

Approximating Optimal Social Choice under Metric Preferences*

Elliot Anshelevich
Rensselaer Polytechnic Institute
Troy, NY, USA
eanshel@cs.rpi.edu

Onkar Bhardwaj
Rensselaer Polytechnic Institute
Troy, NY, USA
bhardo@rpi.edu

Edith Elkind
University of Oxford
Oxford, UK
elkind@cs.ox.ac.uk

John Postl
Rensselaer Polytechnic Institute
Troy, NY, USA
postlj@rpi.edu

Piotr Skowron
University of Warsaw
Warsaw, Poland
p.skowron@mimuw.edu.pl

July 17, 2018

Abstract

We consider voting under metric preferences: both voters and alternatives are associated with points in a metric space, and each voter prefers alternatives that are closer to her to ones that are further away. In this setting, it is often desirable to select an alternative that minimizes the sum of distances to the voters, i.e., the utilitarian social cost, or other similar measures of social cost. However, common voting rules operate on voters’ preference rankings and therefore may be unable to identify an optimal alternative. A relevant measure of the quality of a voting rule is then its *distortion*, defined as the worst-case ratio between the performance of an alternative selected by the rule and that of an optimal alternative. Thus, distortion measures how good a voting rule is at approximating an alternative with minimum social cost, while using only ordinal preference information. The underlying costs can be arbitrary, implicit, and unknown; our only assumption is that they form a metric space. The goal of our paper is to quantify the distortion of well-known voting rules. We first establish a lower bound on the distortion of any deterministic voting rule. We then show that the distortion of positional scoring rules cannot be bounded by a constant, and for several popular rules in this family distortion is linear in the number of alternatives. On the other hand, for Copeland and similar rules the distortion is bounded by a factor of 5. These results hold both for the sum of voters’ cost and the median voter cost. For Single Transferable Vote (STV), we obtain an upper bound of $O(\ln m)$ with respect to the sum of voters’ costs, where m is the number of alternatives, as well as a lower bound of $\Omega(\sqrt{\ln m})$; thus, STV is a reasonable, though not a perfect rule from this perspective. Our results for median voter cost extend to more general objective functions.

1 Introduction

Voting rules aggregate preferences of multiple voters (also referred to as agents) over a set of available candidates (also referred to as alternatives), enabling the voters to choose an option that reflects their collective opinion. Often, voters’ preferences are determined by the candidates’ positions on several issues, such as the level of taxation or military spending. In this case, each voter and each candidate can be identified with a point in the *issue space*, and voters tend to prefer candidates who are close to them to the ones that are further away. This setting can be formally modeled by embedding the input election into a *metric space*, i.e., a set of points S endowed with a distance measure d : voters’ preferences are consistent with this embedding if voter i prefers candidate X to candidate Y whenever $d(i, X) < d(i, Y)$. The spatial model of preferences has received a considerable amount of attention in the social

*This paper combines and extends results presented at the 29th Conference on Artificial Intelligence, AAAI-2015 (paper titled “Approximating Optimal Social Choice under Metric Preferences”) and at the 31st Conference on Artificial Intelligence, AAAI-2017 (paper titled “Social Choice Under Metric Preferences: Scoring Rules and STV”).

choice literature due to its intuitive appeal [Davis and Hinich, 1966; Plott, 1967; Enelow and Hinich, 1984, 1990; McKelvey and Ordeshook, 1990; Merrill and Grofman, 1999; Schofield, 2007] and has also been considered by AI researchers [Elkind *et al.*, 2017].

When preferences are driven by distances, it is natural to measure the *quality* of an alternative X by computing the sum of distances from X to the voters; in some cases, other objective functions, such as the median distance or the maximum distance, may be more appropriate. We may then want to select an alternative that minimizes this objective function. Of course, this task is not difficult if we are given access to voters’ and alternatives’ locations. However, typically voters are unable to precisely pinpoint their position with respect to each issue, and even the issue space itself may not be known to the designer of the voting mechanism (e.g., see the discussion by Boutilier *et al.* [2015]). Thus, it is more realistic to expect the voters to simply provide their rankings of alternatives, which are determined by the underlying metric space: each voter ranks the alternatives by the distance from her. We then find ourselves in a setting that is well-studied in the social choice literature: we are given a collection of ranked ballots (linear orders) and have to select a single alternative based on these ballots. We can therefore use one of the many voting rules that have been proposed for this scenario (see, e.g., the survey by Zwicker [2015]).

Of course, we cannot expect a voting rule that operates on ranked ballots to always identify an alternative that minimizes the total distance to the voters. Therefore, one can think of a voting rule as an approximation algorithm that attempts to choose the best possible alternative (one that minimizes the social cost, such as, e.g., the total distance to the voters) given access to limited information (to ordinal preferences instead of distances). This perspective on voting with ranked ballots was proposed by Procaccia and Rosenschein [2006] in the context of a more general setting where voters’ utilities for alternatives may be arbitrary (i.e., not determined by distances). They introduced the term *distortion* to refer to the quality of approximation provided by a voting rule; we will continue to use this term, although we will formally define it in terms of voters’ costs rather than their utilities. Intuitively, the distortion of a voting rule is the worst-case ratio of the social cost (such as the sum of distances to voters) of an alternative selected by the voting rule over the cost of an optimal alternative.

Distortion can be viewed as a quantitative measure allowing for a normative comparison of various voting rules; in settings where preferences can be assumed to be driven by distances, low distortion is a highly desirable feature. A comparison of voting rules based on their distortion can be seen as an instantiation of the utilitarian approach, which is common in welfare economics and algorithmic mechanism design. This approach, which has recently received renewed attention in the study of social choice [Procaccia and Rosenschein, 2006; Caragiannis and Procaccia, 2011; Boutilier *et al.*, 2015; Branzei *et al.*, 2013], assumes that every agent has (possibly latent or implicit) utility or cost values over the alternatives. As argued by Boutilier *et al.* [2015], although not all social choice problems are amenable to the utilitarian approach (especially the ones where it is unnatural to assume that voters’ utilities or costs can be compared), there are many real-life settings that fit the utilitarian view. For example, in recommender systems and many similar domains from mechanism design and e-commerce, the computational agents typically assign real-valued utilities to alternatives rather than have ordinal preferences over the set of alternatives. Thus, our work complements the vast literature on normative comparison of voting rules, by applying the utilitarian approach to one of the oldest and most fundamental questions in social choice theory.

In this work, we study the quality of outcomes chosen by common voting rules, as measured by their distortion. The objective function that we are primarily interested in is the sum of distances to the voters, though we also consider the median distance and other similar objective functions. Our results show that, while some commonly used rules have high distortion, there are important voting rules for which distortion is bounded by a small constant or grows slowly with the number of alternatives.

Our Contributions

We bound the worst-case distortion of many well-known voting rules. In other words, we show how closely these rules, which only have access to ordinal preferences, approximate an optimal alternative with respect to the metric costs.¹ We consider two general objective functions to quantify the quality of alternatives, and, for most of the rules we consider, we give distortion bounds with respect to both of these functions. The first is the sum objective function, which defines the social cost of an alternative as the sum of all agent costs for that alternative. This function is very natural, and is the most common measure of social cost. Our other objective function defines the quality of an

¹Throughout this paper, we assume that agents submit their true ordinal preferences. We leave questions about non-truthful agents as future work; see also the work of Feldman *et al.* [2016].

	Sum		Median	
	Upper bound	Lower bound	Upper bound	Lower bound
Plurality	$2m - 1$	$2m - 1$	∞	∞
Borda	$2m - 1$	$2m - 1$	∞	∞
k -Approval, $k > 1$	∞	∞	∞	∞
Veto	∞	∞	∞	∞
Harmonic rule	$O\left(\frac{m}{\sqrt{\ln m}}\right)$	$\Omega\left(\frac{m}{\ln m}\right)$	∞	∞
STV	$O(\ln m)$	$\Omega(\sqrt{\ln m})$?	$\Omega(\sqrt{\ln m})$
Uncovered set	5	5	5	5
Copeland	5	5	5	5
Scoring rules	∞	$1 + 2\sqrt{\ln m - 1}$	∞	∞
Deterministic voting rules	∞	3	∞	5

Table 1: The worst-case distortion of various voting rules for both the sum and the median objective functions. The numbers in the columns labeled ‘Lower bound’ correspond to example instances where the voting rule cannot achieve a better distortion with respect to the relevant objective function.

alternative as the median of agent costs for that alternative: focusing on the median agent is quite common as well, as this reduces the impact of outliers, i.e., agents with very high or very low costs.

Most of our results are summarized in Table 1. First, we provide a lower bound on the performance of all deterministic voting rules that operate on ranked ballots. We show that no such rule can have worst-case distortion that is better than 3 (for the sum objective), or better than 5 (for the median objective). With these lower bounds established, we can nevertheless ask: do there exist voting rules that meet this lower bound? Are there rules that obtain the minimum possible distortion?

We begin with the bad news: for common positional scoring rules, such as Plurality and Borda, the worst-case distortion can be as high as $2m - 1$, where m is the number of alternatives. For k -Approval with $k > 1$ and for Veto the worst-case distortion can be unbounded. Moreover, there is no family of positional scoring rules whose worst-case distortion is bounded by a constant independent of m : for each value of m , the distortion of every positional scoring rule for m alternatives is at least $1 + 2\sqrt{\ln m - 1}$. Interestingly, neither Plurality nor Borda achieve the best possible distortion in the class of positional scoring rules: we show that Harmonic rule of Boutilier *et al.* [2015] has asymptotically better distortion than either Plurality or Borda. However, the distortion of Harmonic rule is still almost linear.

In contrast, some of the Condorcet-consistent rules have considerably lower distortion. In particular, we prove that the distortion of every voting rule that always outputs a subset of the *uncovered set* [Moulin, 1986] does not exceed 5; this upper bound holds both for the sum objective and for the median objective, though different techniques are required to prove it for the two cases. There are several voting rules that have this property, including the well-known Copeland rule. This means that, although these rules know nothing about the metric costs other than the ordinal preferences induced by them, and cannot possibly find the true optimal alternative, they nevertheless *always* select an alternative whose quality is only a factor of 5 away from optimal! Moreover, our lower bounds show that no deterministic voting rule can do better than Copeland for the median objective, and no deterministic voting rule can do much better than Copeland for the sum objective. To complement these results, we provide a lower bound on the distortion of the Copeland rule with respect to the sum objective, showing that it cannot achieve the best possible distortion. We also consider the Ranked Pairs rule and identify a range of scenarios where the distortion of this rule with respect to the sum objective does not exceed 3, thereby matching our lower bound.

Another rule that we consider is Single Transferable Vote (STV), which is one of the very few non-trivial voting rules used in real-life elections; for instance, STV is used to elect members of governing bodies (at local or national level) in several countries, including Australia, New Zealand, United Kingdom and United States. It turns out that for the sum objective the distortion of STV is upper-bounded by $O(\ln m)$, i.e., STV performs much better than Plurality or Borda, and, in particular, it offers acceptable distortion when the number of alternatives is not too large. On the other hand, we show that the distortion of STV is lower-bounded by $\Omega(\sqrt{\ln m})$, i.e., STV does not perform quite as

well as the Copeland rule.

In addition to the results in Table 1, we also analyze more general objective functions. Specifically, instead of the median objective, which sets the quality of an alternative W to be the cost to the median voter, we consider more general *percentile* objectives, where the quality of an alternative W is set to be the cost of the voter at the x -th percentile. We show how the distortion of various rules changes with x , and establish that Copeland remains the rule with the best possible distortion for most values of x .

Related Work

The social choice theory offers several strategies to identify a voting rule that is suitable for a particular application. The *normative approach* proceeds by formulating axiomatic properties that a rule should satisfy, and then characterizing the set of rules that have these properties or proving that no such rule exists. Early influential papers in this stream of research include those of May [1952], Arrow [1963], Gibbard [1973], Satterthwaite [1975] and Young [1975]; see the survey by Zwicker [2015]. Another approach is based on *maximum likelihood estimation*: it is assumed that there is an objectively optimal alternative or a ranking of alternatives, but voters make errors of judgement, so their rankings are noisy estimates of the truth. The goal of a voting rule is then to identify the most likely ‘ground truth’ for a given model of noise, and different noise models correspond to different voting rules. This approach dates back to Condorcet [1785], and can be used to justify a variety of voting rules, such as the Kemeny rule [Young, 1988] or positional scoring rules [Conitzer and Sandholm, 2005]; see the survey by Elkind and Slinko [2016]. A related idea is that of *distance rationalizability*, which identifies each voting rule with a consensus (a class of elections with a clear winner) and a distance measure on rankings [Campbell and Nitzan, 1986; Elkind *et al.*, 2015; Elkind and Slinko, 2016].

In contrast, in this work we adopt a utilitarian approach to voting. The utilitarian perspective on social choice has its advocates in welfare economics [Roemer, 1998; Ng, 1997], and has recently received attention from the artificial intelligence and algorithmic game theory communities [Procaccia and Rosenschein, 2006; Caragiannis and Procaccia, 2011; Boutilier *et al.*, 2015; Filos-Ratsikas and Miltersen, 2013; Branzei *et al.*, 2013; Chakrabarty and Swamy, 2014]. While the assumption that voters’ costs can be compared is not universally applicable [Harsanyi, 1976], it is nevertheless reasonable in many applications of interest; see the work of Boutilier *et al.* [2015] for further discussion. The notion of *distortion* as a measure of performance of voting rules in utilitarian settings was introduced in an influential early paper of Procaccia and Rosenschein [2006], and later used by Boutilier *et al.* [2015]. However, in these works it is assumed that each voter may assign arbitrary utilities to the alternatives, so, in particular, the utilities are not determined by metric considerations. In this more general setting, all deterministic voting rules have unbounded distortion. In contrast, in the metric scenario considered in this paper, we obtain meaningful bounds on the distortion of many voting rules.

Voting with spatial preferences has a strong tradition in social choice and political science [Davis and Hinich, 1966; Plott, 1967; Enelow and Hinich, 1984, 1990; McKelvey and Ordeshook, 1990; Merrill and Grofman, 1999; Schofield, 2007]. While some of this work assumes preferences to be one-dimensional, we consider metric spaces of arbitrary dimension. Also, we note that our use of distances is very different from that in the distance rationalizability framework: in the latter, distances are defined over preference profiles or, sometimes, preference orders, while in our model we consider distances between voters and alternatives.

Finally, the concept of distortion is related to many other notions of approximation, as it compares the optimal solution with the solution obtained given only limited information. In this regard, it is similar, for example, to the concept of competitive ratio in online algorithms, which is a measure of how an algorithm performs with limited information (not knowing the future), compared to how an all-knowing algorithm would perform [Borodin and El-Yaniv, 1998; Oren and Lucier, 2014].

After a preliminary version of our paper was published, several authors have made progress towards understanding distortion in voting with metric costs. Specifically, Anshelevich and Postl [2017] extend the analysis of distortion to randomized voting rules. Feldman *et al.* [2016] study distortion under the additional constraint of strategyproofness; in particular, for the case where the underlying metric space is a line they propose a universally truthful randomized mechanism whose distortion does not exceed 2. Goel *et al.* [2017] provide bounds on the distortion of randomized tournament rules, and disprove a conjecture made in the conference version of our paper (see Section 3.2).

2 Preliminaries

Given a positive integer k , let $[k] = \{1, \dots, k\}$.

Social Choice with Ordinal Preferences

Let $N = \{1, 2, \dots, n\}$ be the set of *agents* (also called the *voters*), and let $\mathcal{A} = \{A_1, A_2, \dots, A_m\}$ be the set of *alternatives* (also called the *candidates*). Let \mathcal{S} be the set of all total orders on the set of alternatives \mathcal{A} . We will typically use i, j to refer to agents and W, X, Y, Z to refer to alternatives. Every agent $i \in N$ has a *preference ranking* $\sigma_i \in \mathcal{S}$. By $X \succ_i Y$ we will mean that X is preferred over Y by voter i ; we write $X \succeq_i Y$ if $X \succ_i Y$ or $X = Y$. We denote the position of alternative X in the preference ranking of voter i by $\sigma_i(X)$: if $\sigma_i(X) = 1$ then X is i 's most preferred alternative. We call the vector $\sigma = (\sigma_1, \dots, \sigma_n) \in \mathcal{S}^n$ a *preference profile*. The following notation will be used throughout our proofs:

$$\begin{aligned} XY &= \{i \in N : X \succ_i Y\} \\ XYZ &= \{i \in N : X \succ_i Y \succ_i Z\} \end{aligned}$$

We say that alternative X *Pareto-dominates* alternative Y in σ if $|XY| = n$. We say that X *defeats* Y if $|XY| > \frac{n}{2}$; X *weakly defeats* Y if $|XY| \geq \frac{n}{2}$. An alternative X is a *Condorcet winner* (respectively, a *weak Condorcet winner*) if it defeats (respectively, weakly defeats) all other alternatives. The (weak) *majority graph* of σ is a directed graph with node set \mathcal{A} where there is an arc from X to Y if and only if X (weakly) defeats Y .

Once we are given a preference profile, we want to aggregate the preferences of the agents in order to select a subset of winning alternatives (usually, we would like to identify a single winner, but we allow for ties). A *voting rule* $f: \mathcal{S}^n \rightarrow \mathcal{A}$ is a mapping that, given a preference profile over \mathcal{A} , outputs a non-empty subset \mathcal{W} of \mathcal{A} . The alternatives in \mathcal{W} are called the *winners under f* . In this paper, we consider the following well-known voting rules.

Positional scoring rules. Fix a set of alternatives \mathcal{A} , $|\mathcal{A}| = m$. A *positional scoring rule* for \mathcal{A} is determined by a scoring vector $\vec{s} = (s_1, s_2, \dots, s_m)$ with $s_1 \geq s_2 \geq \dots \geq s_m$, $s_1 > s_m$, $s_i \in \mathbb{Q}$ for $i \in [m]$. If an agent ranks an alternative in position ℓ , then this alternative receives s_ℓ points from that agent. The total score $s(X, \sigma)$ of an alternative X with respect to a preference profile σ is the total number of points that X receives from all agents. The positional scoring rule based on vector \vec{s} is $f_{\vec{s}}(\sigma) = \arg \max_{X \in \mathcal{A}} s(X, \sigma)$; that is, it selects the alternatives with the highest total score. Many well-known voting rules can be thought of as positional scoring rules, for example:

- **Plurality:** $\vec{s} = (1, 0, \dots, 0)$,
- **Veto:** $\vec{s} = (1, 1, \dots, 1, 0)$,
- **Borda:** $\vec{s} = (m-1, m-2, \dots, 1, 0)$,
- **Harmonic rule:** $\vec{s} = (1, 1/2, 1/3, \dots, 1/m)$,
- **k -Approval** ($1 \leq k < m$): $\vec{s} = (\underbrace{1, 1, \dots, 1}_k, 0, \dots, 0)$.

Copeland. The Copeland score of an alternative X is the number of alternatives that X defeats. The Copeland rule outputs all alternatives with the maximum Copeland score.

Ranked Pairs. Given a profile σ over a set of alternatives \mathcal{A} , construct a graph G with node set \mathcal{A} in the following manner. Sort the pairs $(X, Y) \in \mathcal{A} \times \mathcal{A}$ in non-increasing order of weights $|XY|$ and iterate over them. For each directed edge (X, Y) in the sorted list, add it to G if it will not create a directed cycle, and do nothing otherwise. The winners are the source nodes of the resulting directed acyclic graph.

Single Transferable Vote (STV). STV is an iterative rule that works as follows. In each round one of the alternatives with the lowest Plurality score (the one that is ranked first by the fewest voters) is removed from the set of alternatives and from the rankings of the voters; the Plurality scores are then recalculated. After $m-1$ rounds, only one alternative survives; this alternative is declared to be a winner.

Note that we described Ranked Pairs and STV as non-deterministic voting rules: there are multiple ways to order the pairs (X, Y) in non-increasing order of weights, and in any given round, there may be multiple alternatives

with the lowest Plurality score. There are several ways to make them deterministic; in this work, we will consider the so-called *parallel-universe model* [Conitzer *et al.*, 2009]. In this model, an alternative is said to be a Ranked Pairs winner if it is a source node in the resulting acyclic graph for some way of ordering the node pairs in non-increasing order of weights; similarly, an alternative is an STV winner if it survives after $m - 1$ rounds for *some* sequence of choices at each elimination step.

The execution of STV is illustrated by the following example.

Example 1 Consider the following preference profile:

$$\begin{array}{ll} 1: A_1 \succ_1 A_2 \succ_1 A_3 \succ_1 A_4, & 2: A_1 \succ_2 A_4 \succ_2 A_3 \succ_2 A_2, \\ 3: A_2 \succ_3 A_4 \succ_3 A_1 \succ_3 A_3, & 4: A_2 \succ_4 A_4 \succ_4 A_3 \succ_4 A_1, \\ 5: A_3 \succ_5 A_2 \succ_5 A_4 \succ_5 A_1, & 6: A_3 \succ_6 A_2 \succ_6 A_1 \succ_6 A_4, \\ 7: A_1 \succ_7 A_2 \succ_7 A_4 \succ_7 A_3, & 8: A_4 \succ_8 A_1 \succ_8 A_2 \succ_8 A_3. \end{array}$$

In the first round A_4 is eliminated. In the second round, after A_4 is removed, the Plurality scores of the alternatives A_1 , A_2 , and A_3 , are equal to 4, 2, and 2, respectively. If we choose to eliminate A_3 , then in the next round the scores of A_1 and A_2 are equal to 4, so either of them can be eliminated. Alternatively, we can eliminate A_2 in the second round; in the next round, the score of A_1 is 5, while the score of A_3 is 3, so A_3 is eliminated. Thus, the set of STV winners is $\{A_1, A_2\}$.

Cardinal Metric Costs

In our work we study settings where the ordinal preferences σ are induced by underlying cardinal agent costs. Formally, we assume that there exists a pseudo-metric $d : (N \cup \mathcal{A})^2 \rightarrow \mathbb{R}_{\geq 0}$ on the set of agents and alternatives (we only require d to be a pseudo-metric rather than a metric, since we allow agents to be at distance 0 from each other). Here $d(i, X)$ is the cost incurred by agent i when alternative X is selected as the winner; these costs can be arbitrary, but are assumed to obey the triangle inequality. The metric costs d naturally give rise to a preference profile: we say that a profile σ over a set of alternatives \mathcal{A} is *consistent* with d if for all $i \in N$ and for all $X, Y \in \mathcal{A}$ it holds that $d(i, X) < d(i, Y)$ implies $X \succ_i Y$. In other words, if the cost of X is less than the cost of Y for an agent, then the agent should prefer X over Y . Let $p(d)$ denote the set of preference profiles consistent with d ($p(d)$ need not be a singleton, as agent costs may have ties); consequently, given a profile σ , we denote by $p^{-1}(\sigma)$ the set of all pseudo-metrics d such that $\sigma \in p(d)$. All of our results continue to hold if we restrict ourselves to the set of pseudo-metrics d such that $d(X, Y) > 0$ for all distinct $X, Y \in \mathcal{A}$.

In general, by assuming that the preference rankings of the agents are generated in a specific way, we may rule out some preference profiles. For instance, if we further restrict our model by assuming that agents and alternatives are associated with points on the real line, the metric is the usual Euclidean metric, and no agent is equidistant from two alternatives, the resulting profile is guaranteed to be single-peaked and single-crossing and hence its majority graph has no directed cycles [Black, 1948; Gans and Smart, 1996; Saporiti, 2009]. However, having arbitrary metric costs in our model does not restrict the set of possible profiles σ in any way: metrics are general enough that any preference profile in \mathcal{S}^n can be induced, even if we require that all pairwise distances between alternatives are strictly positive; in fact, Bogomolnaia and Laslier [2007] establish that this is the case even if our metric space is \mathbb{R}^t with the usual Euclidean metric, for $t \geq \min\{n, m - 1\}$.

Social Cost and Distortion

We measure the quality of each alternative using the costs incurred by all agents when this alternative is chosen. We use two different notions of social cost. First, we study the sum objective function, defined as $\text{SC}_{\Sigma}(X, d) = \sum_{i \in N} d(i, X)$; this is the most common notion of social cost. We also study the median objective function, $\text{SC}_{\text{med}}(X, d)$: this function sorts the agent costs $d(i, X)$ in non-decreasing order, and outputs the $\lfloor \frac{n}{2} + 1 \rfloor$ -th cost. As described in the introduction, we can view voting rules in our setting as attempting to find an optimal alternative (one that minimizes the social cost), despite only having access to the ordinal preference profile σ , instead of the full underlying costs d . The following proposition establishes that this is impossible to do: we can only eliminate an alternative from consideration of being optimal if it is Pareto-dominated by another alternative.

Proposition 2 *Given a preference profile σ , let \mathcal{A}' denote the set of alternatives that are not Pareto-dominated in σ . Then for every alternative $W \in \mathcal{A}'$ there exists a metric $d \in p^{-1}(\sigma)$ such that W is the unique optimal alternative with respect to the social cost function $\text{SC}_\Sigma(\cdot, d)$.*

Proof. For every alternative $W \in \mathcal{A}'$ we will define a metric $d \in p^{-1}(\sigma)$ such that W is optimal for σ with respect to d . For each agent i , set $d(i, X) = 1$ for all X such that $X \succeq_i W$ and $d(i, X) = 2$ for all X such that $W \succ_i X$. For every pair of distinct alternatives X, Y , set $d(X, Y) = 2$. For every pair of distinct agents i, j , set $d(i, j) = 2$. It is easy to verify that d is a metric; in particular, the triangle inequality holds since d takes values in $\{0, 1, 2\}$.

Consider any alternative $X \neq W$. We observe that

$$\begin{aligned} \text{SC}_\Sigma(X, d) - \text{SC}_\Sigma(W, d) &= \sum_{i \in N} d(i, X) - \sum_{i \in N} d(i, W) \\ &= \sum_{i \in XW} (d(i, X) - d(i, W)) + \sum_{i \in WX} (d(i, X) - d(i, W)) \\ &= 0 + 1 \cdot |WX| \geq 1. \end{aligned}$$

Thus, we conclude that W is the unique optimal alternative for σ with respect to d . ■

A simple corollary of Proposition 2 is that if there is no alternative ranked first by all agents then we cannot determine an optimal alternative with respect to the sum objective based on ordinal preferences only: every alternative that is ranked first by at least one agent may turn out to be the unique optimal alternative.

Since all voting rules defined earlier in this section operate on preference profiles, they cannot be expected to find an optimal alternative. Nevertheless, we would like to understand how good are the alternatives selected by these voting rules in terms of their social cost. We focus on the worst-case performance. To formalize this question, we use the notion of *distortion* [Procaccia and Rosenschein, 2006; Boutilier *et al.*, 2015].

Definition 3 *The distortion of a voting rule f on a profile σ with respect to a social cost function $\text{SC} \in \{\text{SC}_\Sigma, \text{SC}_{\text{med}}\}$ is given by*

$$\text{dist}_{\text{SC}}(f, \sigma) = \sup_{d \in p^{-1}(\sigma)} \frac{\max_{W \in f(\sigma)} \text{SC}(W, d)}{\min_{X \in \mathcal{A}} \text{SC}(X, d)}.$$

To deal with the possibility that $\min_{X \in \mathcal{A}} \text{SC}(X, d)$ may be equal to 0 for some metric d , we use the following convention: if for some $d \in p^{-1}(\sigma)$ we have $\text{SC}(X, d) = 0$ for some $X \in \mathcal{A}$, we set

$$\frac{\max_{W \in f(\sigma)} \text{SC}(W, d)}{\min_{X \in \mathcal{A}} \text{SC}(X, d)} = \begin{cases} 1 & \text{if } \text{SC}(W, d) = 0 \text{ for all } W \in f(\sigma) \\ +\infty & \text{otherwise} \end{cases}$$

For readability, we write dist_Σ and dist_{med} in place of $\text{dist}_{\text{SC}_\Sigma}$ and $\text{dist}_{\text{SC}_{\text{med}}}$, respectively.

In other words, the distortion of a voting rule f on a profile σ is the worst-case ratio between the social cost of an alternative in $f(\sigma)$ and the social cost of a true optimal alternative. The worst case is taken over all possible outputs of $f(\sigma)$ and over all metrics d that may have induced σ , since the voting rule does not and cannot know which of these metrics is the true one.

3 Distortion of Total Agent Cost

In this section, we study the sum objective function, which measures the quality of an alternative as the total agent cost when this alternative is chosen. We prove tight upper bounds for distortion of several well-known voting rules. The main result of this section is that the Copeland rule exhibits a distortion of at most 5; this guarantee is independent of the number of agents or alternatives, and the underlying metric space is allowed to be completely arbitrary (and unknown). Further, we show that scoring rules do not behave well with respect to distortion; in particular, the distortion of Plurality and Borda scales linearly with the number of alternatives. This conclusion is important because scoring rules are used in a variety of applications due to their simplicity and intuitive appeal. Finally, we establish that the

distortion of STV is upper-bounded by $O(\ln m)$ and lower-bounded by $\Omega(\sqrt{\ln m})$. Thus, while STV does not perform quite as well as the Copeland rule, its distortion is much better than that of Plurality or Borda. This result is particularly important because STV is a rule that is often used in practice.

Note that it is possible to directly compute an alternative that minimizes the worst-case distortion, by using a linear program (the construction is similar to the one in the proof of Theorem 3.4 in the work of Boutilier *et al.* [2015]). However, there are two reasons why we believe that our results are important, and why we focus on rules such as Plurality and Copeland. *First*, one of the main questions that we are interested in is how well a voting rule can perform without knowledge of the underlying metric, as compared to the “true” optimum, which could only be computed by an omniscient rule. To prove an upper bound on the distortion of the “LP rule” directly seems difficult; instead it makes sense to focus on bounding the performance of simpler rules, and then use the fact that the LP rule performs even better. In fact, one way to interpret our results is: “The worst-case distortion of the LP rule (and thus of the best outcome that can be achieved without knowing the metric) is at least 3 and at most 5; if the circumference of the weak majority graph is small, then it is exactly 3.” *Second*, and perhaps more importantly, although the worst-case distortion of the LP rule is at least as good as that on any other deterministic voting rule, this rule is computationally intense (although still computable in polynomial time) and it may be difficult to convince people to use it, as compared to such “natural” and well-established voting rules as Plurality, Borda, STV, and Copeland. Furthermore, nothing is known about the axiomatic properties that the LP rule satisfies, while the voting rules we consider are quite well understood. Thus, our work focuses on distortion of very computationally efficient, simple-to-describe, and commonly used voting rules which satisfy many desirable axioms.

Before proceeding with showing upper bounds on the distortion of specific voting rules, we ask the question: how well can *any* voting rule perform? The following simple result tells us that we cannot hope to approximate the optimal alternative within a factor better than 3.

Theorem 4 *No (deterministic) voting rule has worst-case distortion less than 3 for the sum objective.*

Proof. Suppose there are only two alternatives X and W . Half of the agents prefer X over W , and the other half prefer W over X . Suppose without loss of generality that the given voting rule picks W as one of the winners. Consider then the following metric. All $n/2$ agents who prefer X to W are located exactly at X , i.e., $d(i, X) = 0$ and $d(i, W) = 2$. All $n/2$ agents who prefer W are approximately halfway between X and W , i.e., $d(i, X) = 1 + \epsilon$ and $d(i, W) = 1 - \epsilon$ for some small $\epsilon > 0$. Then $\text{SC}_\Sigma(X, d) = \sum_{i \in N} d(i, X) = (1 + \epsilon) \cdot \frac{n}{2}$ and $\text{SC}_\Sigma(W, d) = \sum_{i \in N} d(i, W) = 2 \cdot \frac{n}{2} + (1 - \epsilon) \cdot \frac{n}{2}$. Thus, the distortion approaches 3 as $\epsilon \rightarrow 0$. ■

For two alternatives, this bound is tight: it can be shown that for $|\mathcal{A}| = 2$ the voting rule that outputs all alternatives ranked first by a weak majority of agents has a distortion of 3 (see Corollary 7).

Let us now prove two lemmas that will be useful in our further analysis. Our first lemma provides bounds on the costs of agents incurred by alternatives; these bounds will be used repeatedly throughout all of our proofs. Specifically, we prove two lower bounds on the cost of an alternative, which will be used to lower-bound the cost of an optimal alternative. We also establish an upper bound, which will be used to bound the cost of a winning alternative W for agents who prefer an optimal alternative X over W ; for agent i who prefers W to X , we can simply use $d(i, W) \leq d(i, X)$. Since we assume that agent costs are metric, all of these bounds crucially rely on the triangle inequality.

Lemma 5 *Let W, X, Y, Z be alternatives. Then the following bounds hold:*

$$\forall i \in WX, \quad d(i, X) \geq \frac{1}{2} \cdot d(X, W) \quad (1)$$

$$\forall i \in WY, \quad d(i, X) \geq \frac{1}{2} \cdot (d(X, W) - d(X, Y)) \quad (2)$$

$$\forall i, \quad d(i, W) \leq d(i, X) + \min_{W \succeq_i Z} d(X, Z) \quad (3)$$

Proof. (1) Fix an agent $i \in WX$. Since $W \succ_i X$, we have $d(i, W) \leq d(i, X)$. Combining this with the triangle inequality, we have $d(X, W) \leq d(i, X) + d(i, W) \leq 2 \cdot d(i, X)$.

(2) Fix an agent $i \in WY$. By the triangle inequality we have $d(i, X) \geq d(X, W) - d(i, W)$. Since i prefers W over Y , we have $d(i, W) \leq d(i, Y)$ and hence $d(i, X) \geq d(X, W) - d(i, Y)$. By adding this inequality to $d(i, X) \geq d(i, Y) - d(X, Y)$, we obtain the desired result.

- (3) Fix an agent i and alternatives W and X . Let Z be some alternative such that $W \succeq_i Z$ (note that there always exists such an alternative since $W \succeq_i W$). By the triangle inequality, $d(i, W) \leq d(i, Z) \leq d(i, X) + d(X, Z)$. Since this holds for every Z with $W \succeq_i Z$, we conclude that (3) holds. ■

Our next lemma parameterizes the distortion by the number of agents who prefer a (winning) alternative W over an (optimal) alternative X .

Lemma 6 *For every preference profile σ , every pseudo-metric $d \in p^{-1}(\sigma)$, and every pair of alternatives W, X we have*

$$\frac{\text{SC}_\Sigma(W, d)}{\text{SC}_\Sigma(X, d)} \leq 1 + \frac{2(n - |WX|)}{|WX|} = \frac{2n}{|WX|} - 1.$$

Proof. First, we want to upper-bound the agent cost incurred by alternative W . We do this by dividing the agents into two groups, WX and XW . For an agent $i \in WX$, we know that $d(i, W) \leq d(i, X)$. For an agent $i \in XW$ we use the triangle inequality to obtain $d(i, W) \leq d(i, X) + d(X, W)$. We obtain

$$\begin{aligned} \frac{\text{SC}_\Sigma(W, d)}{\text{SC}_\Sigma(X, d)} &= \frac{\sum_{i \in N} d(i, W)}{\sum_{i \in N} d(i, X)} = \frac{\sum_{i \in WX} d(i, W) + \sum_{i \in XW} d(i, W)}{\sum_{i \in N} d(i, X)} \\ &\leq \frac{\sum_{i \in WX} d(i, X) + \sum_{i \in XW} (d(i, X) + d(X, W))}{\sum_{i \in N} d(i, X)} = 1 + \frac{\sum_{i \in XW} d(X, W)}{\sum_{i \in N} d(i, X)} \\ &= 1 + \frac{|XW| \cdot d(X, W)}{\sum_{i \in N} d(i, X)} = 1 + \frac{(n - |WX|) \cdot d(X, W)}{\sum_{i \in N} d(i, X)}. \end{aligned} \quad (4)$$

It remains to lower-bound the cost of X . It follows from Lemma 5 that $d(i, X) \geq \frac{1}{2} \cdot d(X, W)$ for each agent i with $W \succ_i X$. Thus we have $\sum_{i \in N} d(i, X) \geq \frac{1}{2}|WX| \cdot d(X, W)$. Combining this inequality with equation (4) gives us the desired result. ■

An important corollary of Lemma 6 is that by selecting weak Condorcet winners we can ensure low distortion with respect to the sum objective; in particular, this means that for $m = 2$ the distortion of the weak majority rule (the rule that outputs all alternatives ranked first by at least half of the voters) is at most 3.

Corollary 7 *Consider a profile σ that admits a weak Condorcet winner W , and a pseudo-metric $d \in p^{-1}(\sigma)$. Let X be an optimal alternative for σ with respect to d . Then $\text{SC}_\Sigma(W, d) \leq 3 \cdot \text{SC}_\Sigma(X, d)$.*

Proof. Let n be the number of agents. We have $|WX| \geq n/2$ and hence $\frac{\text{SC}_\Sigma(W, d)}{\text{SC}_\Sigma(X, d)} \leq \frac{2n}{|WX|} - 1 = 3$. ■

We are now ready to show upper bounds on maximum distortion for several important voting rules.

3.1 Distortion of Scoring Rules

We start our discussion by considering positional scoring rules. Unfortunately, it turns out that the distortion of such rules cannot be bounded by a constant: for each constant c there is a value of m_0 such that the worst-case distortion of all positional scoring rules for m_0 or more alternatives is at least c . Moreover, for many well-known families of positional scoring rules the distortion grows linearly with the number of alternatives, or cannot be bounded at all.

Example 8 *Consider a positional scoring rule given by a score vector $\vec{s} = (s_1, \dots, s_m)$ with $s_1 = s_2$. Suppose that for $j = 1, \dots, m$ alternative A_i is located at the point $x = i$ on the real axis \mathbb{R} , all agents are located at $x = 1$, and the distance is the usual Euclidean distance in \mathbb{R} . Then all n agents rank the alternatives as $A_1 \succ A_2 \succ \dots \succ A_m$, and alternatives A_1 and A_2 are among the election winners, as they both achieve the highest possible score $ns_1 = ns_2$. However, the total distance from agents to A_1 is 0, whereas the total distance from agents to A_2 is n , so the distortion of our rule is $+\infty$.*

Example 8 establishes that the distortion of Veto and k -Approval with $k > 1$ is unbounded. We will now prove a lower bound on the distortion of positional scoring rules with $s_1 > s_2$.

Theorem 9 *Let $f_{\vec{s}}$ be a positional scoring rule given by a score vector $\vec{s} = (s_1, \dots, s_m)$ with $s_1 > s_2$. Then there is a profile σ over m alternatives such that $\text{dist}_{\Sigma}(f_{\vec{s}}, \sigma) \geq 1 + 2\sqrt{\ln m - 1}$. Moreover, if $\vec{s} = (m-1, \dots, 1, 0)$ (i.e., if $f_{\vec{s}}$ is the Borda rule) or if $\vec{s} = (1, 0, \dots, 0)$ (i.e., if $f_{\vec{s}}$ is Plurality) there is a profile σ over m alternatives such that $\text{dist}_{\Sigma}(f_{\vec{s}}, \sigma) \geq 2m - 1$. Furthermore, if $\vec{s} = (1, 1/2, \dots, 1/m)$ (i.e., if $f_{\vec{s}}$ is the Harmonic rule) there is a profile σ over m alternatives such that $\text{dist}_{\Sigma}(f_{\vec{s}}, \sigma) = \Omega\left(\frac{m}{\ln m}\right)$.*

Proof. Fix the number of alternatives m and the scoring vector $\vec{s} = (s_1, \dots, s_m)$ with $s_1 > s_2$. As $s_1 > s_2$ implies $s_1 \neq s_m$, it can be assumed without loss of generality that $s_1 = 1, s_m = 0$. For every $z = 1, \dots, m$, let

$$F_z = s_1 + \dots + s_z, \quad L_z = s_{m-z+1} + \dots + s_m,$$

and set

$$f_z = \frac{F_z}{z}, \quad \ell_z = \frac{L_z}{z}.$$

We will construct a profile σ over m alternatives where all alternatives and all agents are associated with points in \mathbb{R} ; the description of the instance depends on the parameters z, n_1 , and n_2 whose values will be chosen later. In our construction, we place several alternatives in the same position; towards the end of the proof, we will show how to fix our construction so that $d(X, Y) > 0$ whenever X and Y are distinct alternatives.

There are z alternatives, $0 < z < m$, at $x = 1$; we denote these alternatives by B_1, \dots, B_z and write $\mathcal{B} = \{B_1, \dots, B_z\}$. The remaining $m - z$ alternatives are at $x = -1$; we denote them by C_1, \dots, C_{m-z} , and write $\mathcal{C} = \{C_1, \dots, C_{m-z}\}$. There are $z \cdot n_1$ agents at $x = 0$ and $z \cdot n_2$ agents at $x = 1$.

We assume that all agents at $x = 0$ prefer alternatives in \mathcal{C} to alternatives in \mathcal{B} , all agents rank the alternatives in \mathcal{C} in the same way, and in aggregate the agents are indifferent among alternatives in \mathcal{B} . Specifically, for each $i = 1, \dots, z$ there are n_1 agents at $x = 0$ who rank the alternatives as

$$C_1 \succ \dots \succ C_{m-z} \succ B_i \succ \dots \succ B_z \succ B_1 \succ \dots \succ B_{i-1}$$

and n_2 agents at $x = 1$ who rank the alternatives as

$$B_i \succ \dots \succ B_z \succ B_1 \succ \dots \succ B_{i-1} \succ C_1 \succ \dots \succ C_{m-z}.$$

For this instance, the sum of distances from agents to any alternative in \mathcal{B} is zn_1 , while the sum of distances from agents to any alternative in \mathcal{C} is $zn_1 + 2zn_2$, so as long as z, n_1 and n_2 are all positive, it is optimal to select an arbitrary alternative in \mathcal{B} . However, the score of C_1 is $zn_1 + zn_2s_{z+1}$, whereas the score of any alternative in \mathcal{B} is

$$n_1(s_{m-z+1} + \dots + s_m) + n_2(s_1 + \dots + s_z) = n_1L_z + n_2F_z.$$

Thus, if

$$zn_1 + zn_2s_{z+1} > n_1L_z + n_2F_z, \tag{5}$$

our rule outputs C_1 and we have

$$\text{dist}_{\Sigma}(f_{\vec{s}}, \sigma) \geq \frac{zn_1 + 2zn_2}{zn_1} = 1 + 2\frac{n_2}{n_1}.$$

We will now show that we can choose $z \in \{1, \dots, m-1\}$ and positive integers n_2, n_1 so that we get a lower bound on $\frac{n_2}{n_1}$ and condition (5) is satisfied. It will be convenient to rewrite condition (5) as follows:

$$z + z \cdot s_{z+1} \cdot \frac{n_2}{n_1} > F_z \cdot \frac{n_2}{n_1} + L_z$$

or, equivalently,

$$\frac{n_2}{n_1} < \frac{z - L_z}{F_z - z \cdot s_{z+1}};$$

note that the denominator of the fraction in the right-hand side is not zero since we assume $s_1 = 1, s_2 < 1$. Let

$$R(z) = \frac{z - L_z}{F_z - z \cdot s_{z+1}},$$

and observe that the distortion of $f_{\vec{s}}$ can be lower-bounded by $1 + 2 \max_{z \in \{1, \dots, m-1\}} R(z)$. Indeed, for a fixed value of $z \in \{1, \dots, m-1\}$ the quantity $R(z)$ is a rational fraction²; let us denote this fraction by $\frac{\alpha}{\beta}$. We can set $n_2 = \alpha$, $n_1 = \beta$ to obtain a profile on which the distortion of our rule is at least $1 + 2R(z)$. It remains to show how to choose z so as to maximize $R(z)$ for a given score vector \vec{s} .

For the Borda rule, after normalizing the score vector, we have $\vec{s} = (1, \frac{m-2}{m-1}, \frac{m-3}{m-1}, \dots, 0)$, so we take $z = 1$ and obtain

$$R(1) = \frac{1 - L_1}{F_1 - s_2} = \frac{1}{1 - \frac{m-2}{m-1}} = m - 1,$$

which gives the lower bound of $2m - 1$ on the distortion of the Borda rule.

For Plurality, we take $z = m - 1$ and get

$$R(m-1) = \frac{m-1 - L_{m-1}}{F_{m-1} - s_m} = \frac{m-1-0}{1-0} = m-1,$$

which proves the lower bound of $2m - 1$ on the distortion of Plurality.

For the Harmonic rule we also take $z = m - 1$ and obtain

$$R(m-1) = \frac{m-1 - L_{m-1}}{F_{m-1} - s_m} \geq \frac{m-1 - \ln(m-1) - 1}{\ln m} = \Omega\left(\frac{m}{\ln m}\right).$$

To show a lower bound for the general case, it remains to argue that for every score vector \vec{s} we can choose $z \in \{1, \dots, m-1\}$ so that $R(z) \geq \sqrt{\ln m - 1}$. For readability, we prove this fact in a separate lemma.

Lemma 10 *There exists a $z \in \{1, \dots, m-1\}$ such that $R(z) \geq \sqrt{\ln m - 1}$.*

Proof. Let $\lambda = 1/\sqrt{\ln m - 1}$. We will consider two cases.

1. $s_2 > 1 - \lambda$. In this case we can set $z = 1$. Indeed, for $z = 1$ we have $z - L_z = 1$, $F_z - z \cdot s_{z+1} = s_1 - s_2 < \lambda$, so $R(1) \geq 1/\lambda$, which is what we need to prove.
2. $s_2 \leq 1 - \lambda$. Note that in this case we have $z - L_z \geq z\lambda$ for $z = 1, \dots, m-1$.

We will argue that there exists a $z \in \{1, \dots, m-1\}$ such that $f_z - s_{z+1} \leq \lambda^2$; we then obtain

$$R(z) \geq \frac{z\lambda}{z\lambda^2} = \frac{1}{\lambda}.$$

Indeed, suppose that for each $z \in \{1, \dots, m-1\}$ we have $f_z - s_{z+1} > \lambda^2$. We have

$$(z+1) \cdot f_{z+1} = F_{z+1} = F_z + s_{z+1} = z \cdot f_z + s_{z+1} < (z+1)f_z - \lambda^2;$$

dividing both sides by $z+1$ gives $f_{z+1} < f_z - \frac{\lambda^2}{z+1}$. Inductively, this implies

$$f_m < 1 - \lambda^2 \left(\frac{1}{2} + \dots + \frac{1}{m} \right).$$

But $\frac{1}{2} + \dots + \frac{1}{m} > \ln m - 1$, and hence

$$\lambda^2 \left(\frac{1}{2} + \dots + \frac{1}{m} \right) > 1,$$

whereas $f_m = \frac{1}{m}(s_1 + \dots + s_m)$ is necessarily positive, a contradiction. Thus, $f_z - s_{z+1} \leq \lambda^2$ for some $z \in \{1, \dots, m-1\}$, and the proof is complete.

²Note that we assume that the entries of the scoring vector are rational numbers.

As argued above, Lemma 10 immediately implies the desired lower bound. ■

We can modify our construction so as to induce the same profile σ using a distance d such that $d(X, Y) > 0$ for all distinct alternatives X, Y and there is no dependence on tie-breaking. Briefly, we can move from \mathbb{R} to \mathbb{R}^{m+1} endowed with the ℓ_1 distance. We pick a small $\epsilon > 0$ and for each $i = 1, \dots, m - z$ we place C_i at $(-1 - i\epsilon, 0, \dots)$. Further, for each $i = 1, \dots, z$ we place B_i at $(1, 0, \dots, 0, \epsilon, 0, \dots, 0)$, where ϵ appears in position $i + 1$. There are $n_1 z$ agents whose location is of the form $(0, \delta_1, \dots, \delta_m)$ and $n_2 z$ agents whose location is of the form $(1, \delta_1, \dots, \delta_m)$, where $0 < \delta_i \leq \epsilon$ for each $i = 1, \dots, m$; for each agent the values of δ_i are chosen so as to induce the desired preference order. For $\epsilon \rightarrow 0$, all calculations in the proof of Theorem 9 remain valid. ■

We will now show upper bounds on the distortion of Plurality and Borda, which match the lower bounds for these rules given by Theorem 9, and an upper bound on the distortion of the Harmonic rule, which almost matches the respective lower bound. Thus, not all positional scoring rules are equally bad from the perspective of distortion: in particular, the Harmonic rule always provides sublinear distortion.

Theorem 11 *For every profile σ over m alternatives the distortion of Plurality and the Borda rule on σ with respect to the sum objective is at most $2m - 1$.*

Proof. Fix an n -voter profile σ over m alternatives and a pseudo-metric $d \in p^{-1}(\sigma)$, and let X be an optimal alternative with respect to d . Let W denote a winning alternative. We will argue that both for Plurality and for the Borda rule we have $|WX| \geq \frac{n}{m}$; we will then invoke Lemma 6.

For Plurality, this claim is immediate, as at least $\frac{n}{m}$ agents rank W first. Now, consider the Borda rule. For each agent $i \in N$ let $\delta_i = \sigma_i(X) - \sigma_i(W)$; as W is a Borda winner, we have $\sum_{i \in N} \delta_i \geq 0$. For each $i \in XW$ we have $\delta_i \leq -1$ and for each $i \in WX$ we have $\delta_i \leq m - 1$. Now, if $|WX| < \frac{n}{m}$, we have $|XW| > n - \frac{n}{m}$ and hence

$$\sum_{i \in N} \delta_i \leq -1|XW| + (m - 1)|WX| < -\frac{n(m - 1)}{m} + \frac{n(m - 1)}{m} = 0,$$

a contradiction. Thus, $|WX| \geq n/m$ for the Borda rule as well. Combining these bounds on $|WX|$ with Lemma 6 gives us the desired distortion bounds. ■

Theorem 12 *The distortion of the Harmonic rule with respect to the sum objective is asymptotically bounded by $O\left(\frac{m}{\sqrt{\ln m}}\right)$.*

Proof. It will be convenient to extend the definition of the harmonic weight vector to non-integer “indices” and set $s_t = 1/t$ for all $t > 0$. As in the proof of Theorem 9, given a positive integer z , we write $F_z = 1 + \frac{1}{2} + \dots + \frac{1}{z}$.

Fix a preference profile σ over m alternatives and a pseudo-metric $d \in p^{-1}(\sigma)$. Let X be an optimal alternative for σ with respect to d , and let W be a winner under the Harmonic rule. Set $\delta = d(W, X)$. Let

$$U = \left\{ i \in N : d(i, X) < \frac{\delta}{6} \right\}, \quad u = |U|.$$

We have

$$\begin{aligned} \frac{\sum_{i \in N} d(i, W)}{\sum_{i \in N} d(i, X)} &\leq \frac{\sum_{i \in N} (d(i, X) + d(X, W))}{\sum_{i \in N} d(i, X)} = 1 + \frac{n\delta}{\sum_{i \in N} d(i, X)} \\ &\leq 1 + \frac{n\delta}{\sum_{i \in N \setminus U} d(i, X)} \leq 1 + \frac{n\delta}{(n - u)\delta/6} = 1 + \frac{6n}{n - u}. \end{aligned}$$

It remains to argue that $n - u = n \cdot \Omega\left(\frac{\sqrt{\ln m}}{m}\right)$. Let

$$\mathcal{G} = \left\{ Y \in \mathcal{A} : d(Y, X) < \frac{\delta}{3} \right\}, \quad z = |\mathcal{G}|.$$

For each $i \in U$, let $s(i)$ be the score that X receives from i ; we have

$$s(i) = s_{\sigma_i(X)} = \frac{1}{\sigma_i(X)}.$$

Also, let

$$\zeta(i) = F_{\sigma_i(X)} = s_1 + \dots + s_{\sigma_i(X)}.$$

Fix a voter $i \in U$. For each $Y \in \mathcal{G}$ we have

$$\begin{aligned} d(i, Y) &\leq d(i, X) + d(X, Y) < \frac{\delta}{6} + \frac{\delta}{3} = \frac{\delta}{2}, \\ d(i, W) &\geq d(X, W) - d(i, X) > \delta - \frac{\delta}{6} = \frac{5\delta}{6}, \end{aligned}$$

so i prefers each alternative from \mathcal{G} over W . On the other hand, for each $Z \in \mathcal{A} \setminus \mathcal{G}$ we have

$$d(i, Z) \geq d(Z, X) - d(i, X) > \frac{\delta}{3} - \frac{\delta}{6} = \frac{\delta}{6} > d(i, X),$$

so i prefers X to each alternative from $\mathcal{A} \setminus \mathcal{G}$. Consequently, i 's preference order is of the following form:

$$i: \underbrace{A_{i_1} \succ \dots \succ A_{i_\ell}}_{\text{subset of } \mathcal{G} \setminus \{X\}} \succ X \succ \dots \succ W \succ \dots$$

Thus, the total score of W is at most

$$us_{z+1} + (n - u),$$

whereas the total score of X is at least

$$\sum_{i \in U} s(i).$$

Moreover, the total score that the alternatives in \mathcal{G} get from a voter $i \in U$ is at least $\zeta(i)$ (observe that $X \in \mathcal{G}$), so by the pigeonhole principle the total score of some alternative in \mathcal{G} is at least

$$\frac{1}{z} \sum_{i \in U} \zeta(i).$$

Since W is a winner under the Harmonic rule, we have

$$u \cdot s_{z+1} + (n - u) \geq \sum_{i \in U} s(i), \text{ and} \tag{6}$$

$$u \cdot s_{z+1} + (n - u) \geq \frac{1}{z} \sum_{i \in U} \zeta(i). \tag{7}$$

Let $\xi = \frac{1}{u} \sum_{i \in U} \sigma_i(X)$. By the inequality between the harmonic mean and the arithmetic mean, we have

$$\frac{u}{\sum_{i \in U} s(i)} = \frac{u}{\sum_{i \in U} \frac{1}{\sigma_i(X)}} \leq \frac{\sum_{i \in U} \sigma_i(X)}{u} = \xi,$$

so $\sum_{i \in U} s(i) \geq \frac{u}{\xi} = u \cdot s_\xi$, and inequality (6) implies

$$u \cdot s_{z+1} + (n - u) \geq u \cdot s_\xi. \tag{8}$$

As we have $\sigma_i(X) \leq z$ for each $i \in U$ and $\frac{F_x}{x}$ is a decreasing function of x , we obtain $F_{\sigma_i(X)} \geq \sigma_i(X) \frac{F_z}{z}$ for each $i \in U$ and hence

$$\frac{\sum_{i \in U} \zeta(i)}{u \cdot \xi} = \frac{\sum_{i \in U} \zeta(i)}{\sum_{i \in U} \sigma_i(X)} = \frac{\sum_{i \in U} F_{\sigma_i(X)}}{\sum_{i \in U} \sigma_i(X)} \geq \frac{F_z}{z}.$$

Therefore, inequality (7) implies

$$u \cdot s_{z+1} + (n - u) \geq \frac{u\xi \cdot F_z}{z^2}. \quad (9)$$

Now, if $\xi \leq z(\ln z)^{-1/2}$ it holds that $s_\xi \geq (\ln z)^{1/2}/z$, whereas if $\xi > z(\ln z)^{-1/2}$, we have $F_z > \ln z$, and hence $\xi \cdot F_z/z^2 > (\ln z)^{1/2}/z$. Together with inequalities (8) and (9), this implies

$$\frac{u}{z+1} + (n - u) \geq u \frac{\sqrt{\ln z}}{z},$$

or, equivalently,

$$n \geq u \left(1 + \frac{\sqrt{\ln z}}{z} - \frac{1}{z+1} \right).$$

As $z \leq m$, it follows that

$$u \leq \frac{n}{1 + \frac{\sqrt{\ln m}}{m} - \frac{1}{m+1}},$$

and, consequently,

$$n - u \geq n \cdot \frac{\frac{\sqrt{\ln m}}{m} - \frac{1}{m+1}}{1 + \frac{\sqrt{\ln m}}{m} - \frac{1}{m+1}} = n \cdot \Omega \left(\frac{\sqrt{\ln m}}{m} \right),$$

which is what we had to prove. ■

Altogether, the results of this section indicate that well-known positional scoring rules have high worst-case distortion; while the Harmonic rule performs slightly better, its distortion is still very substantial. Informally, this is because an optimal alternative can be preferred over the eventual winner by a relatively large fraction of the agents, and yet still lose. In what follows, we will see several voting rules that escape this predicament, resulting in significantly better performance.

3.2 Distortion of the Copeland Rule and Ranked Pairs

In this section we show that for several simple voting rules, including the Copeland rule, the distortion with respect to the sum objective never exceeds 5. We first state and prove some preliminary observations.

Lemma 13 *Consider a vector $\vec{v} \in \mathbb{R}^m$ with $v_1 \geq v_2 \geq \dots \geq v_m \geq 0$ and a pair of vectors $\alpha, \beta \in \mathbb{R}^m$. If for all $k \in [m]$ it holds that $\sum_{i=1}^k \alpha_i \geq \sum_{i=1}^k \beta_i$, then $\sum_{i=1}^m \alpha_i v_i \geq \sum_{i=1}^m \beta_i v_i$.*

Proof. For each $i \in [m]$, let $\gamma_i = \alpha_i - \beta_i$. We use Abel's transform: given two sequences $a_1, \dots, a_m, b_1, \dots, b_m$ it holds that

$$\sum_{i=1}^m a_i b_i = a_m \sum_{i=1}^m b_i - \sum_{i=1}^{m-1} \left(\sum_{j=1}^i b_j \right) (a_{i+1} - a_i).$$

Substituting $b_i = \gamma_i, a_i = v_i$, we obtain

$$\begin{aligned} \sum_{i=1}^m \alpha_i v_i - \sum_{i=1}^m \beta_i v_i &= \sum_{i=1}^m \gamma_i v_i = v_m \sum_{i=1}^m \gamma_i - \sum_{i=1}^{m-1} \left(\sum_{j=1}^i \gamma_j \right) \cdot (v_{i+1} - v_i) \\ &= v_m \sum_{i=1}^m \gamma_i + \sum_{i=1}^{m-1} \left(\sum_{j=1}^i \gamma_j \right) \cdot (v_i - v_{i+1}). \end{aligned} \quad (10)$$

Now, by the conditions of the lemma we have $\sum_{j=1}^i \gamma_j \geq 0$ for all $i \in [m]$ and $v_i - v_{i+1} \geq 0$ for all $i \in [m-1]$, so all summands in (10) are non-negative. Hence, the lemma holds. ■

The next lemma enables us to obtain a lower bound on the cost of a given alternative; we will use it in situations where X is an optimal alternative and W is a winning alternative.

Lemma 14 Consider a profile σ , a pseudo-metric $d \in p^{-1}(\sigma)$, and a pair of alternatives X, W . If

$$\sum_{i \in N} d(i, X) \geq \frac{1}{\gamma} \sum_{i \in XW} \min_{W \succeq_i Z} d(X, Z) \quad (11)$$

for some $\gamma \leq 1$, then $SC_{\Sigma}(W, d) \leq (1 + \gamma) SC_{\Sigma}(X, d)$.

Proof. We use inequality (3) together with the fact that $d(i, W) \leq d(i, X)$ for $i \in WX$ to upper-bound the cost of the alternative W :

$$\begin{aligned} \frac{SC_{\Sigma}(W, d)}{SC_{\Sigma}(X, d)} &= \frac{\sum_{i \in N} d(i, W)}{\sum_{i \in N} d(i, X)} = \frac{\sum_{i \in WX} d(i, W) + \sum_{i \in XW} d(i, W)}{\sum_{i \in N} d(i, X)} \\ &\leq \frac{\sum_{i \in N} d(i, X) + \sum_{i \in XW} \min_{W \succeq_i Z} d(X, Z)}{\sum_{i \in N} d(i, X)} = 1 + \frac{\sum_{i \in XW} \min_{W \succeq_i Z} d(X, Z)}{\sum_{i \in N} d(i, X)}. \end{aligned}$$

Together with assumption (11) this gives the desired result. \blacksquare

Consider a profile σ over a set of alternatives \mathcal{A} . An alternative $W \in \mathcal{A}$ is said to belong to the *uncovered set* of σ if for every $Z \in \mathcal{A}$ it holds that W weakly defeats Z or there is an alternative Y such that W weakly defeats Y and Y weakly defeats Z ; we denote the uncovered set of a profile σ by $UC(\sigma)$ [Moulin, 1986]. For each profile its uncovered set is non-empty; moreover, the Copeland rule always outputs a subset of the uncovered set. To see this, consider a Copeland winner W , and suppose that W defeats k alternatives Y_1, \dots, Y_k , so that its Copeland score is k . If there is an alternative X such that none of the alternatives in $\{W, Y_1, \dots, Y_k\}$ weakly defeats X , then X defeats all these alternatives and hence its Copeland score is higher than that of W , a contradiction. There are several other tournament solution concepts that always output a subset of the uncovered set, such as the Banks set, the bipartisan set, the minimal covering set or the tournament equilibrium set [Laffond *et al.*, 1995].

We will now show that the social cost of every alternative in the uncovered set is within a factor of 5 from optimal.

Theorem 15 Consider a profile σ and a pseudo-metric $d \in p^{-1}(\sigma)$. Let X be an optimal alternative for σ with respect to d . Then for every $W \in UC(\sigma)$ we have $SC_{\Sigma}(W, d) \leq 5 \cdot SC_{\Sigma}(X, d)$.

Proof. Fix a profile σ and a pseudo-metric $d \in p^{-1}(\sigma)$. Let X be an optimal alternative for σ with respect to d , and let W be an alternative in $UC(\sigma)$.

If W weakly defeats X we have $SC_{\Sigma}(W, d) \leq 3 \cdot SC_{\Sigma}(X, d)$ by Lemma 6. Now, suppose that this is not the case. Since W is an element of the uncovered set, it follows that there exists an alternative Y such that W weakly defeats Y and Y weakly defeats X . Thus, we have $|WY| \geq n/2 \geq |YW|$ and $|YX| \geq n/2 \geq |XY|$. We consider the cases where $d(X, Y) \geq d(X, W)$ and $d(X, Y) < d(X, W)$ separately.

1. $d(X, Y) \geq d(X, W)$. We know that at least $n/2$ agents prefer Y over X . By Lemma 5 each of these $n/2$ agents contributes at least $\frac{1}{2} \cdot d(X, Y) \geq \frac{1}{2} \cdot d(X, W)$ to the social cost of X . This observation is all we need to obtain a distortion upper bound of 5. Formally, we have

$$\begin{aligned} \sum_{i \in N} d(i, X) &\geq \sum_{i \in YX} d(i, X) \geq \frac{1}{2} \sum_{i \in YX} d(X, Y) = \frac{1}{2} |YX| \cdot d(X, Y) \geq \frac{n}{4} \cdot d(X, Y) \\ &\geq \frac{n}{4} \cdot d(X, W) \geq \frac{1}{4} |XW| \cdot d(X, W) \geq \frac{1}{4} \sum_{i \in XW} \min_{W \succeq_i Z} d(X, Z). \end{aligned}$$

We now apply Lemma 14 with $\gamma = 4$ to obtain the desired bound.

2. $d(X, Y) < d(X, W)$. In the previous case, we only considered agents in YX when proving a lower bound on $\sum_{i \in N} d(i, X)$. In contrast, now we consider the sets WX , XWY , and YXW .

The following observations follow from basic set theory and our choice of X, Y , and W :

$$|WX| + |XWY| = |YWX| + |WYX| + |WXY| + |XWY|$$

$$\geq |WY| \geq |YW| \geq |YXW| + |XYW| \geq \frac{1}{2}|YXW| + \frac{1}{2}|XYW|; \quad (12)$$

$$\begin{aligned} |WX| + |YXW| &= |YWX| + |WYX| + |WXY| + |YXW| \\ &\geq |YX| \geq \frac{n}{2} \geq \frac{1}{2}(|YXW| + |XYW| + |XWY|). \end{aligned} \quad (13)$$

We will now apply Lemma 13 with

$$\begin{aligned} \alpha_1 &= |WX| + |XWY|, & \alpha_2 &= |YXW|, & \alpha_3 &= -|XWY|, \\ \beta_1 &= \frac{1}{2}|YXW| + \frac{1}{2}|XYW|, & \beta_2 &= \frac{1}{2}|XWY|, & \beta_3 &= 0 \\ v_1 &= d(X, W), & v_2 &= v_3 = d(X, Y). \end{aligned}$$

Note that we have $\alpha_1 \geq \beta_1$ by (12). Moreover, by (13) we have $\alpha_1 + \alpha_2 + \alpha_3 \geq \beta_1 + \beta_2 + \beta_3$; as $\alpha_3 \leq 0$ and $\beta_3 = 0$, this implies $\alpha_1 + \alpha_2 \geq \beta_1 + \beta_2$. Finally, we have $v_1 \geq v_2 = v_3$ by assumption. Thus, the conditions of Lemma 13 are satisfied and we obtain

$$\begin{aligned} \sum_{i=1}^3 \alpha_i v_i &= (|WX| + |XWY|) \cdot d(X, W) + (|YXW| - |XWY|) \cdot d(X, Y) \\ &\geq \sum_{i=1}^3 \beta_i v_i = \frac{1}{2}(|YXW| + |XYW|) \cdot d(X, W) + \frac{1}{2}|XWY| \cdot d(X, Y). \end{aligned} \quad (14)$$

We are now ready to prove a lower bound on the social cost of X . Combining Lemma 5 with (14), we obtain

$$\begin{aligned} \sum_{i \in N} d(i, X) &\geq \sum_{i \in WX} d(i, X) + \sum_{i \in XWY} d(i, X) + \sum_{i \in YXW} d(i, X) \\ &\geq \frac{1}{2}|WX| \cdot d(X, W) + |XWY| \left(\frac{d(X, W) - d(X, Y)}{2} \right) + \frac{1}{2}|YXW| \cdot d(X, Y) \end{aligned} \quad (15)$$

$$\begin{aligned} &= \frac{1}{2}(|WX| + |XWY|) \cdot d(X, W) + \frac{1}{2}(|YXW| - |XWY|) \cdot d(X, Y) \\ &\geq \frac{1}{4}(|YXW| + |XYW|) \cdot d(X, W) + \frac{1}{4}|XWY| \cdot d(X, Y) \end{aligned} \quad (16)$$

$$\geq \frac{1}{4} \sum_{i \in XW} \min_{W \succeq_i Z} d(X, Z). \quad (17)$$

Specifically, line (15) follows from Lemma 5, line (16) follows from (14), and line (17) follows because $XW = YXW \cup XYW \cup XWY$, $W \succeq_i W$ for all $i \in N$ and $W \succeq_i Y$ for all $i \in XWY$.

Given this lower bound on $\sum_{i \in N} d(i, X)$, we can now apply Lemma 14 with $\gamma = 4$, which gives us the desired upper bound of 5 on distortion. ■

Theorem 15 implies that the distortion of the Copeland rule is upper-bounded by 5; moreover, this bound is tight.

Theorem 16 *The distortion of the Copeland rule with respect to the sum objective is at most 5. Moreover, for every $\epsilon > 0$ there is a profile σ such that the distortion of the Copeland rule on σ with respect to the sum objective is at least $5 - \epsilon$.*

Proof. The upper bound holds by Theorem 15, together with the fact that the Copeland rule outputs a subset of the uncovered set. We will now prove that this bound is tight.

Consider a preference profile over three alternatives W, X, Y , where $\frac{n}{2} - 1$ agents rank the alternatives as $Y \succ X \succ W$, $\frac{n}{2} - 1$ agents rank the alternatives as $X \succ W \succ Y$, and the remaining two agents rank the alternatives as $W \succ Y \succ X$. Observe that W defeats Y , Y defeats X , and X defeats W . Thus, the Copeland score of every

alternative is 1 and in particular W is a Copeland winner. Now let the underlying metric be as shown in Figure 1. (The distances not shown in the figure can be chosen to be consistent with the metric and the preference profile.) We get

$$\frac{\sum_{i \in N} d(i, W)}{\sum_{i \in N} d(i, X)} = \frac{\left(\frac{n}{2} - 1\right) \cdot (2 - \epsilon) + \left(\frac{n}{2} - 1\right) \cdot 3 + 20}{\left(\frac{n}{2} - 1\right) \cdot (1 + \epsilon) + 22}$$

Thus as $n \rightarrow \infty$ and $\epsilon \rightarrow 0$, we get instances on which the distortion of the Copeland rule is arbitrarily close to 5. ■




Figure 1: The distortion of any rule that outputs W is arbitrarily close to 5 as $\epsilon \rightarrow 0$ and $n \rightarrow \infty$.

The lower bound in Theorem 16 extends to every neutral voting rule that is defined on the majority graph; in the terminology of Fishburn [1977], such rules are called C1 rules. Note, however, that Ranked Pairs is not a C1 rule. Indeed, while we are not able to prove an upper bound of 3 on the distortion of Ranked Pairs in the general case, we can show that its distortion does not exceed 3 as long as the weak majority graph does not have long cycles. Recall that the *circumference* of a directed graph is the maximum length of a directed cycle in this graph.

Theorem 17 *The distortion of Ranked Pairs with respect to the sum objective is at most 3 as long as the circumference of the weak majority graph does not exceed 4.*

Proof. As usual, fix a profile σ and a pseudo-metric $d \in p^{-1}(\sigma)$, let W be an alternative chosen by Ranked Pairs on σ , and let X be an optimal alternative for σ with respect to d . If W defeats X , the distortion is at most 3 by Lemma 6; thus, let us assume that W does not defeat X , that is, $|XW| \geq n/2 \geq |WX|$. Let G be a graph generated by an execution of Ranked Pairs such that W is a source of G . There must be at least one path in G from W to X : otherwise, when considering the edge (X, W) , we would have added it to G , a contradiction with W being a source. By the same argument, at least one of these paths has the property that the weight of each of its edges is at least $|XW| \geq n/2$. Let P be some such path. Then, P is also a subpath of the weak majority graph, and P together with (X, W) form a cycle in the weak majority graph. By our assumption on the circumference of this graph, this implies that the length of P is at most 3. Assume that P has length 3; the argument for the case when P has length 2 is similar and simpler. Then, P consists of alternatives W, Y, Z, X , where W defeats Y , Y defeats Z , and Z defeats X . with $|WY|, |YZ|, |ZX| \geq |XW|$. Throughout this proof, we let $*W$ denote the set of agents who rank W last out of X, Y, Z , and W ; similarly, $*WZ$ denotes the set of agents who rank Z and W below X and Y , with W ranked above Z , and $*WY$ denotes the set of agents who rank Y and W below X and Z , with W ranked above Y .

We will consider four cases.

1. $d(X, Z) \geq d(X, W)$. Similarly to the first case in the analysis for the uncovered set (Theorem 15), it suffices to consider the agents in ZX . Each of these $|ZX| \geq |XW|$ agents contributes at least $\frac{d(X, Z)}{2} \geq \frac{d(X, W)}{2}$ to the social cost of X , by Lemma 5. This is all we need to obtain an upper bound of 3. Formally, we have

$$\begin{aligned} \sum_{i \in N} d(i, X) &\geq \sum_{i \in ZX} d(i, X) \geq \frac{1}{2} \sum_{i \in ZX} d(X, Z) = \frac{1}{2} |ZX| \cdot d(X, Z) \geq \frac{1}{2} |XW| \cdot d(X, Z) \\ &\geq \frac{1}{2} |XW| \cdot d(X, W) \geq \frac{1}{2} \sum_{i \in XW} \min_{W \succeq_i T} d(X, T). \end{aligned}$$

We obtain the desired result by applying Lemma 14 with $\gamma = 2$.

2. $d(X, W) \geq d(X, Z) \geq d(X, Y)$. We focus on sets WX and ZXW and make the following observations:

$$|WX| = n - |XW| \geq n - |WY| = |YW| \geq |*W|; \quad (18)$$

$$|WX| + |ZXW| \geq |ZX| \geq |XW| \geq |*W| + |*WZ| + |XWY|. \quad (19)$$

We will now invoke Lemma 13 with $\alpha_1 = |WX|$, $\alpha_2 = |ZXW|$, $\beta_1 = |*W|$, and $\beta_2 = |*WZ| + |XWY|$, $v_1 = d(X, W)$, $v_2 = d(X, Z)$; note that $\alpha_1 \geq \beta_1$ by (18) and $\alpha_1 + \alpha_2 \geq \beta_1 + \beta_2$ by (19). Using Lemma 5, the assumption $d(X, Z) \geq d(X, Y)$, and the fact that $XW \subseteq *W \cup *WZ \cup XWY$, we derive that

$$\begin{aligned} \sum_{i \in N} d(i, X) &\geq \sum_{i \in WX \cup ZXW} d(i, X) \\ &\geq \frac{1}{2}|WX| \cdot d(X, W) + \frac{1}{2}|ZXW| \cdot d(X, Z) \\ &\geq \frac{1}{2}|*W| \cdot d(X, W) + \frac{1}{2}|*WZ| \cdot d(X, Z) + \frac{1}{2}|XWY| \cdot d(X, Y) \\ &\geq \frac{1}{2} \sum_{i \in XW} \min_{W \succeq_i T} d(X, T). \end{aligned}$$

We can now apply Lemma 14 with $\gamma = 2$ to obtain the desired result.

3. $d(X, W) \geq d(X, Y) \geq d(X, Z)$. This case is similar to the previous one. We will consider the following pairwise disjoint sets of agents: WX , YXW , $XYZ \cap XW$, and $ZXW \cap ZXY$. We require the following bounds:

$$\begin{aligned} |WX| + |YXW| + |XYZ \cap XW| &\geq |WX \cap YZ| + |YXW \cap YZ| + |XYZ \cap XW| \\ &= |WX \cap YZ| + |XW \cap YX \cap YZ| + |XW \cap XY \cap YZ| \\ &\geq |YZ| \geq |XW| \geq |*W| + |*WY|; \end{aligned} \quad (20)$$

$$\begin{aligned} |WX| + |YXW| + |ZXW \cap ZXY| &\geq |WX \cap ZX| + |YXW \cap ZX| + |ZXW \cap ZXY| \\ &= |WX \cap ZX| + |XW \cap YX \cap ZX| + |XW \cap XY \cap ZX| \\ &\geq |ZX| \geq |XW| \geq |*W| + |*WY| + |XWZ|. \end{aligned} \quad (21)$$

We can now apply Lemma 13 with

$$\begin{aligned} \alpha_1 &= |WX|, \quad \alpha_2 = |YXW| + |XYZ \cap XW|, \quad \alpha_3 = |ZXW \cap ZXY| - |XYZ \cap XW|, \\ \beta_1 &= |*W|, \quad \beta_2 = |*WY|, \quad \beta_3 = |XWZ|; \\ v_1 &= d(X, W), v_2 = d(X, Y), \quad v_3 = d(X, Z). \end{aligned}$$

we have $\alpha_1 \geq \beta_1$ by (18), $\alpha_1 + \alpha_2 \geq \beta_1 + \beta_2$ by (20), and $\alpha_1 + \alpha_2 + \alpha_3 \geq \beta_1 + \beta_2 + \beta_3$ by (21). We obtain

$$\begin{aligned} \sum_{i \in N} d(i, X) &\geq \sum_{i \in WX} d(i, X) + \sum_{i \in YXW} d(i, X) + \sum_{i \in ZXW \cap ZXY} d(i, X) + \sum_{i \in XYZ \cap XW} d(i, X) \\ &\geq \frac{1}{2}|WX| \cdot d(X, W) + \frac{1}{2}|YXW| \cdot d(X, Y) + \frac{1}{2}|ZXW \cap ZXY| \cdot d(X, Z) \end{aligned} \quad (22)$$

$$\begin{aligned} &+ \frac{1}{2}|XYZ \cap XW| \cdot (d(X, Y) - d(X, Z)) \\ &= \frac{1}{2}|WX| \cdot d(X, W) \\ &+ \frac{1}{2}(|YXW| + |XYZ \cap XW|) \cdot d(X, Y) \\ &+ \frac{1}{2}(|ZXW \cap ZXY| - |XYZ \cap XW|) \cdot d(X, Z) \end{aligned} \quad (23)$$

$$\geq \frac{1}{2} |*W| \cdot d(X, W) + \frac{1}{2} |*WY| \cdot d(X, Y) \quad (24)$$

$$+ \frac{1}{2} |XWZ| \cdot d(X, Z) \quad (25)$$

$$\geq \frac{1}{2} \sum_{i \in XW} \min_{W \succeq_i T} d(X, T), \quad (26)$$

where the inequality in line (22) follows from Lemma 5 (inequalities (1) and (2)) and the inequality in line (24) follows from Lemma 13. Again, we invoke Lemma 14 with $\gamma = 2$ to obtain the desired result.

4. $d(X, Y) \geq d(X, W) \geq d(X, Z)$. The analysis is very similar to the previous case. We use Lemma 13 with

$$\begin{aligned} \alpha_1 &= |WX| + |YXW| + |XYZ \cap XW|, \beta_1 = |*W| + |*WY| \\ \alpha_2 &= |ZXW \cap ZXY| - |XYZ \cap XW|, \beta_2 = |XWZ|, \\ v_1 &= d(X, W), v_2 = d(X, Z); \end{aligned}$$

we have $\alpha_1 \geq \beta_1$ by (20) and $\alpha_1 + \alpha_2 \geq \beta_1 + \beta_2$ by (21). Using Lemma 5 and Lemma 13 and the assumption $d(X, Y) \geq d(X, W)$, we obtain

$$\begin{aligned} \sum_{i \in N} d(i, X) &\geq \frac{1}{2} |WX| \cdot d(X, W) \\ &\quad + \frac{1}{2} (|YXW| + |XYZ \cap XW|) \cdot d(X, Y) \\ &\quad + \frac{1}{2} (|ZXW \cap ZXY| - |XYZ \cap XW|) \cdot d(X, Z) \\ &\geq \frac{1}{2} (|WX| + |YXW| + |XYZ \cap XW|) \cdot d(X, W) \\ &\quad + \frac{1}{2} (|ZXW \cap ZXY| - |XYZ \cap XW|) \cdot d(X, Z) \\ &\geq \frac{1}{2} (|*W| + |*WY|) \cdot d(X, W) + \frac{1}{2} |XWZ| \cdot d(X, Z) \\ &\geq \frac{1}{2} \sum_{i \in XW} \min_{W \succeq_i T} d(X, T). \end{aligned}$$

Lemma 14 with $\gamma = 2$ gives us the desired result.

This concludes the proof that the distortion of Ranked Pairs with respect to the sum objective is at most 3 when there is a “heavy” W – X path of length 3 in the weak majority graph. The case where the length of such path is 2 is similar (and simpler), and the case where the length of such path is 1 is simply the case where W defeats X . ■

In the conference version of this paper, we conjectured that the distortion of Ranked Pairs is always bounded by 3. However, this conjecture was recently disproved by Goel *et al.* [2017], who showed that the distortion of Ranked Pairs can be arbitrarily close to 5.

3.3 Distortion of Single Transferable Vote

In this section, we focus on STV and show that it has fairly low distortion. Specifically, we demonstrate that its distortion grows at most logarithmically with m ; note that this is a much better upper bound than the bounds for the Harmonic rule provided by Theorems 9 and 11. However, we also show that STV is not as good as the Copeland rule or other rules that choose from the uncovered set, by proving a non-constant lower bound.

Theorem 18 *The distortion of STV with respect to the sum objective is asymptotically bounded by $O(\ln m)$.*




Figure 2: The sequence of balls from the the proof of Theorem 18.

Proof. Fix a preference profile σ with n voters and m alternatives and a pseudo-metric $d \in p^{-1}(\sigma)$. Let X be an optimal alternative for σ with respect to d , and let W be a winning alternative under STV. We set $d = d(X, W)$. Pick a value of γ in the open interval $(\frac{2}{3}, 1)$. Let

$$x = 2 \left\lceil \log_{\frac{\gamma}{1-\gamma}} \left(m \cdot \frac{2\gamma-1}{3\gamma-2} \right) \right\rceil + 1;$$

note that x is an odd integer and $x = O(\ln m)$. Set $r = \frac{d}{2x}$.

For $i = 1, \dots, x+1$, let \mathcal{B}_i be a ball with center X and radius $(2i-1)r$ (see Figure 2). Note that $W \in \mathcal{B}_{x+1} \setminus \mathcal{B}_x$. We will now argue that \mathcal{B}_1 contains at most γn agents.

For the sake of contradiction, assume that this is not the case, i.e., that $|\mathcal{B}_1 \cap N| > \gamma n$. Fix an elimination sequence that results in W being the last surviving alternative. We will say that an alternative Y is *supported by an agent* j at some stage of the STV elimination procedure if Y is the closest not-yet-removed alternative to j . For each $i \in \{1, \dots, x\}$, let Z_{i-1} be the last alternative from \mathcal{B}_{i-1} to be removed by STV, and let y_i denote the number of alternatives in $\mathcal{B}_i \setminus \mathcal{B}_{i-1}$ just before Z_{i-1} is removed.

Consider $i \leq x-2$. For every agent $j \in \mathcal{B}_1$ and every alternative $Y \notin \mathcal{B}_{i+1}$ it holds that

$$d(j, Z_i) \leq d(j, X) + d(X, Z_i) \leq r + (2i-1)r = (2i+1)r - r < d(Y, X) - d(j, X) \leq d(j, Y).$$

Hence, just before Z_i is removed, each agent in \mathcal{B}_1 supports an alternative in \mathcal{B}_{i+1} . Thus, from the pigeonhole principle we infer that at this moment there exists an alternative $Y \in \mathcal{B}_{i+1}$ that is supported by more than $\frac{\gamma n}{y_{i+1}+1}$ agents from \mathcal{B}_1 . Consequently, when STV decides to remove Y , all surviving alternatives in $\mathcal{B}_{i+3} \setminus \mathcal{B}_{i+2}$ are supported by more than $\frac{\gamma n}{y_{i+1}+1}$ agents. None of these agents is in \mathcal{B}_1 , as all agents in \mathcal{B}_1 prefer Y to every $Y' \in \mathcal{B}_{i+3} \setminus \mathcal{B}_{i+2}$: indeed, for $j \in \mathcal{B}_1$ we have

$$d(j, Y) \leq d(j, X) + d(X, Y) \leq 2(i+1)r < d(Y', X) - d(j, X) \leq d(j, Y').$$

Thus, we get

$$y_{i+3} \frac{\gamma n}{y_{i+1}+1} < n(1-\gamma),$$

and hence $y_{i+1} > \frac{\gamma}{1-\gamma} y_{i+3} - 1$. Set $\xi = \frac{\gamma}{1-\gamma}$; observe that $\gamma \in (\frac{2}{3}, 1)$ implies $\xi > 1$ and hence $\frac{1}{1-\xi} < 0$ (this fact will be useful in the sequence of inequalities below). Recall that $W \in \mathcal{B}_{x+1} \setminus \mathcal{B}_x$, so $y_{x+1} \geq 1$. We have

$$\begin{aligned} y_1 &> \xi y_3 - 1 \geq \xi^2 y_5 - \xi - 1 \geq \dots \\ &\geq \xi^{\frac{x-1}{2}} - \xi^{\frac{x-1}{2}-1} - \dots - 1 = \xi^{\frac{x-1}{2}} - \frac{1 - \xi^{\frac{x-1}{2}}}{1 - \xi} \\ &= \xi^{\frac{x-1}{2}} \left(1 + \frac{1}{1 - \xi} \right) - \frac{1}{1 - \xi} \geq \xi^{\frac{x-1}{2}} \cdot \frac{2 - \xi}{1 - \xi} \\ &= \left(\frac{\gamma}{1 - \gamma} \right)^{\frac{x-1}{2}} \cdot \frac{3\gamma - 2}{2\gamma - 1} \geq m \frac{2\gamma - 1}{3\gamma - 2} \cdot \frac{3\gamma - 2}{2\gamma - 1} = m. \end{aligned}$$




Figure 3: The metric space in the proof of Theorem 19 for $h = 4$. The number in a node denotes the number of agents in that node.

Thus, we obtain $y_1 > m$, a contradiction. We conclude that \mathcal{B}_1 contains at most γn agents. Let us now assess the distortion:

$$\begin{aligned} \frac{\sum_{j \in N} d(j, W)}{\sum_{j \in N} d(j, X)} &\leq \frac{\sum_{j \in N} (d(j, X) + d(X, W))}{\sum_{j \in N} d(j, X)} = 1 + \frac{nd}{\sum_{j \in N} d(j, X)} \\ &\leq 1 + \frac{nd}{\sum_{j \in N \setminus \mathcal{B}_1} d(j, X)} \leq 1 + \frac{nd}{n(1-\gamma)r} = 1 + \frac{2x}{1-\gamma}. \end{aligned}$$

Since γ is a constant and $x = O(\ln m)$, the distortion is upper-bounded by $O(\ln m)$. ■

Our next theorem establishes a lower bound of the distortion of STV.

Theorem 19 *The maximum distortion of STV with respect to the sum objective over all profiles with m alternatives is $\Omega(\sqrt{\ln m})$.*

Proof. Given a positive integer h , we construct a perfectly balanced tree of height h and then connect all leaves to one additional node (see Figure 3). We say that all leaves belong to the first layer; for $i > 1$ layer i consists of parents of the nodes at level $i-1$. For $i = 2, \dots, h$, each node at level i has y_i children, where $y_i = 2^i + 2^{i-2} - 2$. We denote by z_i the number of nodes at level i .

We define the length of each edge of our graph to be one, and define the distance between a pair of nodes to be the length of a shortest path between these nodes.

We place one alternative in each node, including the node that is connected to all leaves. As we have $y_i = 2^i + 2^{i-2} - 2$ for $2 \leq i \leq h$, and hence $y_i \leq 2^{i+1}$ for $i \in [h]$, the total number of leaves can be upper-bounded as

$$z_1 \leq 2^{h+1} \cdot 2^h \cdot \dots \cdot 2 = 2^{\frac{(h+1)(h+2)}{2}} \leq 2^{(h+1)^2};$$

as the degree of each internal node is at least 2, this implies that the number of alternatives m is at most $2 \cdot 2^{(h+1)^2}$. From this, we conclude that $h \geq \sqrt{\log_2 m} - 2$.

The agents' positions are defined as follows. We place one agent in each leaf node. Let S_i denote the total number of agents in a subtree rooted in a level- i node; we have $S_1 = 1$. Now, for each $i = 2, \dots, h$, we compute S_{i-1} and place S_{i-1} agents in each node of layer i . Thus, we have $S_i = S_{i-1}(y_i + 1)$. Note that there are exactly z_1 agents in the bottom layer.

There are $z_i S_{i-1}$ agents at level i , and we have $z_i = z_{i+1} y_{i+1}$. Thus,

$$\frac{z_{i+1} S_i}{z_i S_{i-1}} = \frac{S_i}{y_{i+1} S_{i-1}} = \frac{y_i + 1}{y_{i+1}} = \frac{2^i + 2^{i-2} - 1}{2^{i+1} + 2^{i-1} - 2} = \frac{1}{2},$$

i.e., layer i contains twice as many agents as layer $i + 1$. As layer 1 has z_1 agents, the number of agents in layer i is equal to $z_1 2^{-(i-1)}$.

Let Y_0 be the alternative located in the node connected to all the leaves. STV would first remove Y_0 , as no agent ranks it first, and every other alternative is ranked first by at least one agent. STV can then remove all alternatives located in the leaves, one by one: initially, each such alternative is ranked first by exactly one agent, and no leaf alternative gains additional votes as other leaf alternatives are removed. Inductively, suppose that STV has removed all alternatives in layers $1, \dots, i - 1$, and all other alternatives are still present. Then an alternative in layer i is ranked first among the remaining alternatives by the S_i agents in the respective subtree, and each alternative in layer j , $j > i$, is ranked first by the $S_{j-1} \geq S_i$ agents who are located in the same node as that alternative. Thus, STV can remove the alternatives in layer i one by one. We conclude that the root of the tree can be selected as the winner. As there is an agent in each leaf, the total distance of the agents to the root, which we will denote by d_{stv} , is at least $z_1 h$.

In contrast, the total distance of the agents to the alternative in the node connecting all the leaves can be upper-bounded as

$$\begin{aligned} d_{\text{bot}} &= z_1 + \frac{z_1}{2} + \dots + \frac{z_1}{2^{h-1}} h = z_1 \sum_{i=1}^h i 2^{-(i-1)} \\ &= 4z_1 (1 - (h+1)2^{-h} + h2^{-h-1}) \leq 4z_1. \end{aligned}$$

Thus, we can lower-bound the distortion in our example as

$$\frac{d_{\text{stv}}}{d_{\text{bot}}} \geq \frac{h}{4} \geq \frac{\sqrt{\log_2 m} - 2}{4}.$$

This completes the proof. ■

3.4 Improving the Worst-Case Distortion: $(m - 1)$ -simplex

We showed that for the general metric setting, among all voting rules that we have considered, the Copeland rule has the lowest worst-case distortion. However, the worst-case distortion of the Copeland rule can be arbitrarily close to 5, so there is still a gap between the performance of this rule and our lower bound of 3 (Theorem 4). In this section we consider an interesting special setting for which this gap disappears, and Plurality, STV, and Ranked Pairs all have best-possible worst-case distortion.

Specifically, instead of considering arbitrary metric spaces, we now assume that the m alternatives are the vertices of the $(m - 1)$ -simplex endowed with the Euclidean metric (i.e., it holds that $d(X, Y) = d(X, Z)$ for all alternatives X, Y, Z such that $Y, Z \neq X$), and that the agents are located inside this simplex. This choice of metric may be appropriate for scenarios where the alternative space does not have any obvious structure, so it is natural to associate the alternatives with pairwise orthogonal directions. We note that Theorem 4 remains true for the simplex, so the worst-case distortion of every deterministic voting rule with respect to the sum objective is at least 3 even in this case. However, the constructions establishing some of the stronger lower bounds for specific voting rules cannot be embedded into the simplex and, indeed, the worst-case distortion of several voting rules improves considerably. Intuitively, there are two reasons for our positive results in this setting: first, all pairwise distances between distinct alternatives are equal, and second, an agent cannot be too close to an alternative that is not her top choice.

We start by showing that Plurality has best possible distortion in the $(m - 1)$ -simplex setting, i.e., its worst-case distortion with respect to the sum objective matches the lower bound of 3.

Theorem 20 *In the $(m - 1)$ -simplex setting, the distortion of Plurality with respect to the sum objective is at most 3.*

Proof. Fix an n -voter profile σ over a set of alternatives \mathcal{A} , $|\mathcal{A}| = m$. Let $d \in p^{-1}(\sigma)$ be a Euclidean metric that associates each alternative in \mathcal{A} with a vertex of an $(m - 1)$ -simplex. Let X be an optimal alternative for σ with respect to d , and let W be a Plurality winner. For each alternative $Z \in \mathcal{A}$, let N_Z be the set of all agents that rank Z first; note that we have $|N_W| \geq |N_X|$ by definition of the Plurality rule and $|WX| \geq |N_W|$.

We will first derive a strong lower bound on the cost of X in the $(m - 1)$ -simplex setting. Consider an agent $i \in N_Z$ for some $Z \neq X$. By Lemma 5 we have $d(i, X) \geq \frac{d(X, Z)}{2} = \frac{d(X, W)}{2}$. Thus, we obtain $\sum_{i \in N} d(i, X) \geq$

$\frac{d(X,W)}{2}(n - |N_X|)$. Using this lower bound, we can build on inequality (4) in Lemma 6 and upper-bound the distortion of Plurality on σ as

$$1 + \frac{2(n - |WX|)d(X,W)}{d(X,W)(n - |N_X|)} = 1 + \frac{2(n - |WX|)}{n - |N_X|} \leq 1 + \frac{2(n - |WX|)}{n - |N_W|} \leq 1 + \frac{2(n - |WX|)}{n - |WX|} = 3. \quad (27)$$

Ranked Pairs and STV also achieve the best possible distortion with respect to the sum objective in the $(m-1)$ -simplex setting.

Theorem 21 *In the $(m-1)$ -simplex setting, the distortion of Ranked Pairs with respect to the sum objective is at most 3.*

Proof. Fix an n -voter profile σ over a set of alternatives \mathcal{A} , $|\mathcal{A}| = m$. Let $d \in p^{-1}(\sigma)$ be a Euclidean metric that associates each alternative in \mathcal{A} with a vertex of an $(m-1)$ -simplex. Let X be an optimal alternative for σ with respect to d , and let W be a Ranked Pairs winner. In our analysis of Ranked Pairs (Theorem 17) we have established that one of the following statements holds: (i) $W = X$, in which case our claim is trivially true, (ii) W weakly defeats X , in which case we are done by Lemma 6, or (iii) there is another alternative Z that weakly defeats X . In case (iii), in the $(m-1)$ -simplex setting we have $d(X, Z) = d(X, W)$, and we have argued in the proof of Theorem 17 that in this case the distortion of Ranked Pairs with respect to the sum objective does not exceed 3. ■

Theorem 22 *In the $(m-1)$ -simplex setting, the distortion of STV with respect to the sum objective is at most 3.*

Proof. Fix a profile σ over a set of alternatives \mathcal{A} , $|\mathcal{A}| = m$. Let $d \in p^{-1}(\sigma)$ be a Euclidean metric that associates each alternative in \mathcal{A} with a vertex of an $(m-1)$ -simplex. Let X be an optimal alternative for σ with respect to d , and let W be an STV winner. Consider an execution of STV that results in W being elected, and let t be the round in which X is eliminated. For each alternative $Z \in \mathcal{A}$, let N_Z denote the set of agents who rank Z as their most preferred alternative in the beginning of round t , i.e., just before X is eliminated. Clearly, $|N_X| \leq |N_W|$.

Consider an alternative $Y \in \mathcal{A} \setminus \{X, W\}$. We have $N_Y \subseteq YX$, so by Lemma 5 for each agent $i \in N_Y$ we have $2d(i, X) \geq d(Y, X)$ and hence

$$d(i, W) \leq d(i, Y) + d(Y, W) \leq d(i, X) + d(Y, W) = d(i, X) + d(Y, X) \leq 3d(i, X).$$

Consequently, we obtain

$$\sum_{i \notin N_W \cup N_X} d(i, W) \leq 3 \sum_{i \notin N_W \cup N_X} d(i, X). \quad (28)$$

Next, let us consider agents from $N_X \cup N_W$:

$$\begin{aligned} \sum_{i \in N_W \cup N_X} d(i, W) &\leq \sum_{i \in N_W \cup N_X} d(i, X) + |N_X|d(X, W) \\ &\leq \sum_{i \in N_W \cup N_X} d(i, X) + |N_W|d(X, W) \leq \sum_{i \in N_W \cup N_X} d(i, X) + \sum_{i \in N_W} (d(i, X) + d(i, W)) \\ &\leq \sum_{i \in N_W \cup N_X} d(i, X) + 2 \sum_{i \in N_W} d(i, X) \leq 3 \sum_{i \in N_W \cup N_X} d(i, X). \end{aligned} \quad (29)$$

Combining inequalities (28) and (29), we obtain the desired bound. ■

On the other hand, for Borda, Veto, k -Approval with $k > 1$, and the Copeland rule, the worst-case distortion with respect to the sum objective remains unchanged in the $(m-1)$ -simplex setting; the constructions establishing the lower bounds for these rules in the general metric setting can be adjusted to work in the $(m-1)$ -simplex setting. The question for other scoring rules, including the Harmonic rule, remains open.

4 Distortion of Median Agent Cost

In this section we study the distortion of deterministic voting rules as measured by the median agent cost. As a shorthand, when the pseudo-metric d is fixed, we denote the median social cost of an alternative Y by $\text{med}(Y)$. Recall that we define $\text{med}(Y)$ as the $\lfloor \frac{n}{2} + 1 \rfloor$ -th smallest cost, i.e., if $n = 8$, four agents are located at distance 1 from Y and four agents are located at distance 3 from Y then $\text{med}(Y) = 3$.

The following simple lemma, which follows from the triangle inequality, will be useful in our analysis.

Lemma 23 *For every pair of alternatives Y and Z we have $\text{med}(Z) \leq \text{med}(Y) + d(Y, Z)$.*

Proof. There are at least $\lfloor \frac{n}{2} + 1 \rfloor$ agents i with $d(i, Y) \leq \text{med}(Y)$. For every such agent i we have $d(i, Z) \leq d(i, Y) + d(Y, Z) \leq \text{med}(Y) + d(Y, Z)$. Hence the claim follows. ■

For profiles that admit weak Condorcet winners, we obtain the same result as for the sum objective (Corollary 7): by selecting a weak Condorcet winner, we ensure that the distortion does not exceed 3.

Theorem 24 *Consider an n -voter profile σ and a pseudo-metric $d \in p^{-1}(\sigma)$. Let X be an optimal alternative for σ with respect to d . If W weakly defeats X then $\text{SC}_{\text{med}}(W, d) \leq 3 \cdot \text{SC}_{\text{med}}(X, d)$.*

Proof. We split the analysis into two cases: $\text{med}(W) > \frac{3}{2} \cdot d(X, W)$ and $\text{med}(W) \leq \frac{3}{2} \cdot d(X, W)$.

1. $\text{med}(W) = \beta \cdot d(X, W)$, $\beta > \frac{3}{2}$. By Lemma 23, we have

$$\frac{\text{med}(W)}{\text{med}(X)} \leq \frac{\text{med}(W)}{\text{med}(W) - d(W, X)} = \frac{\beta \cdot d(X, W)}{(\beta - 1) \cdot d(X, W)} = 1 + \frac{1}{\beta - 1}.$$

As we assume $\beta > \frac{3}{2}$, the desired bound follows.

2. $\text{med}(W) \leq \frac{3}{2} \cdot d(X, W)$. For each agent i in WX we have $d(i, X) \geq \frac{1}{2} \cdot d(W, X)$ (Lemma 5). As $|WX| \geq n/2$, it follows that $\text{med}(X) \geq \frac{1}{2} \cdot d(X, W)$. This immediately implies the desired bound. ■

For $|\mathcal{A}| = 2$ Theorem 24 implies that the distortion of the weak majority rule with respect to the median objective does not exceed 3. However, this result does not extend to the case $|\mathcal{A}| \geq 3$; indeed, in this case we obtain a lower bound of 5 that applies to all deterministic voting rules.

Theorem 25 *For $m > 2$ no deterministic voting rule has worst-case distortion less than 5 with respect to the median objective.*

Proof. Let $\mathcal{A} = \{X, Y, W\}$. Suppose that there are $n/3$ agents corresponding to each of the preference rankings $W \succ Y \succ X$, $Y \succ X \succ W$ and $X \succ W \succ Y$. Without loss of generality, suppose that the given voting rule picks W as a winner. Consider an underlying metric as shown in Figure 4. (The distances not shown in the figure can be chosen to be consistent with the metric and the preference profile). In this instance, we have $\text{med}(W) = 5 + \epsilon$ and $\text{med}(X) = 1 + \epsilon$. Thus, the distortion approaches 5 as $\epsilon \rightarrow 0$. ■

For the median objective the performance of positional scoring rules is even worse than for the sum objective: we will now show that for $m > 2$ every positional scoring rule has unbounded distortion.

Theorem 26 *For $m > 2$ the worst-case distortion of every positional scoring rule with respect to the median objective is unbounded.*

Proof. Fix the number of alternatives $m > 2$ and the scoring vector $\vec{s} = (s_1, \dots, s_m)$ with $s_1 > s_m$; we can assume without loss of generality that $s_m = 0$. All of our constructions associate agents and alternatives with points in \mathbb{R} and use the standard Euclidean metric. Let ϵ be a small positive value.

Suppose first $s_2 = 0$, i.e., our rule is equivalent to Plurality. We construct an n -voter profile over the set of alternatives \mathcal{A} , $|\mathcal{A}| = m$, where $n \geq 6$ and n is divisible by 3. We place $n/3 - 1$ agents and alternative Y at $x = -\epsilon$,




Figure 4: If we pick W as a winning alternative, the worst-case distortion with respect to the median objective approaches 5 as $\epsilon \rightarrow 0$.

$n/3$ agents and alternative X at $x = 0$, and $n/3 + 1$ agents and alternative W at $x = 1$. All the remaining alternatives are located between $x = 1 + \epsilon$ and $x = 1 + 2\epsilon$. We have $\text{med}(X) = \epsilon$, $\text{med}(W) = 1$. Since Plurality chooses W and ϵ can be arbitrarily small, this proves our claim.

Now, suppose that $s_2 > 0$. We construct an n -voter profile over the set of alternatives \mathcal{A} , $|\mathcal{A}| = m$, where n is divisible by 2 and $n > \frac{4s_1}{s_2}$. We place $n/2 + 1$ agents and alternative X at $x = 0$, $n/2 - 1$ agents and alternative W at $x = 1$, and all the remaining alternatives between $x = 1 + \epsilon$ and $x = 1 + 2\epsilon$. Then the score of X is $(n/2 + 1) \cdot s_1$, whereas the score of W is $(n/2 + 1) \cdot s_2 + (n/2 - 1) \cdot s_1 = (n/2 + 1) \cdot s_1 + (n/2 + 1) \cdot s_2 - 2s_1$. If $n > \frac{4s_1}{s_2}$, the score of W is higher than that of X , so an alternative $Y \in \mathcal{A} \setminus \{X\}$ will be a winner, with $\text{med}(Y) \geq 1$. On the other hand, $\text{med}(X) = 0$, which proves that the distortion is unbounded. ■

We will now show that every voting rule that outputs a subset of the uncovered set (and hence, in particular, the Copeland rule) achieves the best possible distortion with respect to the median objective.

Theorem 27 *Consider a voting rule f such that for every profile σ it holds that $f(\sigma) \subseteq \text{UC}(\sigma)$. The distortion of f with respect to the median objective does not exceed 5.*

Proof. Fix a profile σ , an alternative $W \in \text{UC}(\sigma)$ and a pseudo-metric $d \in p^{-1}(\sigma)$. Let X be an optimal alternative with respect to d . If W weakly defeats X , then by Theorem 24 we have $\text{med}(W)/\text{med}(X) \leq 3$. Thus, assume that X defeats W . As $W \in \text{UC}(\sigma)$, there exist an alternative Y such that W weakly defeats Y and Y weakly defeats X . We split the analysis into two cases: $\text{med}(W) > \frac{5}{4} \cdot d(X, W)$ and $\text{med}(W) \leq \frac{5}{4} \cdot d(X, W)$.

1. $\text{med}(W) = \beta \cdot d(X, W)$, $\beta > \frac{5}{4}$. By Lemma 23, $\text{med}(X) \geq \text{med}(W) - d(X, W) = (\beta - 1) \cdot d(X, W)$. Hence the distortion in this case is at most $\frac{\beta}{\beta - 1} \leq 5$.
2. $\text{med}(W) \leq \frac{5}{4} \cdot d(X, W)$. At least half of the agents prefer Y to X and by Lemma 5 for each $i \in YX$ we have $d(i, X) \geq \frac{1}{2} \cdot d(X, Y)$. Similarly, at least half of the agents prefer W to Y , and by Lemma 5 for each $i \in WY$ we have $d(i, X) \geq \frac{1}{2} \cdot (d(X, W) - d(X, Y))$. Thus,

$$\begin{aligned} \text{med}(X) &\geq \frac{1}{2} \cdot \max \{d(X, Y), d(X, W) - d(X, Y)\}, \text{ and hence} \\ \frac{\text{med}(W)}{\text{med}(X)} &\leq \frac{\frac{5}{4} \cdot d(X, W)}{\frac{1}{2} \cdot \max \{d(X, Y), d(X, W) - d(X, Y)\}}. \end{aligned}$$

The above term achieves its maximum value when $d(X, W) = 2 \cdot d(X, Y)$, thus giving us the desired bound. ■

Finally, we note that the construction in the proof of Theorem 19 also gives a lower bound of $\Omega(\sqrt{\ln m})$ for the distortion of STV with respect to the median objective. However, it is not clear if the upper bound on the distortion of STV (Theorem 18) can be adapted to the median objective; we leave this question for future work.

4.1 Generalizing the Median: Percentile Distortion

While in many applications it is natural to consider the cost of the median agent, sometimes we may also be interested in other order statistics, such as, e.g., the satisfaction of the agents in the 25-th or 75-th percentile. We can generalize the median objective function $\text{med}(Y)$ by using percentiles, as follows. Sort the set $\{d(i, Y) : i \in N\}$ in non-decreasing order; for each $\alpha \in [0, 1)$ we let $\alpha\text{-PC}(Y)$ be the $\lfloor \alpha \cdot |\{d(i, Y) : i \in N\}| + 1 \rfloor$ -th element in this order. Thus $\alpha\text{-PC}(Y)$ with $\alpha = 1/2$ is equal to $\text{med}(Y)$. The definition of distortion with respect to the $\alpha\text{-PC}$ objective is similar to the definitions in Section 2: we consider the ratio of $\alpha\text{-PC}(W)$ and $\alpha\text{-PC}(X)$, where W is a winning alternative and X is an optimal alternative.

Theorem 28 provides lower bounds on distortion of all deterministic voting rules with respect to the $\alpha\text{-PC}$ objective, for all values of α . Then in Theorems 30 and 31 we identify voting rules that achieve these bounds for some ranges of values of α .

Theorem 28 *The worst-case distortion of every deterministic voting rule with respect to $\alpha\text{-PC}$ is*

- at least 3 for $\alpha \in [\frac{2}{3}, 1)$;
- at least 5 for $\alpha \in [\frac{1}{2}, \frac{2}{3})$;
- unbounded for $\alpha \in [0, \frac{1}{2})$.

Proof. 1. $\alpha \in [\frac{2}{3}, 1)$. The construction in Figure 5 shows that for every deterministic voting rule its distortion with respect to $\alpha\text{-PC}$ can be arbitrarily close to 3 for all $\alpha \in [\frac{1}{2}, 1)$, so in particular for all $\alpha \in [\frac{2}{3}, 1)$. Specifically, we pick a small $\epsilon > 0$, place alternatives X and W at $x = 0$ and $x = -2$, respectively, and place $n/2$ agents at $x = 1 + \epsilon$ and $n/2$ agents at $x = -1 - \epsilon$. Then half of the agents prefer W to X and half of the agents prefer X to W , so without loss of generality we can assume that the given voting rule picks W as a winner. For $\alpha \geq \frac{1}{2}$, we have $\alpha\text{-PC}(X) = 1 + \epsilon$ and $\alpha\text{-PC}(W) = 3 + \epsilon$, making distortion arbitrarily close to 3 as $\epsilon \rightarrow 0$.




Figure 5: For a voting rule that chooses W the distortion with respect to $\alpha\text{-PC}$ approaches 3 as $\epsilon \rightarrow 0$, for all $\alpha \geq \frac{1}{2}$.

2. $\alpha \in [\frac{1}{2}, \frac{2}{3})$. The argument in Theorem 25 extends to all values of α in $[\frac{1}{2}, \frac{2}{3})$.
3. $\alpha \in [0, \frac{1}{2})$. We can use the construction in the proof of Theorem 4: we have $\alpha\text{-PC}(X) = 0$, $\alpha\text{-PC}(W) = 1 - \epsilon$ for all α in $[0, \frac{1}{2})$. ■

We now give upper bounds of the distortion of Plurality and the Copeland rule to show they are optimal for certain ranges of values of α . The following simple lemma, which is a generalization of Lemma 23, will be used in our proofs.

Lemma 29 *For every pair of alternatives Y and Z and every $\alpha \in [0, 1)$ we have $\alpha\text{-PC}(Z) \leq \alpha\text{-PC}(Y) + d(Y, Z)$.*

Proof. For each agent i with $d(i, Y) \leq \alpha\text{-PC}(Y)$ we have $d(i, Z) \leq d(i, Y) + d(Y, Z) \leq \alpha\text{-PC}(Y) + d(Y, Z)$. Hence the result follows. ■

Theorem 30 *The distortion of Plurality with respect to the $\alpha\text{-PC}$ objective with $\alpha \geq \frac{m-1}{m}$ is at most 3.*

Proof. This proof is similar to that of Theorem 24. Fix $\alpha \in [\frac{m-1}{m}, 1)$, an n -voter profile σ over m alternatives and a pseudo-metric $d \in p^{-1}(\sigma)$. Let W be a Plurality winner and let X be an optimal alternative with respect to d . We consider two cases: $\alpha\text{-PC}(W) > \frac{3}{2} \cdot d(X, W)$ and $\alpha\text{-PC}(W) \leq \frac{3}{2} \cdot d(X, W)$.

1. $\alpha\text{-PC}(W) = \beta \cdot d(X, W)$, $\beta \geq \frac{3}{2}$. By Lemma 29, we obtain

$$\frac{\alpha\text{-PC}(W)}{\alpha\text{-PC}(X)} \leq \frac{\alpha\text{-PC}(W)}{\alpha\text{-PC}(W) - d(X, W)} = \frac{\beta \cdot d(X, W)}{(\beta - 1) \cdot d(X, W)} = \frac{\beta}{\beta - 1}.$$

As we assume $\beta \geq 3/2$, the desired bound follows.

2. $\alpha\text{-PC}(W) \leq \frac{3}{2} \cdot d(X, W)$. Since W is a Plurality winner, at least $\frac{n}{m}$ agents rank W first. Each such agent prefers W to X , and for each agent $i \in WX$ by Lemma 5 we have $d(i, X) \geq \frac{1}{2} \cdot d(W, X)$. As $\alpha \geq \frac{m-1}{m}$, it follows that $\alpha\text{-PC}(X) \geq \frac{1}{2} \cdot d(X, W)$. As $\alpha\text{-PC}(W) \leq \frac{3}{2} \cdot d(X, W)$, the desired distortion bound is immediate. ■

Theorem 31 *The distortion of the Copeland rule with respect to the $\alpha\text{-PC}$ objective with $\alpha \in [\frac{1}{2}, 1)$ is at most 5, and this bound is tight.*

Proof. For the upper bound, we use the fact that the Copeland rule outputs a subset of the uncovered set. We then observe that the proof of Theorem 27 works verbatim, with median replaced by $\alpha\text{-PC}$. To prove that this bound is tight, let $k = \lceil \frac{1}{1-\alpha} \rceil$ and consider a profile over $\{X, Y, W\}$ where two agents rank the alternatives as $W \succ Y \succ X$, $k + 2$ agents rank the alternatives as $Y \succ X \succ W$ and $k + 1$ agents rank the alternatives as $X \succ W \succ Y$. Then W defeats Y , Y defeats X , and X defeats W , so W is a winner under the Copeland rule. Now let the underlying metric be as shown in Figure 6, where all voters that do not rank X last are located at distance $1 - \epsilon$ from X . (The distances not shown in the figure can be chosen to be consistent with the metric and the preference profile.)




Figure 6: For the Copeland rule for a given $\alpha \in [\frac{1}{2}, 1)$, we can choose a value of $k > 1$ so that $\alpha\text{-PC}(W) = 5 + \epsilon$ and $\alpha\text{-PC}(X) = 1 - \epsilon$. Thus the distortion with respect to the $\alpha\text{-PC}$ objective approaches 5 as $\epsilon \rightarrow 0$.

Now, consider $\alpha\text{-PC}(W)$. Since $\alpha \geq \frac{1}{2}$, we know that $\frac{k+1}{2k+5} < \alpha$, and so less than α fraction of the agents are within distance $3 + \epsilon$ of W . On the other hand, $\frac{2k+3}{2k+5} = 1 - \frac{2}{2k+5} \geq 1 - \frac{1}{k} \geq \alpha$, so at least α fraction of the agents are within distance $5 + \epsilon$ of W and hence $\alpha\text{-PC}(W) = 5 + \epsilon$. Similarly, $\alpha\text{-PC}(X) = 1 - \epsilon$. Thus, the distortion with respect to the $\alpha\text{-PC}$ objective approaches 5 as $\epsilon \rightarrow 0$. ■

Together with the lower bound from Theorem 28, Theorems 30 and 31 show that, for distortion with respect to the $\alpha\text{-PC}$ objective, Plurality is optimal for $\frac{m-1}{m} \leq \alpha \leq 1$, whereas the Copeland rule is optimal for $\frac{1}{2} \leq \alpha < \frac{2}{3}$.

5 Conclusion and Future Directions

We analyzed the distortion of many common voting rules in the setting where the agent costs form a metric space. We showed that despite the process of winner determination having absolutely no extra information about the underlying metric space except the induced ordinal agent preferences, rules that select from the uncovered set achieve a small constant-factor approximation to the optimal alternative (and in fact, for median objective function they achieve the best approximation to an optimal alternative that a deterministic voting rule can ever hope to achieve); this class of

rules includes the well-known Copeland rule. We have also established that all positional scoring rules have distortion that grows with the number of alternatives, with popular rules in this family having distortion that is linear or almost linear in the number of alternatives, and that STV exhibits acceptable performance with respect to distortion.

Nevertheless, some important open questions remain. Foremost among them is whether there exists a voting rule that beats Copeland, and maybe achieves the best possible distortion of 3. In the conference version of this paper, we conjectured that this may be the case for Ranked Pairs; however, this conjecture was recently disproved by Goel *et al.* [2017]. Further, while there are some promising results on the distortion of randomized voting rules [Anshelevich and Postl, 2017; Feldman *et al.*, 2016; Gross *et al.*, 2017], our understanding of the randomized setting is far from complete. It particular, it would be interesting to see if, by adding a small amount of randomization to Copeland and Ranked Pairs, we can improve their distortion properties. Finally, we observe that for the median objective function many of the entries in Table 1 are ∞ . A natural research direction is to develop a more fine-grained approach that would enable us to compare different voting rules in this model.

6 Acknowledgments

This work began as a project in a class taught by Lirong Xia, and we would like to thank him for many interesting discussions on the topic of social choice. This work was partially supported by NSF awards CCF-1101495, CNS-1218374, by an ERC Starting Grant ACCORD (639945), and by the Foundation for Polish Science within the Homing programme (Project title: "Normative Comparison of Multiwinner Election Rules").

References

- E. Anshelevich and J. Postl. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research*, 58:797–827, 2017.
- K. J. Arrow. *Social Choice and Individual Values*. Wiley, New York, 2nd edition, 1963.
- D. Black. On the rationale of group decision-making. *The Journal of Political Economy*, pages 23–34, 1948.
- A. Bogomolnaia and J.-F. Laslier. Euclidean preferences. *Journal of Mathematical Economics*, 43(2):87–98, 2007.
- A. Borodin and R. El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 1998.
- C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- S. Branzei, I. Caragiannis, J. Morgenstern, and A. Procaccia. How bad is selfish voting? In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pages 138–144, 2013.
- D. Campbell and S. Nitzan. Social compromise and social metrics. *Social Choice and Welfare*, 3(1):1–16, 1986.
- I. Caragiannis and A. D. Procaccia. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9):1655–1671, 2011.
- D. Chakrabarty and C. Swamy. Welfare maximization and truthfulness in mechanism design with ordinal preferences. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pages 105–120, 2014.
- J.-A.-N. de Caritat, Marquis de Condorcet. *Essai sur l'Application de L'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. 1785. Facsimile reprint of original published in Paris, 1972, by the Imprimerie Royale.
- V. Conitzer and T. Sandholm. Common voting rules as maximum likelihood estimators. In *Proceedings of the 21st Conference in Uncertainty in Artificial Intelligence*, pages 145–152, 2005.
- V. Conitzer, M. Rognlie, and L. Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 109–115, 2009.

- O. A. Davis and M. J. Hinich. A mathematical model of preference formation in a democratic society. In J. L. Bernd, editor, *Mathematical Applications in Political Science II*, pages 175–208. Southern Methodist University Press, 1966.
- E. Elkind and A. Slinko. Rationalizations of voting rules. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- E. Elkind, P. Faliszewski, and A. Slinko. Distance rationalization of voting rules. *Social Choice and Welfare*, 45(2):345–377, 2015.
- E. Elkind, P. Faliszewski, J. F. Laslier, P. Skowron, A. Slinko, and N. Talmon. What do multiwinner voting rules do? An experiment over the two-dimensional euclidean domain. In *Proceedings of the 31st Conference on Artificial Intelligence*, pages 494–501, 2017.
- J. M. Enelow and M. J. Hinich. *The spatial theory of voting: An introduction*. CUP Archive, 1984.
- J. M. Enelow and M. J. Hinich. *Advances in the spatial theory of voting*. Cambridge University Press, 1990.
- M. Feldman, A. Fiat, and I. Golomb. On voting and facility location. In *Proceedings of the 17th ACM Conference on Economics and Computation*, pages 269–286, 2016.
- A. Filos-Ratsikas and P. B. Miltersen. Truthful approximations to range voting. In *Proceedings of the 10th International Conference on Web and Internet Economics*, pages 175–188, 2013.
- P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on applied Mathematics*, 33(3):469–489, 1977.
- J. S. Gans and M. Smart. Majority voting with single-crossing preferences. *Journal of Public Economics*, 59(2):219–237, 1996.
- A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–601, 1973.
- A. Goel, A. K. Krishnaswamy, and K. Munagala. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 18th ACM Conference on Economics and Computation*, pages 287–304, 2017.
- S. Gross, E. Anshelevich, and L. Xia. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In *Proceedings of the 31st Conference on Artificial Intelligence*, 2017.
- J. C. Harsanyi. *Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility*. Springer, 1976.
- G. Laffond, J.-F. Laslier, and M. Le Breton. Condorcet choice correspondences: A set-theoretical comparison. *Mathematical Social Sciences*, 30(1):23–35, 1995.
- K. O. May. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, 20(4):680–684, 1952.
- R. D. McKelvey and P. C. Ordeshook. A decade of experimental research on spatial models of elections and committees. In J. M. Enelow and M. J. Hinich, editors, *Advances in the spatial theory of voting*, pages 99–144. Cambridge University Press, 1990.
- S. Merrill and B. Grofman. *A unified theory of voting: Directional and proximity spatial models*. Cambridge University Press, 1999.
- H. Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.
- Y.-K. Ng. A case for happiness, cardinalism, and interpersonal comparability. *The Economic Journal*, 107(445):1848–1858, 1997.
- J. Oren and B. Lucier. Online (budgeted) social choice. In *Proceedings of 28th AAAI Conference on Artificial Intelligence*, pages 1456–1462, 2014.
- C. R. Plott. A notion of equilibrium and its possibility under majority rule. *The American Economic Review*, 57(4):787–806, 1967.

- A. D. Procaccia and J. S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of 10th International Workshop on Cooperative Information Agents*, pages 317–331. Springer, 2006.
- J. E. Roemer. *Theories of distributive justice*. Harvard University Press, 1998.
- A. Saporiti. Strategy-proofness and single-crossing. *Theoretical Economics*, 4(2):127–163, 2009.
- M. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.
- N. Schofield. *The spatial model of politics*. Routledge, 2007.
- H. P. Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.
- H. P. Young. Condorcet’s theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988.
- W. S. Zwicker. Introduction to the theory of voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 2. Cambridge University Press, 2015.