The Distortion of Distributed Voting*

Aris Filos-Ratsikas[1], Evi Micha[2], and Alexandros A. Voudouris[3]

[1] École polytechnique fédérale de Lausanne, Switzerland
aris.filosratsikas@epfl.ch
[2] University of Toronto, Canada
emicha@cs.toronto.edu
[3] University of Oxford, UK
alexandros.voudouris@cs.ox.ac.uk

Abstract. Voting can abstractly model any decision-making scenario and as such it has been extensively studied over the decades. Recently, the related literature has focused on quantifying the impact of utilizing only limited information in the voting process on the societal welfare for the outcome, by bounding the *distortion* of voting rules. Even though there has been significant progress towards this goal, all previous works have so far neglected the fact that in many scenarios (like presidential elections) voting is actually a distributed procedure. In this paper, we consider a setting in which the voters are partitioned into disjoint districts and vote locally therein to elect local winning alternatives using a voting rule; the final outcome is then chosen from the set of these alternatives. We prove tight bounds on the distortion of well-known voting rules for such distributed elections both from a worst-case perspective as well as from a best-case one. Our results indicate that the partition of voters into districts leads to considerably higher distortion, a phenomenon which we also experimentally showcase using real-world data.

Keywords: Distributed voting · District-based elections · Distortion

1 Introduction

In a decision-making scenario, the task is to aggregate the opinions of a group of different people into a common decision. This process is often distributed, in the sense that smaller groups first reach an agreement, and then the final outcome is determined based on the options proposed by each such group. This can be due to scalability issues (e.g., it is hard to coordinate a decision between a very large number of participants), due to different roles of the groups (e.g., when each group represents a country in the European Union), or simply due to established institutional procedures (e.g., electoral systems).

For example, in the US presidential elections, the voters in each of the 50 states cast their votes within their regional district, and each state declares a

* This work has been supported by the Swiss National Science Foundation under contract number 200021_165522 and by the European Research Council (ERC) under grant number 639945 (ACCORD).

winner; the final winner is taken as the one that wins a weighted plurality vote over the state winners, with the weight of each state being proportional to its size. Another example is the Eurovision Song Contest, where each participating country holds a local voting process (consisting of a committee vote and an Internet vote from the people of the country) and then assigns points to the 10 most popular options, on a 1-12 scale (with 11 and 9 omitted). The winner of the competition is the participant with the most total points.

The foundation of utilitarian economics, which originated near the end of the 18th century, revolves around the idea that the outcome of a decision making process should be one that maximizes the well-being of the society, which is typically captured by the notion of the *social welfare*. A fundamental question that has been studied extensively in the related literature is whether the rules that are being used for decision making actually achieve this goal, or to what extend they fail to do so. This motivates the following question: *What is the effect of distributed decision making on the social welfare?*

The importance of this investigation is highlighted by the example of the 2016 US presidential election [24]. While 48.2% of the US population (that participated in the election) viewed Hillary Clinton as the best candidate, Donald Trump won the election with only 46.1% of the popular vote. This was due to the district-based electoral system, and the outcome would have been different if there was a single pool of voters instead. A similar phenomenon occurred in the 2000 presidential election as well, when Al Gore won the popular vote, but George W. Bush was elected president.

1.1 Our Setting and Contribution

For concreteness, we use the terminology of voting as a proxy for any distributed decision-making scenario. A set of voters are called to vote on a set of alternatives through a district-based election. In other words, the set of voters is partitioned into *districts* and each district holds a local election, following some voting rule. The winners of the local elections are then aggregated into the single winner of the general election. Note that this setting models many scenarios of interest, such as those highlighted in the above discussion.

We are interested in the effect of the distributed nature of elections on the social welfare of the voters (the sum of their valuations for the chosen outcome). Typically, this effect is quantified by the notion of *distortion* [22], which is defined as the worst-case ratio between the maximum social welfare for any outcome and the social welfare for the outcome chosen through voting. Concretely, we are interested in bounding the distortion of voting rules for district-based elections.

We consider three cases when it comes to the district partition: (a) *symmetric districts*, in which every district has the same number of voters and contributes the same weight to the final outcome, (b) *unweighted districts*, in which the weight is still the same, but the sizes of the districts may vary, and (c) *unrestricted districts*, where the sizes and the weights of the districts are unconstrained. For each case, we show upper and lower bounds on the distortion of voting rules.

First, in Section 3, we consider general voting rules (which might have access to the numerical valuations of the voters) and provide distortion guarantees for any voting rule as a function of the worst-case distortion of the voting rule when applied to a single district. As a corollary, we obtain distortion bounds for *Range Voting*, the rule that outputs a welfare-maximizing alternative, and prove that it is optimal among all voting rules for the problem. Then, in Section 4, we consider ordinal rules and provide a general lower bound on the distortion of any such rule. For the widely-used Plurality voting rule, we provide *tight* distortion bounds, proving that it is asymptotically the best ordinal voting rule in terms of distortion. In Section 5, we provide experiments based on real data to evaluate the distortion on "average case" and "average worst case" district partitions. Finally, in Section 6, we explore whether *districting* (i.e., manually partitioning the voters into districts in the best-way possible) can allow to recover the winner of Plurality or Range Voting in the election without districts. We conclude with possible avenues for future work in Section 7. Due to space constraints, most proofs are omitted; see the full version of the paper [12].

1.2 Related Work

The distortion framework was first proposed by Procaccia and Rosenschein [22] and subsequently it was adopted by a series of papers; for instance, see [1, 2, 4, 5, 7, 8, 13]. The original idea of the distortion measure was to quantify the loss in performance due to the lack of *information*, meaning how well an ordinal voting rule (i.e., one that has access only to the preference orderings induced by the numerical values of the voters) can approximate the cardinal objective. In our paper, the distortion will be attributed to two factors: *always* the fact that the election is being done in districts, and *possibly* also the fact that the voting rules employed are ordinal. Our setting follows closely that of Boutilier et al. [7] and Caragiannis et al. [8], with the novelty of introducing district-based elections and measuring their distortion. The worst-case distortion bounds of voting rules in the absence of districts can be found in the aforementioned papers.

The ill effects of district-based elections have been highlighted in a series of related articles, mainly revolving around the issue of *gerrymandering* [23], i.e., the systematic manipulation of the geographical boundaries of an electoral constituency in favor of a particular political party. The effects of gerrymandering have been studied in the related literature before [6, 9, 19], but never in relation to the induced distortion of the elections. While our district partitions are not necessarily geographically-based, our worst-case bounds capture the potential effects of gerrymandering on the deterioration of the social welfare. Other works on district-based elections and distributed decision-making include [3, 10].

Related to our results in Section 6 is the paper by Lewenberg et al. [20], which explores the effects of districting with respect to the winner of Plurality, when ballot boxes are placed on the real plane, and voters are partitioned into districts based on their nearest ballot box. The extra constraints imposed by the geological nature of the districts in their setting leads to an NP-hardness result for the districting problem, whereas for our unconstrained districts, making the

Plurality winner the winner of the general election is always possible in polynomial time. In contrast, the problem becomes NP-hard when we are interested in the winner of Range Voting instead of Plurality.

2 Preliminaries

A *general election* $\mathcal{E}$ is defined as a tuple $(\mathcal{M}, \mathcal{N}, \mathcal{D}, \mathbf{w}, \mathbf{v}, \mathbf{f})$, where

- $\mathcal{M}$ is a set of m *alternatives*;
- $\mathcal{N}$ is a set of n *voters*;
- $\mathcal{D}$ is a set of $k \geq 2$ districts, with district $d \in \mathcal{D}$ containing n_d voters such that $\sum_{d \in \mathcal{D}} n_d = n$ (i.e., the districts define a partition of the set of voters);
- $\mathbf{w} = (w_d)_{d \in \mathcal{D}}$ is a *weight-vector* consisting of a weight $w_d \in \mathbb{R}_{>0}$ for each district $d \in \mathcal{D}$;
- $\mathbf{v} = (\mathbf{v}_1, \ldots, \mathbf{v}_n)$ is a *valuation profile* for the n voters, where $\mathbf{v}_i = (v_{ij})_{j \in \mathcal{M}}$ contains the *valuation* of voter i for all alternatives, and $\mathcal{V}^n$ is the set of all such valuation profiles;
- $\mathbf{f} = (f_d)_{d \in \mathcal{D}}$ is a set of *voting rules* (one for each district), where $f_d : \mathcal{V}^{n_d} \to \mathcal{M}$ is a map of valuation profiles with n_d voters to alternatives.

For each voter $i \in \mathcal{N}$, we denote by $d(i)$ the district she belongs to. For each district $d \in \mathcal{D}$, a *local* or *district* election between its members takes place, and the winner of this election is the alternative $j_d = f_d((\mathbf{v}_i)_{i:d(i)=d})$ that gets elected according to f_d. The outcome of the general election $\mathcal{E}$ is an alternative

$$j(\mathcal{E}) \in \arg\max_{j \in M} \sum_{d \in \mathcal{D}} w_d \cdot \mathbf{1}\left\{j_d = j\right\},$$

where $\mathbf{1}\{X\}$ is equal to 1 if the event X is true, and 0 otherwise. In simple words, the winner $j(\mathcal{E})$ of the general election is the alternative with the highest weighted approval score, breaking ties arbitrarily. For example, when all weights are 1, $j(\mathcal{E})$ is the alternative that wins the most local elections.

Following the standard convention, we adopt the unit-sum representation of valuations, according to which $\sum_{j \in \mathcal{M}} v_{ij} = 1$ for every voter $i \in \mathcal{N}$. For a given valuation profile $\mathbf{v}$, the *social welfare* of alternative $j \in \mathcal{M}$ is defined as the total value the agents have for her:

$$\mathrm{SW}(j|\mathbf{v}) = \sum_{i \in \mathcal{N}} v_{ij}.$$

Throughout the paper, we assume that the same voting rule is applied in every local election (possibly for a different number of voters though, depending on how the districts are defined); we denote this voting rule by f and also let $f(\mathbf{v})$ be the alternative that is chosen by f when the voters have the valuation profile $\mathbf{v}$.

The distortion of a voting rule f *in a local election* with η voters is defined as the worst-case ratio, over all possible valuation profiles of the voters participating in that election, between the maximum social welfare of any alternative and the social welfare of the alternative chosen by the voting rule:

$$\texttt{dist}(f) = \sup_{\mathbf{v} \in \mathcal{V}^\eta} \frac{\max_{j \in \mathcal{M}} \mathrm{SW}(j|\mathbf{v})}{\mathrm{SW}(f(\mathbf{v})|\mathbf{v})}.$$

The distortion of a voting rule f *in a general election* is defined as the worst-case ratio, over all possible general elections $\mathcal{E}$ that use f as the voting rule within the districts, between the maximum social welfare of any alternative and the social welfare of the alternative chosen by the general election:

$$\texttt{gdist}(f) = \sup_{\mathcal{E}: f \in \mathcal{E}} \frac{\max_{j \in \mathcal{M}} \mathrm{SW}(j|\mathbf{v})}{\mathrm{SW}(j(\mathcal{E})|\mathbf{v})}.$$

Again, in simple words, the distortion of a voting rule f is the worst-case over all the possible valuations that voters can have and over all possible ways of partitioning these voters into districts. When $k = 1$, we recover the standard definition of the distortion.

Next, we define some standard properties of voting rules.

Definition 1 (Properties of voting rules). A voting rule f is

- *ordinal*, if the outcome only depends on the preference orderings induced by the valuations and not the actual numerical values themselves. Formally, given a valuation profile $\mathbf{v}$, let $\Pi_{\mathbf{v}}$ be the ordinal preference profile formed by the values of the agents for the alternatives (assuming some fixed tie-breaking rule). A voting rule is ordinal if for any two valuation profiles $\mathbf{v}$ and $\mathbf{v}'$ such that $\Pi_{\mathbf{v}} = \Pi_{\mathbf{v}'}$, it holds that $f(\mathbf{v}) = f(\mathbf{v}')$.
- *unanimous*, if whenever all agents agree on an alternative, that alternative gets elected. Formally, whenever there exists an alternative $a \in \mathcal{M}$ for whom $v_{ia} \geq v_{ij}$ for all voters $i \in \mathcal{N}$ and all alternatives $j \in \mathcal{M}$, then $f(\mathbf{v}) = a$.
- *(strictly) Pareto efficient*, if whenever all agents agree that an alternative a is better than b, then b cannot be elected. Formally, if $v_{ia} > v_{ib}$ for all $i \in \mathcal{N}$, then $f(\mathbf{v}) \neq b$.[4]

Remark. It is not hard to see that we can assume that the best voting rule in terms of distortion is Pareto efficient, without loss of generality. Indeed, for any voting rule f that is not Pareto efficient, we can construct the following Pareto efficient rule f': for every input on which f outputs a Pareto efficient alternative, f' outputs the same alternative; for every input on which f outputs an alternative that is not Pareto efficient, f' outputs a maximal Pareto improvement, that is, a Pareto efficient alternative which all voters (weakly) prefer more than the alternative chosen by f. Clearly, f' is Pareto efficient and achieves a social

[4] Pareto efficiency usually requires that there is no other alternative who all voters *weakly prefer* and who one voter *strictly prefers*. We use the strict definition in our proofs, as it is also without loss of generality with respect to distortion.

welfare at least as high as f. Note also that Pareto efficiency implies unanimity. In our proofs, we will use both of these properties without loss of generality. Finally, most of the voting rules that are being employed in practice are ordinal, with the notable exception of *Range Voting*, which is the voting rule that outputs the alternative that maximizes the social welfare.

We consider the following three basic cases for the general elections, depending on the size and the weight of the districts:

- *Symmetric Elections*: all districts consist of n/k voters and have the same weight, i.e., $n_d = n/k$ and $w_d = 1$ for each $d \in \mathcal{D}$.
- *Unweighted Elections*: all districts have the same weight, but not necessarily the same number of voters, i.e., $w_d = 1$ for each $d \in \mathcal{D}$.
- *Unrestricted Elections*: there are no restrictions on the sizes and weights of the districts.

Of course, the class of symmetric elections is a subclass of that of unweighted elections which in turn is a subclass of the class of unrestricted elections.

3 The Effect of Districts for General Voting Rules

Our aim in this section is to showcase the immediate effect of using districts to distributively aggregate votes. To this end, we present tight bounds on the distortion of all voting rules in a general election. We will first state a general theorem relating the distortion $\texttt{gdist}(f)$ of any general election that uses a voting rule f for the local elections, with the distortion $\texttt{dist}(f)$ of the voting rule.

Theorem 1. *Let f be a voting rule with $\texttt{dist}(f) = \gamma$. Then, the distortion $\texttt{gdist}(f)$ of f in the general election is at most*

(i) $\gamma + \frac{\gamma m k}{2}$ *for symmetric elections;*
(ii) $\gamma + \frac{\gamma m}{2}\left(\frac{n+\max_{d\in\mathcal{D}} n_d}{\min_{d\in\mathcal{D}} n_d} - 1\right)$ *for unweighted elections;*
(iii) $\gamma + \gamma m\left(\frac{n}{\min_{d\in\mathcal{D}} n_d} - 1\right)$ *for unrestricted elections.*

We now turn to concrete voting rules and consider perhaps the most natural such rule: *Range Voting* (RV).

Definition 2 (Range Voting (RV)). *Given a valuation profile $\mathbf{v} = (\mathbf{v}_1, ..., \mathbf{v}_\eta)$ with η voters, Range Voting elects the alternative that maximizes the social welfare of the voters.*

Note that the rule is both unanimous and Pareto efficient. Immediately from the definition of the rule and Theorem 1, we have the following corollary.

Corollary 1. *The distortion* $\texttt{gdist}(\mathrm{RV})$ *of* RV *in the general election is at most*

(i) $1 + \frac{mk}{2}$ *for symmetric elections;*
(ii) $1 + \frac{m}{2}\left(\frac{n+\max_{d\in\mathcal{D}} n_d}{\min_{d\in\mathcal{D}} n_d} - 1\right)$ *for unweighted elections;*
(iii) $1 + m\left(\frac{n}{\min_{d\in\mathcal{D}} n_d} - 1\right)$ *for unrestricted elections.*

We continue by presenting matching lower bounds on the distortion of any voting rule in a general election. The high-level idea in the proof of the following theorem is that the election winner is chosen arbitrarily among the alternatives with equal weight, which might lead to the cardinal information within the districts to be lost.

Theorem 2. *The distortion of all voting rules in a general election is at least*

(i) $1 + \frac{mk}{2}$ *for symmetric elections;*
(ii) $1 + \frac{m}{2}\left(\frac{n+\max_{d\in\mathcal{D}} n_d}{\min_{d\in\mathcal{D}} n_d} - 1\right)$ *for unweighted elections;*
(iii) $1 + m\left(\frac{n}{\min_{d\in\mathcal{D}} n_d} - 1\right)$ *for unrestricted elections.*

4 Ordinal Voting Rules and Plurality

Although Range Voting is quite natural, its documented drawback is that it requires a very detailed informational structure from the voters, making the elicitation process rather complicated. For this reason, most voting rules that have been applied in practice are ordinal (see Definition 1), as such rules present the voters with the much less demanding task of reporting a preference ordering over the alternatives, rather than actual numerical values.

Thus, a very meaningful question, from a practical point of view, is "What is the distortion of ordinal voting rules?" The most widely used such rule is *Plurality Voting*. Besides its simplicity, the importance of this voting rule also comes from the fact that it is used extensively in practice. For instance, it is used in presidential elections in a number of countries like the USA and the UK.

Definition 3 (Plurality Voting (PV)). *Given a valuation profile* $\mathbf{v}$ *and its induced ordinal preference profile* $\Pi_{\mathbf{v}}$*,* PV *elects the alternative with the most first position appearances in* $\Pi_{\mathbf{v}}$*, breaking ties arbitrarily.*

It is known that the distortion $\texttt{dist}(\text{PV})$ of Plurality Voting is $O(m^2)$ [8]. Therefore, if we plug-in this number to our general bound in Theorem 1, we obtain corresponding upper bounds for PV. However, in the following we obtain much better bounds, taking advantage of the structure of the mechanism; these bounds are actually tight.

Theorem 3. *The distortion* $\texttt{gdist}(\text{PV})$ *of* PV *is* exactly

(i) $1 + \frac{3m^2k}{4}$ *for symmetric elections;*
(ii) $1 + \frac{m^2}{4}\left(\frac{3n+\max_{d\in\mathcal{D}} n_d}{\min_{d\in\mathcal{D}} n_d} - 1\right)$ *for unweighted elections;*

(iii) $1 + m^2 \left(\frac{n}{\min_{d \in \mathcal{D}} n_d} - \frac{1}{2} \right)$, *for unrestricted elections.*

Proof. We prove only the upper bounds for the first two parts here; the upper bound for the third part as well as the matching lower bounds can be found in the full version.

Consider a general unweighted election $\mathcal{E}$ with a set $\mathcal{M}$ of m alternatives, a set $\mathcal{N}$ of n voters, a set $\mathcal{D}$ of k districts such that each district d consists of n_d voters and has weight $w_d = 1$. Let $\mathbf{v}$ be the valuation profile consisting of the valuations of all voters for all alternatives, which induces the ordinal preference profile $\Pi_{\mathbf{v}}$. To simplify our discussion, let $\mathcal{N}_d(j)$ be the set of voters in district d that rank alternative j in the first position, and also set $|\mathcal{N}_d(j)| = n_d(j)$.

Let $a = j(\mathcal{E})$ be the winner of the election and denote by $A \subseteq \mathcal{D}$ the set of districts in which a wins according to PV. Then, we have that

$$\mathrm{SW}(a|\mathbf{v}) = \sum_{i \in \mathcal{N}} v_{ia} \geq \sum_{i:d(i) \in A} v_{ia}. \tag{1}$$

Since a has the plurality of votes in each district $d \in A$, we have that $n_d(a) \geq n_d(j)$ for every $j \in \mathcal{M}$, and by the fact that $\sum_{j \in \mathcal{M}} n_d(j) = n_d$, we obtain that $n_d(a) \geq \frac{n_d}{m}$. Similarly, for each agent $i \in \mathcal{N}_d(a)$ we have that $v_{ia} \geq v_{ij}$ for every $j \in \mathcal{M}$, and by the unit-sum assumption, we obtain that $v_{ia} \geq \frac{1}{m}$. We also have that $\sum_{d \in A} n_d \geq |A| \cdot \min_{d \in \mathcal{D}} n_d$. Hence,

$$\begin{aligned} \sum_{i:d(i) \in A} v_{ia} &\geq \sum_{d \in A} \sum_{i \in \mathcal{N}_d(a)} v_{ia} \geq \frac{1}{m} \cdot \sum_{d \in A} n_d(a) \\ &\geq \frac{1}{m^2} \sum_{d \in A} n_d \geq \frac{1}{m^2} \cdot |A| \cdot \min_{d \in \mathcal{D}} n_d. \end{aligned} \tag{2}$$

Let b the optimal alternative, and denote by $B \subset \mathcal{D}$ the set of districts in which b is the winner. We split the social welfare of b into three parts:

$$\mathrm{SW}(b|\mathbf{v}) = \sum_{i:d(i) \in A} v_{ib} + \sum_{i:d(i) \in B} v_{ib} + \sum_{i:d(i) \notin A \cup B} v_{ib}. \tag{3}$$

We will now bound each term individually. First consider a district $d \in A$. Then, the welfare of the agents in d for b can be written as

$$\sum_{i:d(i)=d} v_{ib} = \sum_{i \in \mathcal{N}_d(a)} v_{ib} + \sum_{i \in \mathcal{N}_d(b)} v_{ib} + \sum_{i \notin \mathcal{N}_d(a) \cup \mathcal{N}_d(b)} v_{ib}.$$

Since a is the favourite alternative of every agent $i \in \mathcal{N}_d(a)$, $v_{ib} \leq v_{ia}$. By definition, the value of every agent $i \in \mathcal{N}_d(b)$ for b is at most 1. The value of every agent $i \notin \mathcal{N}_d(a) \cup \mathcal{N}_d(b)$ for b can be at most $1/2$ since otherwise b would definitely be the favourite alternative of such an agent. Combining these observations, we get

$$\begin{aligned}
\sum_{i:d(i)=d} v_{ib} &\leq \sum_{i\in\mathcal{N}_d(a)} v_{ia} + n_d(b) + \frac{1}{2}\sum_{j\neq a,b} n_d(j) \\
&\leq \sum_{i:d(i)=d} v_{ia} + \frac{1}{2}n_d(b) + \frac{1}{2}\sum_{j\neq a} n_d(j) \\
&\leq \sum_{i:d(i)=d} v_{ia} + \frac{1}{2}n_d(a) + \frac{1}{2}\Big(n_d - n_d(a)\Big) \\
&= \sum_{i:d(i)=d} v_{ia} + \frac{1}{2}n_d,
\end{aligned}$$

where the second inequality follows by considering the value of all agent in d for alternative a, while the third inequality follows by the fact that a wins b by plurality. By summing over all districts in A, we can bound the first term of (3) as follows:

$$\sum_{i:d(i)\in A} v_{ib} \leq \sum_{i:d(i)\in A} v_{ia} + \frac{1}{2}\sum_{d\in A} n_d. \tag{4}$$

For the second term of (3), by definition we have that the value of each agent in the districts of B for alternative b can be at most 1, and therefore

$$\sum_{i:d(i)\in B} v_{ib} \leq \sum_{d\in B} n_d.$$

For the third term of (3), observe that the total value of the agents in a district $d \notin A\cup B$ for b must be at most $\frac{3}{4}n_d$; otherwise b would necessarily be ranked first in strictly more than half of the agents' preferences and therefore win in the district. Hence,

$$\sum_{i:d(i)\notin A\cup B} v_{ib} \leq \frac{3}{4}\sum_{d\notin A\cup B} n_d.$$

By substituting the bounds for the three terms of (3), as well as by taking into account the facts that $|B| \leq |A|$ and $|A| \geq 1$, we can finally upper-bound the social welfare of b as follows:

$$\begin{aligned}
\mathrm{SW}(b|\mathbf{v}) &\leq \sum_{i:d(i)\in A} v_{ia} + \frac{1}{2}\sum_{d\in A} n_d + \sum_{d\in B} n_d + \frac{3}{4}\sum_{d\notin A\cup B} n_d \\
&= \sum_{i:d(i)\in A} v_{ia} + \frac{1}{4}\left(3n + \sum_{d\in B} n_d - \sum_{d\in A} n_d\right) \\
&\leq \sum_{i:d(i)\in A} v_{ia} + \frac{1}{4}\left(3n + |B|\cdot\max_{d\in\mathcal{D}} n_d - |A|\cdot\min_{d\in\mathcal{D}} n_d\right) \\
&\leq \sum_{i:d(i)\in A} v_{ia} + \frac{1}{4}\cdot|A|\cdot\left(3n + \max_{d\in\mathcal{D}} n_d - \min_{d\in\mathcal{D}} n_d\right) \qquad (5)
\end{aligned}$$

By (1), (2) and (5), we can upper-bound the distortion of PV as follows:

$$\begin{aligned}\texttt{gdist}(\mathrm{PV}) &= \frac{\mathrm{SW}(b|\mathbf{v})}{\mathrm{SW}(a|\mathbf{v})}\\ &\leq \frac{\sum_{i:d(i)\in A} v_{ia} + \frac{1}{4}\cdot|A|\cdot(3n+\max_{d\in\mathcal{D}} n_d - \min_{d\in\mathcal{D}} n_d)}{\sum_{i:d(i)\in A} v_{ia}}\\ &\leq 1+\frac{m^2}{4}\left(\frac{3n+\max_{d\in\mathcal{D}} n_d}{\min_{d\in\mathcal{D}} n_d}-1\right).\end{aligned}$$

This completed the proof of part (ii). For part (i), we get the desired bound of $1+\frac{3m^2k}{4}$ by simply setting $\min_{d\in\mathcal{D}} n_d = \max_{d\in\mathcal{D}} n_d = n/k$. □

Our next theorem shows that PV is asymptotically the best possible voting rule among all deterministic ordinal voting rules.

Theorem 4. *The distortion* $\texttt{gdist}(f)$ *of any ordinal Pareto efficient voting rule* f *is*

(i) $\Omega(m^2k)$*, for symmetric elections;*
(ii) $\Omega\left(m^2\frac{n+\max_{d\in\mathcal{D}} n_d}{\min_{d\in\mathcal{D}} n_d}\right)$*, for unweighted elections;*
(iii) $\Omega\left(\frac{m^2n}{\min_{d\in\mathcal{D}} n_d}\right)$*, for unrestricted elections.*

5 Experiments

Thus far, we have studied the worst-case effect of the partition of voters into districts on the distortion of voting rules. In this section, we further showcase this phenomenon experimentally by using real-world utility profiles that are drawn from the Jester dataset [16], which consists of ratings of 100 different jokes in the interval $[-10, 10]$ by approximately 70,000 users; this dataset has been used in a plethora of previous papers, including the seminal work of Boutilier et al. [7]. Following their methodology, we build instances with a set of alternatives that consists of the eight most-rated jokes. For various values of k, we execute 1000 independent simulations as follows: we select a random set of 100 users among the ones that evaluated all eight alternatives, rescale their ratings so that they are non-negative and satisfy the unit-sum assumption, and then divide them into k districts.

For the partition into districts, we consider both *random* partitions as well as *bad* partitions in terms of distortion. For the construction of the latter, for each instance consisting of a specific value of k and a set of voters, we create 100 random partitions of the voters into k districts, simulate the general election (based on the voting rules we consider) and then keep the partition with maximum distortion.

We compare the average distortion of four rules: Range Voting, Plurality, Borda, and Harmonic. Borda and Harmonic are two well-known positional scoring rules defined by the scoring vectors $(m-1, m-2, ..., 0)$ and $(1, 1/2, ..., 1/m)$,

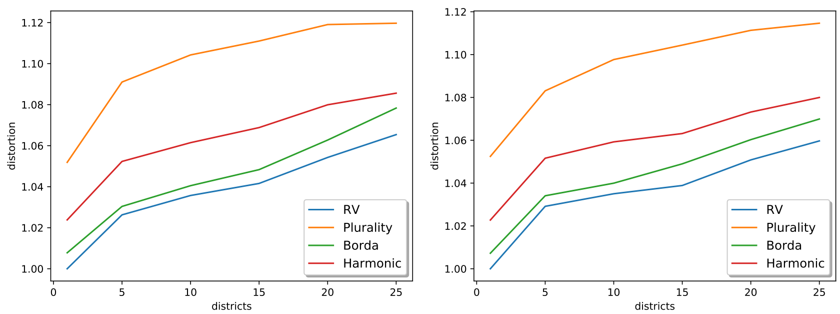

(a) Unweighted (b) Weighted

Fig. 1: Average distortion from 1000 simulations as a function of the number of districts k with random partitions of voters into districts.

Table 1: Average distortion from 1000 simulations with bad partitions of voters into districts.

	unweighted					weighted				
k	1	2	3	4	5	1	2	3	4	5
Range Voting	1	4.82	4.51	4.50	4.60	1	4.46	4.96	5.14	5.14
Plurality	1.05	5.03	4.66	4.71	4.81	1.05	4.77	5.29	5.47	5.49
Borda	1.01	4.83	4.47	4.50	4.61	1.01	4.51	4.98	5.16	5.18
Harmonic	1.02	4.97	4.60	4.62	4.72	1.02	4.64	5.16	5.35	5.36

respectively. According to these rules, each voter assigns points to the alternatives based on the positions she ranks them, and the alternative with the most points is the winner; Plurality can also be defined similarly by the scoring vector $(1, 0, ..., 0)$.

Figure 1 depicts the results of our simulations for unweighted and weighted districts when the partition into districts is random and $k \in \{1, 5, 10, 15, 20, 25\}$; for weighted districts, the weights are drawn uniformly at random from a given interval. As one can observe, the behaviour of the four voting rules is very similar in both cases, and it is evident that as the number of districts increases, the distortion increases as well. For instance, the distortion of Plurality increased by 3.71% for $k = 5$ compared to $k = 1$ (i.e., when there are no districts) and by 6.44% for $k = 25$; these values are similar for the other rules as well, although a bit lower. Table 1 contains the results of our simulations for unweighted and weighted districts when the partition into districts is bad (in terms of distortion) and $k \in \{1, 2, 3, 4, 5\}$. As in the case of random districts, we can again observe that the distortion increases as k increases, but now the difference between the cases with districts ($k \geq 2$) and without districts ($k = 1$) is more clear; the distortion is almost five times higher.

6 Best-case Partitions via Districting

In this section we turn out attention to a somewhat different setting. We assume that the k districts are not a priori defined, and instead we are free to decide the partition of the voters into the districts so as to minimize their effect on the distortion of the underlying voting rule; we refer to the process of partitioning the voters into k districts as *k-districting*. We consider symmetric districts, and start our analysis with the question of whether it is possible to define the districts so that the optimal alternative (i.e., the one that maximizes the social welfare of the voters) wins the general election when RV is used as the voting rule. Unfortunately, as we show with our next theorem, this is not always possible.

Theorem 5. *For every $k \geq 2$, there exists an instance such that no symmetric k-districting allows the optimal alternative to win the general election when RV is the voting rule.*

Proof. Consider a general election with $n+1$ alternatives $\mathcal{M} = \{a_1, ..., a_n, b\}$ and let k be such that n/k is an integer for simplicity; then, each district must consist of exactly n/k voters. Let $\varepsilon \in \left(0, \frac{1}{2(n+1)}\right)$ and let $\mathbf{v}$ be the valuation profile according to which voter i has value $\frac{n}{n+k} + \varepsilon$ for alternative a_i and value $\frac{k}{n+k} - \varepsilon$ for alternative b; her value for the remaining alternatives is zero.

Since $\text{SW}(a_i|\mathbf{v}) = \frac{n}{n+k} + \varepsilon$ for every $i \in [n]$ and $\text{SW}(b|\mathbf{v}) = \frac{nk}{n+k} - n\varepsilon$, alternative b is clearly the optimal alternative. However, observe that all possible sets of n/k voters that can be included together in a district cannot make b the winner of the district when the voting rule is RV. Indeed, the welfare of such a set of voters for b is only $\frac{n}{n+k} - \frac{n\varepsilon}{k}$, while their welfare for the alternatives they rank first is $\frac{n}{n+k}$. Therefore, there is no symmetric k-districting that can make b the winner of the general election with RV. $\square$

In fact, the instance used in the above proof indicates that even the best-case distortion of RV may be at least k. We continue the bad news by showing that the problem of deciding whether it is possible to define the districts such that the optimal alternative wins the general election with RV is NP-hard for $k = 2$.

Theorem 6. *Deciding whether there is a symmetric 2-districting such that the optimal alternative is the winner of the general election with RV is NP-hard.*

In contrast to the above result for the optimal alternative and RV, we next show that we can always find a symmetric k-districting so that the PV winner without districts can be made the winner of the general election when PV is used as the voting rule within the districts. Since the voting rule is PV, we assume that the only knowledge which we can leverage in order to define the districts is about the favourite alternatives of the voters (i.e., for each voter, we know the alternative she approves).

Theorem 7. *For any $k \geq 2$, there always exists a symmetric k-districting that allows the winner of PV without districts to win the general election with k districts, and this districting can be computed in polynomial time.*

We conclude this section by showing that the above result for PV is essentially tight. This follows by the existence of instances where any partition of the voters into any number of districts yields distortion for the general election with PV that is asymptotically equal to the distortion of PV without districts.

Theorem 8. *There exist instances where any symmetric districting yields distortion* $\texttt{gdist}(\text{PV}) = \Omega(m^2)$.

7 Conclusion and Possible Extensions

In this paper, we have initiated the study of the distortion of distributed voting. We showcased the effect of districting on the social welfare both theoretically from a worst- and a best-case perspective, as well as experimentally using real-world data. Even though we have painted an almost complete picture, our work reveals many interesting avenues for future research.

In terms of our results, possibly the most obvious open question is whether we can strengthen the weak intractability result of Theorem 6 using a reduction from a strongly NP-hard problem, and also extend it to $k \geq 2$. Moving away from the unconstrained normalized setting that we considered here, it would be very interesting to analyze the effect of districts in the case of *metric preferences* [1], a setting that has received considerable attention in the recent related literature on the distortion of voting rules without districts [2, 11, 14, 15, 17, 21]. Other important extensions include settings in which the partitioning of voters into districts is further constrained by natural factors such as geological locations [20] or connectivity in social networks [18].

References

1. Anshelevich, E., Bhardwaj, O., Elkind, E., Postl, J., Skowron, P.: Approximating optimal social choice under metric preferences. Artificial Intelligence **264**, 27–51 (2018)
2. Anshelevich, E., Postl, J.: Randomized social choice functions under metric preferences. Journal of Artificial Intelligence Research **58**, 797–827 (2017)
3. Bachrach, Y., Lev, O., Lewenberg, Y., Zick, Y.: Misrepresentation in district voting. In: Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI). pp. 81–87 (2016)
4. Benade, G., Nath, S., Procaccia, A.D., Shah, N.: Preference elicitation for participatory budgeting. In: Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI). pp. 376–382 (2017)
5. Bhaskar, U., Dani, V., Ghosh, A.: Truthful and near-optimal mechanisms for welfare maximization in multi-winner elections. In: Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI). pp. 925–932 (2018)
6. Borodin, A., Lev, O., Shah, N., Strangway, T.: 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). pp. 98–104 (2018)

7. Boutilier, C., Caragiannis, I., Haber, S., Lu, T., Procaccia, A.D., Sheffet, O.: Optimal social choice functions: A utilitarian view. Artificial Intelligence **227**, 190–213 (2015)
8. Caragiannis, I., Nath, S., Procaccia, A.D., Shah, N.: Subset selection via implicit utilitarian voting. Journal of Artificial Intelligence Research **58**, 123–152 (2017)
9. Cohen-Zemach, A., Lewenberg, Y., Rosenschein, J.S.: Gerrymandering over graphs. In: Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS). pp. 274–282 (2018)
10. Erdélyi, G., Hemaspaandra, E., Hemaspaandra, L.A.: More natural models of electoral control by partition. In: Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT). pp. 396–413 (2015)
11. Feldman, M., Fiat, A., Golomb, I.: On voting and facility location. In: Proceedings of the 17th ACM Conference on Economics and Computation (EC). pp. 269–286 (2016)
12. Filos-Ratsikas, A., Micha, E., Voudouris, A.A.: The distortion of distributed voting. CoRR **abs/1905.01882** (2019)
13. Filos-Ratsikas, A., Miltersen, P.B.: Truthful approximations to range voting. In: Proceedings of the 10th Conference on Web and Internet Economics (WINE). pp. 175–188 (2014)
14. Goel, A., Hulett, R., Krishnaswamy, A.K.: Relating metric distortion and fairness of social choice rules. In: Proceedings of the 13th Workshop on the Economics of Networks, Systems and Computation (NetEcon). p. 4:1 (2018)
15. Goel, A., Krishnaswamy, A.K., Munagala, K.: Metric distortion of social choice rules: Lower bounds and fairness properties. In: Proceedings of the 18th ACM Conference on Economics and Computation (EC). pp. 287–304 (2017)
16. Goldberg, K., Roeder, T., Gupta, D., Perkins, C.: Eigentaste: A constant time collaborative filtering algorithm. Information Retrieval **4**, 133–151 (2001)
17. Gross, S., Anshelevich, E., Xia, L.: Vote until two of you agree: Mechanisms with small distortion and sample complexity. In: Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI). pp. 544–550 (2017)
18. Lesser, O., Naamani-Dery, L., Kalech, M., Elovici, Y.: Group decision support for leisure activities using voting and social networks. Group Decision and Negotiation **26**(3), 473–494 (2017)
19. Lev, O., Lewenberg, Y.: "reverse gerrymandering": a decentralized model for multi-group decision making. In: Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI) (2019)
20. Lewenberg, Y., Lev, O., Rosenschein, J.S.: 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). pp. 624–632 (2017)
21. Pierczynski, G., Skowron, P.: Approval-based elections and distortion of voting rules. CoRR **abs/1901.06709** (2019)
22. Procaccia, A.D., Rosenschein, J.S.: The distortion of cardinal preferences in voting. In: Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA). pp. 317–331 (2006)
23. Schuck, P.H.: The thickest thicket: Partisan gerrymandering and judicial regulation of politics. Columbia Law Review **87**(7), 1325–1384 (1987)
24. Wikipedia: 2016 United States presidential election (2016), https://en.wikipedia.org/wiki/2016_United_States_presidential_election