(f_{t_i})}{\alpha_{t_i}} \cdot \min_{a \in A_{t_i} \cup \{a^*_t : t \in N(t_i)\}} C(a) \leqslant \frac{1 - \alpha_{t_i} + \phi(f_{t_i})}{\alpha_{t_i}} \cdot C(a^*_{t_{i+1}}), \tag{7}$$

where the last transition holds because $t_{i+1} \in N(t_i)$. Similarly, applying Lemma 2 to the election at node t_k, and noticing that $a_{\mathrm{OPT}} \in A_{t_k}$, we get

$$C(a^*_{t_k}) \leqslant \frac{1 - \alpha_{t_k} + \phi(f_{t_k})}{\alpha_{t_k}} \cdot C(a_{\mathrm{OPT}}). \tag{8}$$

Combining Equations (7) and (8), we get the bound with a_{OPT}. The second bound, when the path to a_{OPT} is not known is a trivial consequence of the first bound, simply going over all possible paths. $\square$

To understand this result, let us assume that for some constant $\alpha > 0$, $\alpha_t \geqslant \alpha$ for each node t, i.e., at least a constant fraction of the voters participate in each election. Then, using the fact that $\phi(f) \geqslant 1$ implies $\frac{1-\alpha+\phi(f)}{\alpha} \leqslant \frac{2}{\alpha} \cdot \phi(f)$, we get that the distortion of the tree-structured primary process is at most $\max_{(t_1,\ldots,t_k)\in\mathcal{P}} (2/\alpha)^k \prod_{i=1}^{k} \phi(f_{t_i})$. Note that while the distortion is exponential in the *height* of the tree (a.k.a. the number of stages), in practice it is common to use a process with few stages. On the other hand, the important observation is that the distortion is independent of the *width* of the tree.

Another important observation is what Theorem 11 implies for the two-stage primary process illustrated in Figure 3a, which allows independent voters and candidates and any number of parties. Suppose we use the same voting rule f in all the primaries and in the general election, and suppose that at least α (resp. γ) fraction of the voters participate in each primary (resp. in the general election). Then, the distortion of the overall process is at most $\frac{4}{\alpha\gamma}\phi(f)^2$. While, somewhat surprisingly, this is independent of the number of parties, we would expect α to be a decreasing function of the number of parties, and hence the distortion bound would linearly grow in the number of parties.

9. 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 focus on the two-stage, two-party primary process which has been the main setting in 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 (i.e., do not vote in any party primary): default = 0%, range = 0% to 90% in increments of 10%.

Independent candidates The percentage of candidates who are independent (i.e., are not a candidate in any party primary): 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 $x : (1-x)$, we find a hyperplane dividing voters into this $x : (1-x)$ ratio. Due to symmetry, we simply find a threshold t on the k^{th} coordinate such that the locations of x

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

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

fraction of the voters (call this set V_{-1}) have k^{th} coordinate at most t, while the locations of the rest (call this set V_1) have k^{th} coordinate at least t. We do not set 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) – rounded down – 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 the separable instance, we assign $V_{-1} \cup A_{-1}$ to party -1, and assign $V_1 \cup A_1$ to party 1. In the other instance (called "random"), we assign $|V_{-1}|$ voters 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 the distortion of 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-separable, 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" [54, 55, 56, 57, 58].

9.1. Primary versus Direct: Summary

	Primary is better	Direct is better	No significant difference
plurality separable	183	7	18
plurality random	203	0	5
STV separable	8	178	22
STV random	152	13	48
Borda separable	2	202	4
Borda random	8	179	21
Maximin separable	0	207	1
Maximin random	0	208	0
Copeland separable	0	207	1
Copeland random	0	208	0

Table 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-value equal to 0.05.

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 (separable or random), and evaluate which system results in a better average distortion. For statistical significance, we use the paired t-test with p-value equal to 0.05. The results are presented in Table 1.

Without party separability (i.e. in the random case), our theoretical results indicate that the primary system is no better than the direct system in the worst case (Theorem 7). While this is also true in our experimentally generated average cases for Borda, Copeland, and Maximin, we see that for plurality and STV, the primary system almost always outperforms direct voting in the average case (showing the downside of distortion – just like price of anarchy and other metrics – as a worst-case metric).

With party separability, 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 8),

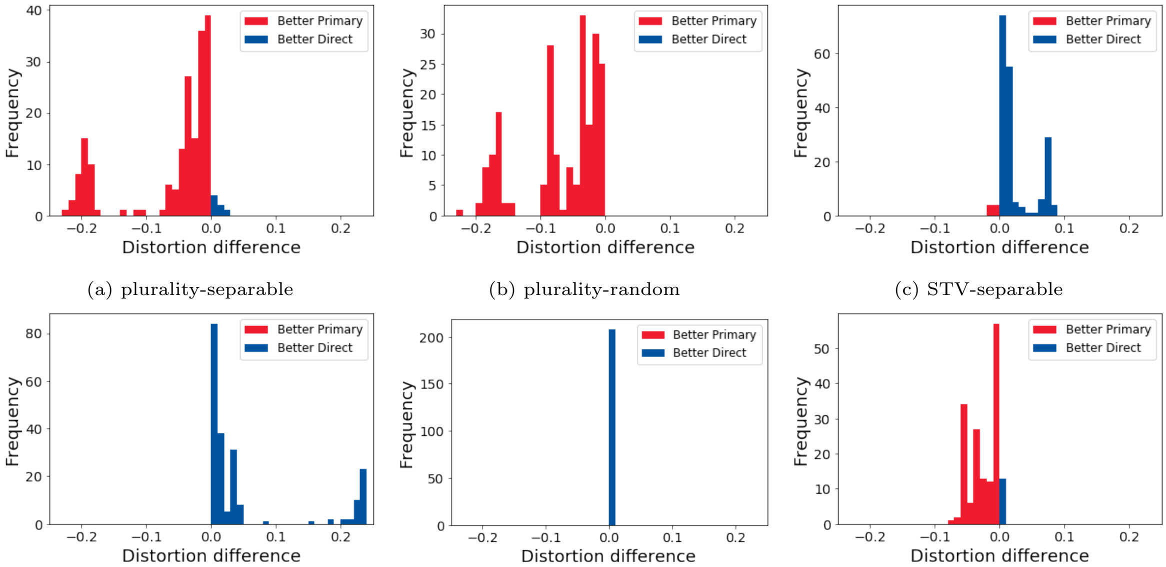

(a) plurality-separable (b) plurality-random (c) STV-separable

(d) Maximin-separable (e) Maximin-random (f) STV-random

Figure 4: These histograms show the difference between the average distortion under the primary system and that under the direct system.

we see that this is not true in the average case.

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

9.2. Primary versus Direct: Margins

Table 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 4. 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-separable 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-separable). Hence, neither system seems to have a significant advantage over the other in terms of the difference in average distortions.

9.3. 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 [10], 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.

9.3.1. The Number of Voters

Figure 5 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.

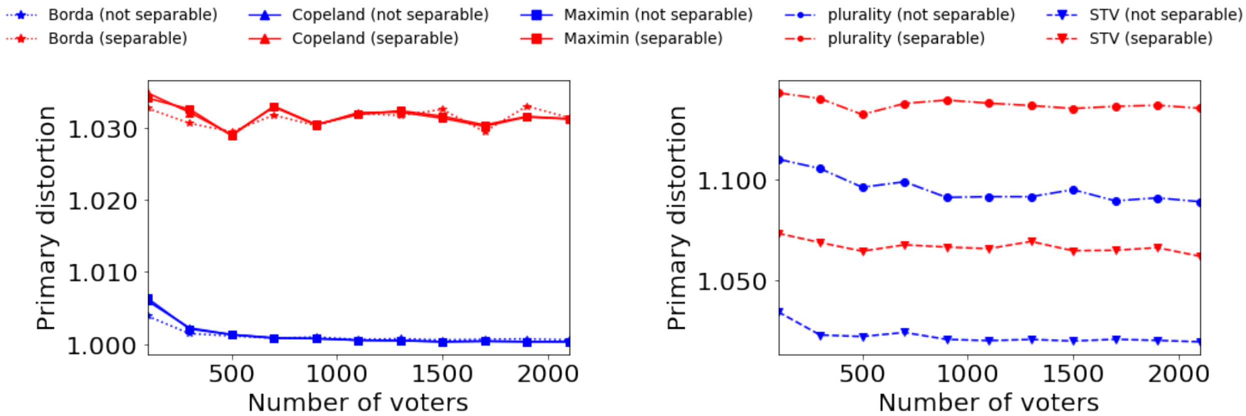

(a) Borda, Copeland, and Maximin

(b) Plurality and STV

Figure 5: The average distortion under the primary system as a function of the number of voters.

9.3.2. The Number of Candidates

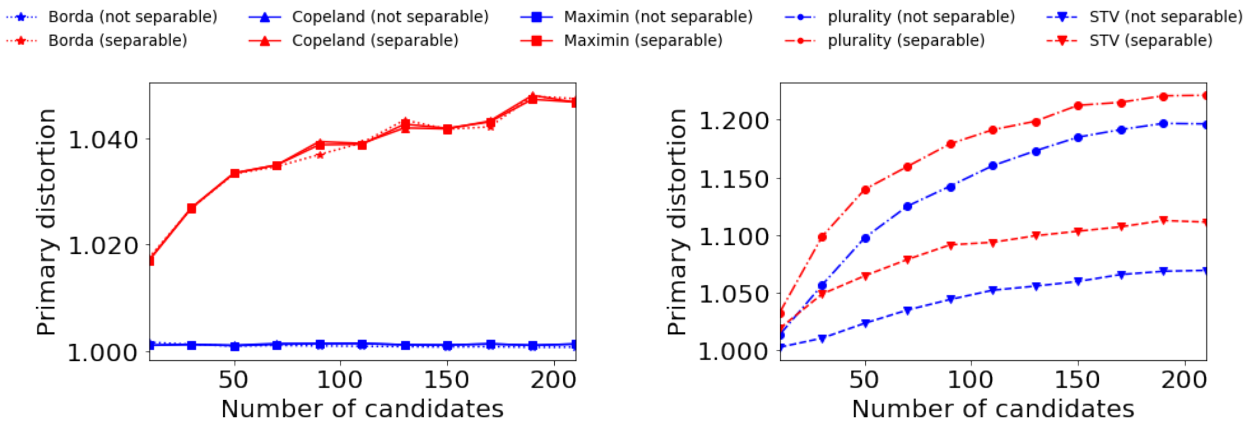

(a) Borda, Copeland, and Maximin

(b) Plurality and STV

Figure 6: 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 6. 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 seprable 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.

9.3.3. The relative number of voters in each party

The relative size of the parties has an interesting effect on the average distortion. As both sub-figures of Figure 7 show, there is 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.

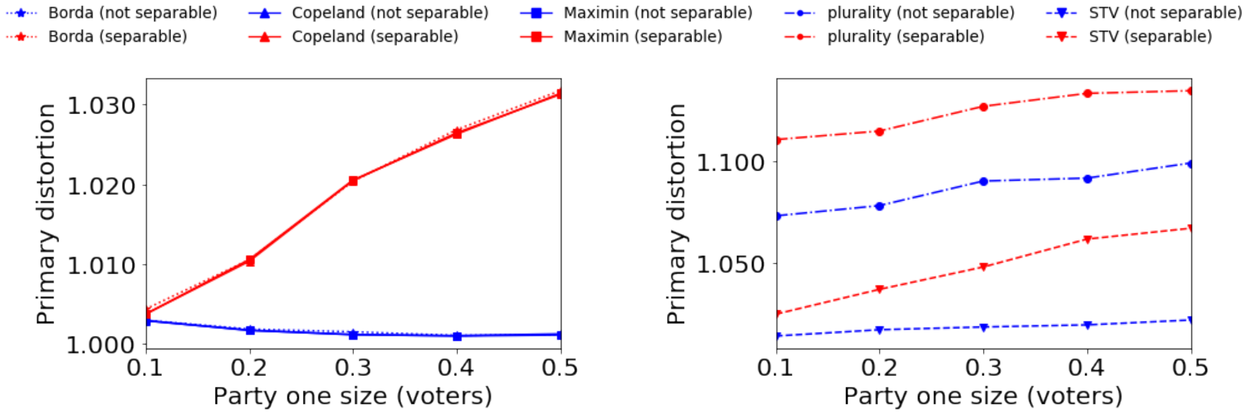

(a) Borda, Copeland, and Maximin

(b) Plurality and STV

Figure 7: The average distortion under the primary system as a function of the fraction of the 500 voters in the first party.

9.3.4. The number of candidates in each party

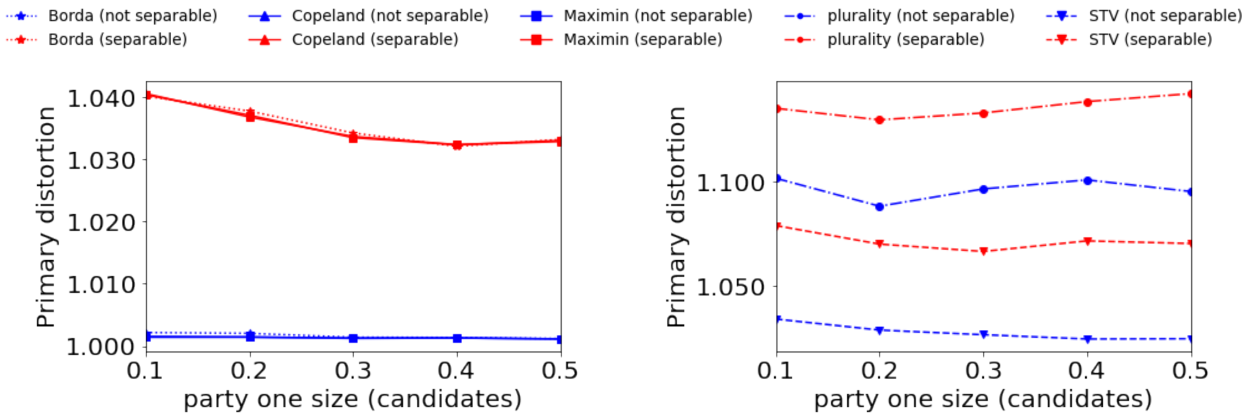

(a) Borda, Copeland, and Maximin

(b) Plurality and STV

Figure 8: The average distortion under the primary system as a function of the fraction of the 50 candidates in the first party.

As Figure 8 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 were the Condorcet consistent rules (and Borda), which show a slight decrease in distortion as the parties became closer in terms of number of candidates.

9.3.5. 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 9, 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.

9.3.6. 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

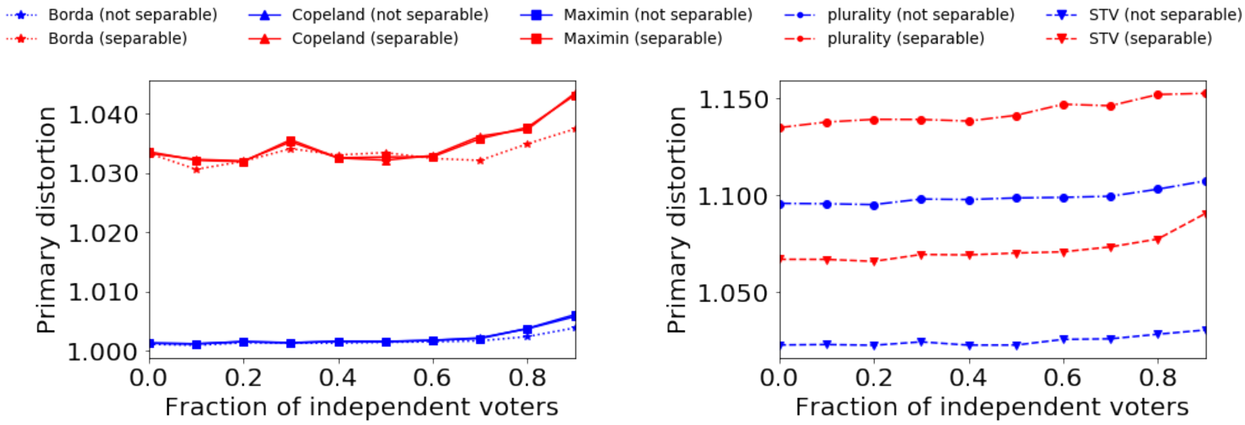

(a) Borda, Copeland, and Maximin

(b) Plurality and STV

Figure 9: The average distortion under the primary system as a function of the percentage of independent voters.

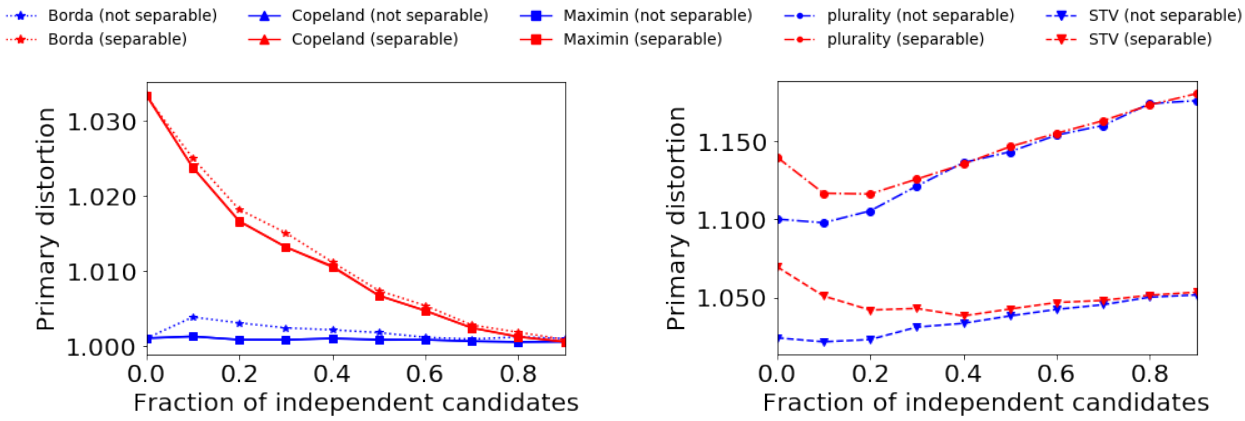

(a) Borda, Copeland, and Maximin

(b) Plurality and STV

Figure 10: The average distortion under the primary system as a function of the percentage of independent candidates.

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 10, 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 separable and random cases.

9.3.7. The Dimension of the Metric Space

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 11 shows as the dimension of the embedding increases the distortion goes down in all. While Figure 11 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.

10. Discussion

In this paper we explored the topic of multi-stage elections, in which, prior to the general elections, a subset of voters can decide which candidates will continue to the next stage. In particular, we focus on the

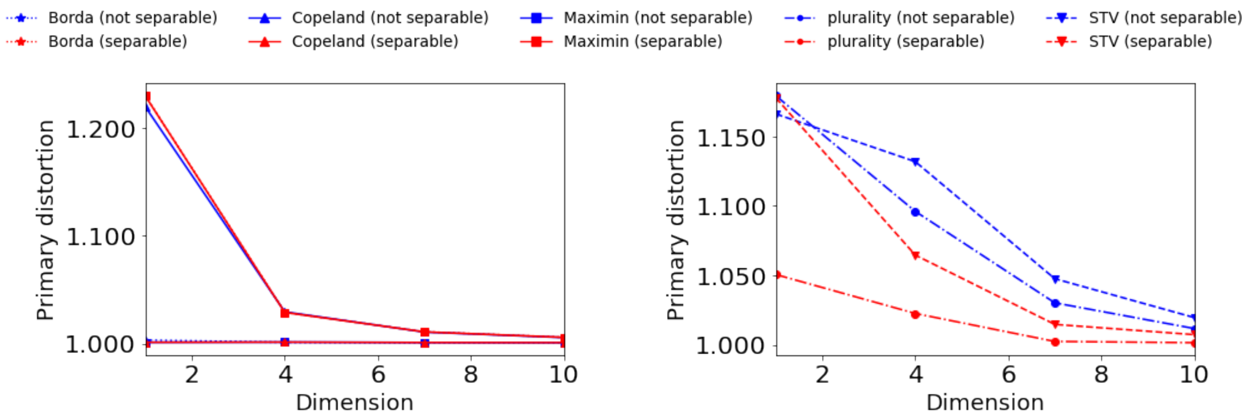

(a) Borda, Copeland, and Maximin

(b) Plurality and STV

Figure 11: The average distortion under the primary system as a function of the dimension of the party embedding.

two-stage primary model, used in various party systems throughout the world, in which parties vote on who their leader shall be, and subsequently these parties' leaders contest the general election.

Using this model we show primary elections cannot produce significantly worse results than direct elections, if the parties contain a large enough share of the population. In addition, for many natural voting rules we also find that primary systems might produce a worst case distortion which is about the same as in the direct system. However, we are able to construct a voting rule and setting in which the primary system has an unbounded advantage (again in the worst case) over the corresponding direct system. All in all, our theoretical results show primaries can in some instances improve the election system's distortion, and they cannot hurt the distortion too much.

Examining voting rules empirically, we see that plurality – the most widely used voting rule – benefits most from using primaries. This might be since in plurality, a winner can emerge from advantages stemming from a small clustering of supporters in the direct elections, when plenty of candidates mean a winner can have very small support. When candidates are eliminated in the stages, a very small basis of support is much harder to sustain against winners of other parties, which can each command a significant portion of voters. This does not hold for several other voting rules, mainly those which are Condorcet consistent, in which parties can eliminate the Condorcet winner.

Our paper initiates the novel quantitative study of multi-stage elections (and their comparison to single-stage elections), but leaves plenty to explore. 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 results 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 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 2) in Section 8, what tree-graph structures produce better results? What are voting rule combinations that work well together? Examining the use of multiple and different voting rules as Narodytska and Walsh [59] do 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.

Another potential avenue for comparing primaries with direct elections is the axiomatic approach, where one identifies qualitative desiderata and seeks voting rules satisfying them. This approach may help us identify the exact properties that make primary elections more desirable than direct ones (or vice versa), and also separate different types of primary elections.

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.

11. Acknowledgements

We thank the anonymous reviewers of this paper for their careful and constructive comments, which have greatly improved the quality of this paper.

This work was partially supported by NSERC under the Discovery Grants program, and Israel Science Fund grants #1965/20 and #3152/20.

Appendix A. Reduction to Primaries

In Section 5, we show that the distortion of any voting rule is not much worse under the primary system (given that the primaries are large) than under the direct system. However, this analysis assumes a *restricted primary system* in which each voter participates in exactly one of two primaries and all voters participate in the general election. In Section 8, we introduce a more *relaxed primary system* and show that the result continues to hold. However, this result bounds the distortion of a voting rule under the relaxed primary system in terms of its distortion under the direct system (see Theorem 11).

In Section 7, we show that for some voting rules, their distortion under the restricted primary system can be much better than their distortion under the direct system. For such voting rules, can we relax the assumptions, but instead of reducing to the direct system, reduce to the restricted primary system? In other words, can we bound their distortion under the relaxed primary system in terms of their distortion under the restricted primary system, given that the latter is known to be much better than their distortion under the direct system? In this section, we show that this is possible for a milder relaxation than in Section 8.

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 αn 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 αn voters participate in each primary election, at least βn of voters participated in the primaries (obviously, $\beta \geqslant 2\alpha$), and at least

γn voters participate in the general election. Formally, we require $|\pi^{-1}(p)| \geqslant \alpha n$ for each $p \in \{-1, 1\}$, $|\pi^{-1}(-1) \cup \pi^{-1}(1)| \geqslant \beta n$ and $|\tau^{-1}(1)| \geqslant \gamma n$. 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.[12]
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 $\phi_{\mathcal{I}^{\alpha,\beta,\gamma}_{m,\mathcal{M}}}(\widehat{f}) \geqslant \phi_{\mathcal{I}^{\alpha}_{m,\mathcal{M}}}(\widehat{f})$ follows trivially from the fact that the relaxed framework is strictly more general than the restricted framework.

We begin by noticing that Lemma 2, applied on the general election in this framework, yields the following comparison between the social cost of the general election winner and the social cost of the two primary winners. Here, we use the fact that the distortion of the majority rule is 3 [13].

Corollary 12. *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,*

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

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

Theorem 13. *For every choice of $\mathcal{M}$, $m \in \mathbb{N}$, affiliation-independent voting rule f, and constants $\alpha, \beta > 0$,*

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

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 $|\overline{V}| \geqslant \beta n$. Let $\overline{\pi} : \overline{V} \cup A \to \{-1, 1\}$ be the restriction of π to $\overline{V} \cup A$ (note that π does not map any voter in $\overline{V}$ to 0, and cannot map any candidate in A to 0). Note that at least $\alpha n \geqslant \alpha|\overline{V}|$ 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.[13] Hence, the primary winners in instances I and $\overline{I}$ must be identical.

[12] While data indicates that almost all voters who participate in the primary elections also participate in the general election [33], we consider this possibility for the sake of completeness.

[13] 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.

Let a_{-1} and a_1 denote the primary winners of parties -1 and 1, respectively, in both I and $\overline{I}$. Let $a_{\text{OPT}} \in \arg\min_{a \in A} C^I(a)$ and $\overline{a_{\text{OPT}}} \in \arg\min_{a\in A} C^{\overline{I}}(a)$ 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, $\phi(\widehat{f}, I) = \frac{C^I(a^*)}{C^I(a_{\text{OPT}})}$ and $\phi(\widehat{f}, \overline{I}) = \frac{C^{\overline{I}}(\overline{a}^*)}{C^{\overline{I}}(\overline{a_{\text{OPT}}})}$.

For any candidate a, we have

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

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

$$d(\overline{a_{\text{OPT}}}, a) \leqslant \frac{2}{\beta n} C^{\overline{I}}(a). \tag{A.1}$$

Now, we have

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

where the first inequality follows from the fact that both a^* and $\overline{a^*}$ are primary winners in I and from Corollary 12, the third and the fourth inequalities follow from the triangle inequality, the fifth inequality uses Equation (A.1), and the final inequality uses the fact that $C^{\overline{I}}(a) \leqslant C^I(a)$ 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

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

where the first inequality uses Equation (A.2), and the second inequality uses the fact that $C^I(a_{\mathrm{OPT}}) \geqslant C^{\overline{I}}(a_{\mathrm{OPT}}) \geqslant C^{\overline{I}}(\overline{a_{\mathrm{OPT}}})$. Thus, we have the desired result. $\square$

Since having independent voters in the general election and having voters who only participate in the primaries cannot improve the distortion of a voting rule under the primary system, our earlier results which establish that primaries are no better than direct elections (Theorems 1 and 7) continue to hold in this extended framework. Additionally, we proved that the existence of such voters will not increase the distortion under the primary system by more than a constant factor (which depends on β and γ). Hence, our earlier results which establish that primaries cannot be much worse than direct elections (Theorem 5) or that primaries can be significantly better than direct elections (Theorem 10) also continue to hold in this extended framework.

References

[1] A. Hamilton, J. Madison, J. Jay, The federalist papers, The Independent Journal, 1787.

[2] B. Simms, Three victories and a defeat: The rise and fall of the first British empire, 1714-1783, Allen Lane, 2007.

[3] S. Wilentz, The Rise of American Democracy, Norton, 2005.

[4] W. P. Cross, A. Blais, Who selects the party leader?, Party Politics 18 (2011) 127–150.

[5] M. Cohen, D. Karol, H. Noel, J. Zaller, The party decides: Presidential nominations before and after reform, University of Chicago Press, 2008.

[6] V. Conitzer, M. Yokoo, Using mechanism design to prevent false-name manipulations, AI Magazine 31 (2010) 65–77.

[7] A. D. Procaccia, J. S. Rosenschein, The distortion of cardinal preferences in voting, in: Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA), Edinburgh, Great Britain, 2006, pp. 317–331.

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

[9] A. Borodin, O. Lev, N. Shah, T. Strangway, Primarily about primaries, in: Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI), Honolulu, Hawaii, 2019, pp. 1804–1811.

[10] F. Brandt, V. Conitzer, U. Endriss, J. Lang, A. D. Procaccia (Eds.), Handbook of computational social choice, Cambridge University Press, 2016.

[11] A. Downs, An economic theory of democracy, Harper and Row, 1957.

[12] N. Schofield, The spatial model of politics, number 95 in Routledge Frontiers of Political Economy, Routledge, 2008.

[13] E. Anshelevich, O. Bhardwaj, J. Postl, Approximating optimal social choice under metric preferences, in: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI), Austin, Texas, 2015, pp. 777–783.

[14] E. Anshelevich, J. Postl, Randomized social choice functions under metric preferences, Journal of Artificial Intelligence Research 58 (2017) 797–827.

[15] P. Skowron, E. Elkind, Social choice under metric preferences: Scoring rules and STV, in: Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI), San Francisco, California, 2017, pp. 706–712.

[16] V. Gkatzelis, D. Halpern, , N. Shah, Resolving the optimal metric distortion conjecture, in: Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS), Durham, North Carolina (online), 2020, pp. 1427–1438.

[17] F. E. Kizilkaya, D. Kempe, Plurality veto: A simple voting rule achieving optimal metric distortion, in: Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI), 2022, pp. 349–355.

[18] F. E. Kizilkaya, D. Kempe, Generalized veto core and a practical voting rule with optimal metric distortion, in: Proceedings of the 24th ACM Conference on Economics and Computation (EC), 2023, pp. 913–936.

[19] M. Feldman, A. Fiat, I. Golomb, On voting and facility location, in: Proceedings of the 17th ACM Conference on Economics and Computation (EC), Maastricht, The Netherlands, 2016, pp. 269–286.

[20] E. Anshelevich, A. Filos-Ratsikas, N. Shah, A. A. Voudouris, Distortion in social choice problems: The first 15 years and beyond, in: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), 2021, pp. 4294–4301.

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

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

[23] M. Brill, V. Conitzer, Strategic voting and strategic candidacy, in: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI), Austin, Texas, 2015, pp. 819–826.

[24] M. Polukarov, S. Obraztsova, Z. Rabinovich, A. Kruglyi, N. R. Jennings, Convergence to equilibria in strategic candidacy, in: Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI), Buenos Aires, Argentina, 2015, pp. 624–630.

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

[26] J. Rothe (Ed.), Economics and computation, Springer, 2015.

[27] O. Kenig, Classifying party leaders' selection methods in parliamentary democracies, Journal of Elections, Public Opinion and Parties 19 (2009) 433–447.

[28] W. P. Cross, R. S. Katz (Eds.), The challenges of intra-party democracy, Oxford University Press, 2013.

[29] R. Jobson, M. Wickham-Jones, Gripped by the past: Nostalgia and the 2010 Labour party leadership contest, British Politics 5 (2010) 525–548.

[30] B. Wauters, Explaining participation in intra-party elections: Evidence from Belgian political parties, Party Politics 16 (2010) 237–259.

[31] R. Y. Hazan, The 1996 intra-party elections in Israel: Adopting party primaries, Electoral Studies 16 (1997) 95–103.

[32] H. Norpoth, From primary to general election: A forecast of the presidential vote, PS: Political Science & Politics 37 (2004) 737–740.

[33] J. Sides, C. Tausanovitch, L. Vavreck, C. Warshaw, On the representativeness of primary electorates, British Journal of Political Science (2018) 1–9.

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

[35] A. Kahng, S. Mackenzie, A. D. Procaccia, Liquid democracy: An algorithmic perspective, in: Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI), New Orleans, Louisiana, 2018, pp. 1095–1102.

[36] Y. Bachrach, O. Lev, Y. Lewenberg, Y. Zick, Misrepresentation in district voting, in: Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI), New York City, New York, 2016, pp. 81–87.

[37] S. Bervoets, V. Merlin, Gerrymander-proof representative democracies, International Journal of Game Theory 41 (2012) 473–488.

[38] R. van Bevern, R. Bredereck, J. Chen, V. Froese, R. Niedermeier, G. J. Woeginger, Network-based vertex dissolution, SIAM Journal on Discrete Mathematics 29 (2015) 888–914.

[39] Y. Lewenberg, O. Lev, J. S. Rosenschein, Divide and conquer: Using geographic manipulation to win district-based elections, in: Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), São-Paulo, Brazil, 2017, pp. 624–632.

[40] W. Pegden, A. D. Procaccia, D. Yu, A partisan districting protocol with provably nonpartisan outcomes, ArXiv:1710.08781, 2017.

[41] A. Borodin, O. Lev, N. Shah, 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), Stockholm, Sweden, 2018, pp. 98–104.

[42] O. Lev, Y. Lewenberg, "Reverse gerrymandering": A decentralized model for multi-group decision making, in: Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI), Honolulu, Hawaii, 2019, pp. 2069–2076.

[43] H.-K. Ahn, S.-W. Cheng, O. Cheong, M. Golin, R. van Oostrum, Competitive facility location: The Voronoi game, Theoretical Computer Science 310 (2004) 457–467.

[44] C. Dürr, N. K. Thang, Nash equilibria in Voronoi games on graphs, in: Proceedings of the 15th Annual European Conference on Algorithms (ESA), Eilat, Israel, 2007, pp. 17–28.

[45] M. Mavronicolas, B. Monien, V. G. Papadopoulou, F. Schoppmann, Voronoi games on cycle graphs, in: Proceedings of the 33rd International Symposium on Mathematical Foundations of Computer Science (MFCS), Torun, Poland, 2008, pp. 503–514.

[46] S. Bandyapadhyay, A. Banik, S. Das, H. Sarkar, Voronoi game on graphs, Theoretical Computer Science 562 (2015) 270–282.

[47] M. Boppana, R. Hod, M. Mitzenmacher, T. Morgan, Voronoi choice games, in: Proceedings of the 43rd International Colloquium on Automata, Languages and Programming (ICALP), Rome, Italy, 2016, pp. 23:1–23:13.

[48] P. Harrenstein, G. Lisowski, R. Sridharan, P. Turrini, A Hotelling-Downs framework for party nominees, in: Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), London, Great Britain (online), 2021, pp. 593–601.

[49] A. Karpov, O. Lev, S. Obraztsova, Hotelling-Downs equilibria: Moving beyond plurality variants, in: Proceedings of the 22nd International Conference on Group Decision and Negotiation (GDN), 2022, pp. 3–16.

[50] N. Nisan, T. Roughgarden, E. Tardos, V. V. Vazirani (Eds.), Algorithmic game theory, Cambridge University Press, 2007.

[51] E. Koutsoupias, C. Papadimitriou, Worst-case equilibria, in: Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS), volume 1563, Trier, Germany, 1999, pp. 404–413.

[52] K. O. May, A set of independent necessary and sufficient conditions for simple majority decision, Econometrica 20 (1952) 680–684.

[53] D. Black, On the rationale of group decision-making, Journal of Political Economy 56 (1948) 23–34.

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

[55] 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, 1985.

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

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

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

[59] N. Narodytska, T. Walsh, Manipulating two stage voting rules, in: Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), St. Paul, Minnesota, 2013, pp. 423–430.