Approximating Optimal Social Choice under Metric Preferences

Elliot Anshelevich   Onkar Bhardwaj

John Postl
Rensselaer Polytechnic Institute
110 8th Street, Troy NY 12180
eanshel@cs.rpi.edu, bhardo@rpi.edu, postlj@rpi.edu

November 23, 2014

Abstract

We examine the quality of social choice mechanisms using a *utilitarian* view, in which all of the agents have costs for each of the possible alternatives. While these underlying costs determine what the optimal alternative is, they may be unknown to the social choice mechanism; instead the mechanism must decide on a good alternative based only on the ordinal preferences of the agents which are induced by the underlying costs. Due to its limited information, such a social choice mechanism cannot simply select the alternative that minimizes the total social cost (or minimizes some other objective function). Thus, we seek to bound the *distortion*: the worst-case ratio between the social cost of the alternative selected and the optimal alternative. Distortion measures how good a mechanism is at approximating the 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 the agent costs form a *metric space*, which is a natural assumption in many settings. We quantify the distortion of many well-known social choice mechanisms. We show that for both total social cost and median agent cost, many positional scoring rules have large distortion, while on the other hand *Copeland* and similar mechanisms perform optimally or near-optimally, always obtaining a distortion of at most 5. We also give lower bounds on the distortion that could be obtained by *any* deterministic social choice mechanism, and extend our results on median agent cost to more general objective functions.

1 Introduction

Social choice theory deals with aggregating agent preferences over a set of alternatives into a collective decision via a *social choice mechanism*. The social choice mechanism takes as input the preferences of agents, which are usually total orderings over the set of alternatives, and typically outputs a single alternative as the winner. It is natural to now ask about the quality of different social choice mechanisms; to do this one needs to define what it means for a chosen alternative to be "good" or to accurately represent the consensus of the agent preferences. A popular way of achieving this is to define criteria or axioms for social choice mechanisms, which guarantee that the alternatives selected by these mechanisms satisfy desirable properties (see Related Work). Another common approach in fields like welfare economics and algorithmic mechanism design, and which we follow in this paper, is to use a *utilitarian* view [4]. Instead of assuming that agents only have ordinal preferences over the alternatives, this approach assumes that every agent has (possibly latent or implicit) utility or cost values over the alternatives. These values are cardinal, and represent how happy the agent is with each alternative. The quality of an alternative can then be defined simply as the sum (or some other objective function) of the utility received by each agent for that alternative. Thus the best, or optimal, alternative is simply the one that maximizes the total social welfare (or minimizes total cost), as measured by the total utility received by the agents.

Utilitarian approach has recently received renewed attention in the study of social choice [23, 6, 4, 5]. Indeed, as argued in [4], although not all social choice problems are amenable to the utilitarian approach (especially the ones where it is unnatural to assume that agent utilities or costs can be compared) there are many real-life settings which fit the utilitarian view. For example, in recommender systems and many similar domains from mechanism design and

	Sum	Median
Plurality	$2m-1$	∞
Borda	$2m-1$	∞
k-approval	$2n-1$	∞
Veto	$2n-1$	∞
Copeland	5	5
Uncovered Set	5	5
Lower Bound	3	5

Table 1: The worst-case distortion of various social choice mechanisms for both the sum and the median objective functions. All of the above bounds are provably tight, meaning that we provide example instances where the social choice function cannot achieve a better bound. The lower bounds of 3 and 5 are for any deterministic social choice functions.

e-commerce, the computational agents typically use real-valued rather than ordinal utilities (see Related Work and [4]).

If the social choice mechanism knew exactly what utilities the agents receive from each alternative, then it could simply pick the alternative maximizing social welfare directly. An important point here, however, is that while we assume that some underlying utility structure exists, it is unreasonable to assume that we (or even the agents themselves) know exactly what it is. As discussed in [4], it is often difficult for agents to determine their exact cardinal utilities, and most social choice mechanisms thus take only the ordinal preference orderings of the agents as input, even when latent utilities exist. Thus, a social choice function will not simply output the alternative that maximizes global utility, but instead may choose another alternative, since it only has access to ordinal preferences. As a result, one can think of a social choice function as an approximation algorithm which attempts to choose the best possible alternative (maximize social welfare or minimize social cost), but only has access to limited information (ordinal preferences instead of cardinal utilities). To denote the approximation factor of a social choice function, [23] introduced the term *distortion* which we will continue to use, although we will define it in terms of social cost instead of social welfare. Informally, the distortion of a social choice function is the worst-case ratio of the social cost of the alternative selected by the social choice function over the cost of the optimal alternative.

In this work, we are primarily interested in determining the quality of outcomes chosen by social choice mechanisms, as measured by their distortion. We prove bounds on the distortion of many well-known social choice functions for both the sum and median objective functions. Our results show that while the distortion is high for some mechanisms, the distortion of many important social choice functions is bounded by a small constant, assuming that the preferences of the agents are *spatial*. Specifically, we assume that the costs of agents for various alternatives form an arbitrary *metric space*. Such *metric costs* have a very natural interpretation – in the context of voting, as described in the classic Downsian proximity model [17], we can think of the cost experienced by voter i due to candidate j being elected as the distance between i and j's beliefs in some high-dimensional space, as the number of issues they disagree on, etc. Such spatial preferences have been extensively studied (see Related Work), although usually the metric space is assumed to be simple, e.g., Euclidean with only one or two dimensions. In contrast, we make no assumptions about the metric space, other than the fact that it is a metric space. To see how general our metric assumption is, note that, unlike many common assumptions on spacial agent preferences, our metric assumption does not restrict the set of possible ordinal preferences in any way (see Proposition 1 and discussion before it).

1.1 Our Contributions

In this work, we bound the worst-case distortion of many well-known social choice functions. In other words, we show how closely social choice functions approximate the optimal alternative when they are given only the ordinal preference orderings, instead of the underlying metric costs which generate these preferences.[1] We consider two general objective functions to quantify the quality of alternatives, and give distortion results for both. 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.

[1]This is assuming that agents submit their true ordinal preferences. We leave questions about non-truthful agents as future work.

This function is very natural, and is the most common measure of social cost. Our other objective function defines the quality of an alternative as the median of agent costs for that alternative: this captures the objective that the best alternative is the one in which the cost of the median voter is minimized, instead of the average voter.

Most of our results are summarized in Table 1. First, we consider how well *any* social choice function could do when it only knows the ordinal preferences, but is supposed to approximate the social optimum. We show that no deterministic social choice mechanism can have worst-case distortion 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 social choice rules which meet this lower bound? Are there rules which obtain the minimum possible distortion?

We begin with the bad news: for common positional scoring rules such as plurality, Borda, k-approval, and veto, we prove that the worst case distortion can be high: either $2m-1$ or $2n-1$ where m is the number of alternatives and n is the number of agents/voters. There is good news as well, however. For the Copeland social choice rule, we prove that the distortion is always at most 5. This means that, although the Copeland social choice mechanism knows nothing about the metric costs other than the ordinal preferences induced by them, and cannot possibly find the true optimal alternative, it nevertheless *always* selects an alternative whose quality is only a factor of 5 away from optimal! Moreover, due to our lower bound, no deterministic mechanism can do better than Copeland for the median objective, and no deterministic mechanism can do much better than Copeland for the sum objective, because the distortion lower bound for any deterministic mechanism is 3.

While this bound of 5 holds for both the sum and median objectives, different techniques are required to prove it for the two cases. In fact, this bound holds not just for Copeland, but for similar voting rules as well, such as uncovered set [19]. Since Copeland does not perform as well as the lower bound for the sum objective, we also analyze the distortion of the ranked pairs mechanism. We show that it performs even better than Copeland, but only when certain conditions on the agent preference profiles are satisfied (see Theorem 8).

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 mechanisms changes with x, and establish that Copeland remains the mechanism with the best possible distortion for most values of x.

1.2 Related Work

The focus of much of the existing literature in social choice theory is the design and analysis of social choice functions with respect to various normative criteria (See for example [11, 1, 8, 16]). Results like Arrow's impossibility theorem and Gibbard-Satterthwaite theorem demonstrate non-existence of social choice functions satisfying certain desirable criteria, and additional assumptions must be made in order to circumvent these results (e.g., [18, 12, 20]).

In this work, we instead adopt a *utilitarian* view of social choice as described in the Introduction. Social choice with utilitarian viewpoint has its advocates in welfare economics [24, 21] and has recently received attention from the AI community [23, 6, 4]. The utilitarian approach has also been investigated in recommender systems [13], information extraction [26], etc. While assuming that agent utilities can be compared does not make sense for all settings [14], it is nevertheless reasonable in many applications of interest: see [4] for much more discussion on this subject.

Distortion as a measure of performance of a social choice function in utilitarian settings was introduced first in [23] and later used in [4]. In both these works, the worst-case distortion of social choice functions was shown to be unbounded or very high. In our work, however, we show that considering agent costs that form an *unknown metric* immediately greatly reduces the distortion of many mechanisms, from unbounded to only a small constant. [6] use an analogous notion of distortion to analyze the worst-case distortion of embeddings into voting rules: these embeddings are functions that take as input an agent's utility function and determine which alternative the agent should select. Apart from the classic normative criteria, other papers have also used related interpretations of what makes a good social choice function, such as distance rationalizability [9], rank approximation [7], and dynamic price of anarchy [5].

In our paper, we assume that agents have spatial preferences resulting from metric agent costs. Spatial preferences and utility theory in the context of voting have a strong tradition in social choice and political science [10, 17]. Common assumptions include single-peaked preferences [18, 27] and single-crossing preferences [25, 20]; often preferences are assumed to be one-dimensional, while we consider metrics with arbitrary dimension.

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. This is similar, for example, to the 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 [3, 22].

2 Preliminaries

Social Choice with Ordinal Preferences. Let $N = \{1, 2, \ldots, n\}$ be the set of agents, and let $\mathcal{A} = \{A_1, A_2, \ldots, A_m\}$ be the set of alternatives. 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 in ranking σ_i. We call the vector $\sigma = (\sigma_1, \ldots, \sigma_n) \in \mathcal{S}^n$ a *preference profile*. We say that an alternative X *pairwise defeats* Y if $|\{i \in N : X \succ_i Y\}| > \frac{n}{2}$.

Once we are given a preference profile, we want to aggregate the preferences of the agents and select a single alternative as the winner. A *social choice function* $f : \mathcal{S}^n \to \mathcal{A}$ is a mapping from a preference profile to an alternative. Some well-known social choice functions which we consider in this paper are as follows.

- **Positional scoring rules.** We are given a scoring vector $\vec{s} = (s_1, s_2, \ldots, s_m)$ with $s_1 \geq s_2 \geq \cdots \geq s_m$. If an agent ranks an alternative in position l, then the alternative receives s_l points. The total score $s(X, \sigma)$ of an alternative X for a preference profile σ is the total number of points that X receives. The positional scoring rule is $f(\sigma) = \arg\max_{X \in \mathcal{A}} s(X, \sigma)$; that is, it selects the alternative with the highest total score. Many well-known voting rules can be thought of as positional scoring rules, for example:
  - **Plurality:** $\vec{s} = (1, 0, \ldots, 0)$
  - **Veto:** $\vec{s} = (1, 1, \ldots, 1, 0)$
  - **Borda:** $\vec{s} = (m-1, m-2, \ldots, 1, 0)$
  - k**-approval** $(1 < k < m)$: $\vec{s} = (\underbrace{1, 1, \ldots, 1}_{k}, 0, \ldots, 0)$
- **Copeland:** The score of an alternative X is $|\{Y \in \mathcal{A} : X \text{ pairwise defeats } Y\}|$. The alternative with the highest score, i.e., the alternative with the largest number of pairwise victories, is the winner.
- **Ranked pairs:** Construct a graph G in the following manner. Let every alternative be a node in G. For every pair of alternatives X, Y, let $w(X, Y) = |\{i \in N : X \succ_i Y\}|$. Sort these $w(X, Y)$ values in non-increasing order and iterate over them. For each $w(X, Y)$ value, add the directed edge (X, Y) to G if it won't create a cycle, and do nothing otherwise. The winner is the source node of the resulting directed acyclic graph.

Cardinal Metric Costs. In our work we take the utilitarian view, and study the case when the ordinal preferences σ are in fact a result of the underlying cardinal agent costs. Formally, we assume that there exists an arbitrary metric $d : (N \cup \mathcal{A})^2 \to \mathbb{R}_{\geq 0}$ on the set of agents and alternatives (or more generally a *semi-metric*, since we allow agents to be identical and have $d(i, j) = 0$). Here $d(i, X)$ is the cost incurred by agent i of alternative X being 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. Formally, we say that σ is *consistent* with d if $\forall i \in N, \forall X, Y \in \mathcal{A}$, if $d(i, X) < d(i, Y)$, then $X \succ_i Y$. 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 ($p(d)$ may include several preference profiles if the agent costs have ties). Similarly, we define $p^{-1}(\sigma)$ to be the set of metrics such that $\sigma \in p(d)$.

When making additional assumptions on how the preference rankings of the agents are generated, the set of possible preference profiles may become restricted. For example, if we restrict agents to one-dimensional single-peaked preferences, or to single-crossing preferences, then preference profiles with the Condorcet paradox can no longer be realized [2, 12, 25]. 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.

Proposition 1 *For every preference profile σ, there exists a metric d such that σ is consistent with d.*

Proof. Consider the metric space $(\mathbb{R}^m, d)$, where d is the Manhattan distance $d(x, y) = \sum_{k=1}^{m} |x_k - y_k|$, for all $x, y \in \mathbb{R}^m$. We will define a mapping $g : N \cup \mathcal{A} \to \mathbb{R}^m$. Each alternative is mapped to a distinct point on the simplex; A_1 is mapped to $(1, 0, \ldots, 0)$, A_2 is mapped to $(0, 1, \ldots, 0)$, etc. For every agent i, we map it to a vector v and set

$v_k = \frac{m-l+1}{m}$, where l denotes the position of A_k in agent i's preference ranking. We can now define a metric d' on $N \cup \mathcal{A}$, where for all $x, y \in N \cup \mathcal{A}$, $d'(x,y) = d(g(x), g(y))$.

We observe that $(g(N \cup A), d)$ is a metric space, because it is a metric subspace of $(\mathbb{R}^m, d)$. We conclude that $(N \cup \mathcal{A}, d')$ is a metric space as well.

Next, we claim that σ is consistent with d'. We observe that for any alternative X ranked in position l by i,

$$\begin{aligned} d(i,X) &= \sum_{Y \neq X} |0 - \frac{m - \sigma_i(Y) + 1}{m}| + \left|1 - \frac{m-l+1}{m}\right| \\ &= \sum_{k=1, k\neq l}^{m} \left[\frac{m-k+1}{m}\right] + \frac{l-1}{m} \\ &= \sum_{k=1}^{m} \frac{m-k+1}{m} - \frac{m-l+1}{m} + \frac{l-1}{m} \\ &= \frac{1}{m}\sum_{k=1}^{m}(m-k+1) + \frac{-m+2l-2}{m} \\ &= m+1 - \frac{1}{m}\sum_{k=1}^{m} k + \frac{-m+2l-2}{m} \\ &= m+1 - \frac{m(m+1)}{2m} + \frac{-m+2l-2}{m} \end{aligned}$$

Thus, $d(i,X)$ strictly increases with l, and so for any agent i, $d'(i,X) < d'(i,Y)$ iff $\sigma_i(X) < \sigma_i(Y)$. ∎

Social Cost and Distortion. We measure the quality of each alternative using the costs incurred by all the agents when this alternative is chosen. We use two different notions of social cost. First, we study the sum objective function, defined as $\mathrm{SC}_{\sum}(X,d) = \sum_{i\in N} d(i,X)$; this is the most common notion of social cost. We also study the median objective function, $\mathrm{SC}_{\mathrm{med}}(X,d) = \mathrm{med}_{i\in N}(d(i,X))$. As described in the Introduction, we can view social choice mechanisms in our setting as attempting to find the optimal alternative (one that minimizes social cost), but 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: the only way one can determine the optimal alternative while only having access to σ is if there is a single alternative that is the top preference for *all* agents. In fact, we cannot even eliminate any alternative from consideration of being optimal, except in trivial cases.

Proposition 2 *For any preference profile σ and alternative X, there exists a metric $d \in p^{-1}(\sigma)$ such that X is optimal with respect to the social cost function $\mathrm{SC}_{\sum}(X,d)$, except when there exists an alternative Y such that for all $i \in N$, $Y \succ_i X$.*

Proof. Let $\mathcal{A}'$ denote the set of alternatives X that do not have an alternative Y such that every agent prefers Y to X. We will define a metric d such that any arbitrary alternative $W \in \mathcal{A}'$ is optimal. For each agent i, set $d(i,X) = \frac{1}{2}$ for all $X \succeq_i W$ and $d(i,X) = 1$ for all $X \prec_i W$. For every pair of alternatives X, Y, set $d(X,Y) = 1$. For every pair of agents i, j, set $d(i,j) = 1$. It can be easily verified that d is a metric.

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

$$\begin{aligned} \mathrm{SC}_{\sum}(X,d) - \mathrm{SC}_{\sum}(W,d) &= \sum_{i\in N} d(i,X) - \sum_{i\in N} d(i,W) \\ &= \sum_{i: X \succ_i W} (d(i,X) - d(i,W)) \\ &+ \sum_{i: X \prec_i W} (d(i,X) - d(i,W)) \\ &= 0 + \frac{1}{2}|\{i \in N : X \prec_i W\}| \\ &\geq \frac{1}{2}. \end{aligned}$$

Thus, we conclude that W is optimal. ■

Since it is impossible to compute the optimal alternative using only ordinal preferences, we would like to determine how well the aforementioned social choice functions select alternatives based on their social costs, despite only being given the preference profiles. In particular, we would like to quantify how the social choice functions perform in the worst-case. To do this, we use the notion of *distortion* from [23, 4], defined as follows.

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

In other words, the distortion of a social choice mechanism f on a profile σ is the worst-case ratio between the social cost of $f(\sigma)$, and the social cost of the true optimum alternative. The worst-case is taken over all metrics d which may have induced σ, since the social choice function does not and cannot know which of these metrics is the true one.

Notation. We introduce the following notation, which will be used throughout our proofs, for describing agents with particular preference rankings.

$$\begin{aligned}
XY &= \{i \in N : X \succ_i Y\} \\
XYZ &= \{i \in N : X \succ_i Y \succ_i Z\} \\
X* &= \{i \in N : X \succ_i Y \text{ for all } Y \neq X\} \\
XY* &= \{i \in N : X \succ_i Y \succ_i Z \text{ for all } Z \neq X, Y\} \\
*X &= \{i \in N : Y \succ_i X \text{ for all } Y \neq X\} \\
*XY &= \{i \in N : Z \succ_i X \succ_i Y \text{ for all } Z \neq X, Y\}
\end{aligned}$$

3 Distortion of Total Agent Cost

In this section, we study the sum objective function, which measures the quality of an alternative to be the total agent cost when this alternative is chosen. We prove tight upper bounds for distortion of several well-known social choice functions. Our main result in this section is that the Copeland voting mechanism (as well as several others) exhibit 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).

Before proceeding with showing upper bounds on possible distortion, we ask the question: how well can *any* social choice function perform? The following simple theorem tells us that we cannot possibly hope to approximate the optimal alternative within a factor better than 3.

Theorem 3 *No (deterministic) social choice function 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 social choice function picks W as the winner. The underlying metric can be as follows. All $n/2$ agents who prefer X 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 $\text{SC}_{\sum}(X, d) = \sum_{i \in N} d(i, X) = (1 + \epsilon) \cdot n/2$ and $\text{SC}_{\sum}(W, d) = \sum_{i \in N} d(i, W) = 2 \cdot n/2 + (1 - \epsilon) \cdot n/2$. Thus, the distortion approaches 3 as $\epsilon \to 0$. ■

In fact, it is easy to show that for only two alternatives, any social choice function that picks the winner preferred by the majority of agents has a distortion of 3, i.e., all such social choice functions achieve the optimal distortion bound for two alternatives. This is a corollary of Theorem 4. Unfortunately, as the number of agents and candidates becomes large, the distortion of many social choice mechanisms increases linearly.

Theorem 4 *For plurality and Borda social choice functions, the distortion is at most $2m - 1$; for k-approval and veto it is at most $2n - 1$. Furthermore, these bounds are tight, i.e., they are achieved exactly in some instances.*

Before proving this theorem, we first prove some useful lemmas. Our first lemma consists of bounds on the costs of agents incurred by alternatives that are used repeatedly throughout all of our proofs. The first two bounds provide lower bounds for the cost of an alternative, which will be used to lower bound the cost of the optimal alternative. The last bound is an upper bound which will be used to bound the cost of the winning alternative for agents who prefer the optimal alternative over the winning alternative. For agents i who prefer the winning alternative W over the optimal alternative X, we can simply use $d(i, X) \leq d(i, W)$. Since we assume the costs of agents come from a 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:*

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

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

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

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

(2). We have $d(i, X) \geq d(X, W) - d(i, W)$, by the triangle inequality. Since W is preferred over Y, we get $d(i, X) \geq d(X, W) - d(i, Y)$. Adding this with $d(i, X) \geq d(i, Y) - d(X, Y)$ gives us the desired result.

(3). Let Z be any alternative such that $W \succeq_i Z$ (note that there always exists such a 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 any $Z \preceq_i W$, we conclude that (3) holds. ■

Our next lemma parameterizes the distortion by the number of agents that prefer the winning alternative W over the optimal alternative X.

Lemma 6 *For any instance σ and social choice function f, $\operatorname{dist}_{\sum}(f, \sigma) \leq 1 + \frac{2(n-|WX|)}{|WX|}$, where $W = f(\sigma)$ is the winning alternative and X is the optimal alternative.*

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 $d(i, W) \leq d(i, X)$, but for an agent $i \in |XW|$, the best we can do is use the triangle inequality to obtain that $d(i, W) \leq d(i, X) + d(X, W)$. Applying these bounds to the distortion allows us to derive

$$\begin{aligned}
\frac{\mathrm{SC}_{\sum}(W, d)}{\mathrm{SC}_{\sum}(X, d)} &= \frac{\sum_{i \in N} d(i, W)}{\sum_{i \in N} d(i, X)} \\
&= \frac{\sum\limits_{i \in WX} d(i, W) + \sum\limits_{i \in XW} d(i, W)}{\sum\limits_{i \in N} d(i, X)} \\
&\leq \frac{\sum\limits_{i \in WX} d(i, X) + \sum\limits_{i \in XW} (d(i, X) + d(X, W))}{\sum\limits_{i \in N} d(i, X)} \\
&= 1 + \frac{\sum_{i \in XW} d(X, W)}{\sum_{i \in N} d(i, X)} \\
&= 1 + \frac{|XW| \cdot d(X, W)}{\sum_{i \in N} d(i, X)} \\
&= 1 + \frac{(n - |WX|) \cdot d(X, W)}{\sum_{i \in N} d(i, X)}
\end{aligned}$$

Finally, we want to lower bound the cost of the optimal alternative X. It follows from Lemma 5 that $d(i, X) \geq d(X, W)/2$ for an agent i with $W \succ_i X$. Thus we have $\sum_{i \in N} d(i, X) \geq |WX| \cdot d(X, W)/2$. Applying this inequality to our last equation gives us the desired result. ■

Since the bound from Lemma 6 decreases with $|WX|$, bounding the smallest possible $|WX|$ for a social choice function will give us an upper bound on the worst-case distortion. This is the technique we use to prove worst-case distortion results for some positional scoring rules, as we see in the proof of Theorem 4.

Proof of Theorem 4. Let W denote the winning alternative, and let X denote an optimal alternative. Since the bound from Lemma 6 decreases with $|WX|$, bounding the smallest possible $|WX|$ for each scoring function will give us an upper bound on the worst-case distortion.

For plurality, W clearly must be the top preference of at least n/m agents, i.e., $|WX| \geq \frac{n}{m}$. For Borda, we claim that $|WX| \geq \frac{n}{m}$. Suppose, by way of contradiction, that $|WX| < n/m$. Then the agents that prefer X to W contribute at least $\frac{(m-1)n}{m}$ more to the score of X than they do to the score of W. The agents that prefer W to X can contribute at most $|WX| \cdot (m-1)$ more to the score of W than X if all of them rank W first and X last. However, $|WX| \cdot (m-1) < \frac{(m-1)n}{m}$, implying that X has a higher score than the winner W which is a contradiction.

For both k-approval and Veto, W being the winner implies that W has strictly higher score than X, implying at least one agent prefers W to X, i.e., $|WX| \geq 1$.

Applying these $|WX|$ lower bounds to Lemma 6 gives us the desired distortion results for the above positional scoring rules.

Figure 1 shows tight examples of these bounds for plurality (left-hand side) and Borda (right-hand side), achieving the distortion bound of $2m-1$ for $m=4$ alternatives (all of these examples can easily be generalized to any $m \geq 3$). In these examples, W is the winner and X is the optimal alternative. For plurality, the dummy alternatives Y and Z split the top preferences of agents preferring X over W. For Borda, these dummy alternatives artificially inflate the score of W. For k-approval and veto, the tight examples are the same as Borda's, except there is only one agent with $W \succ Y \succ Z \succ X$, while the remaining agents have preference $X \succ W \succ Y \succ Z$. □

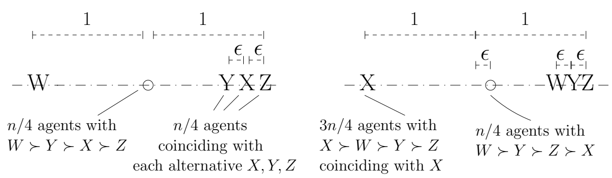

Figure 1: Examples showing tightness of distortion bound for plurality (left) and Borda (right) with $m=4$. Here W is picked as winner and X is the optimal alternative. As $\epsilon \to 0$, distortion approaches $2m-1$. Adding an extra agent coinciding with the center point makes W the unique winner.

Theorem 4 implies that the distortion for plurality and Borda is unbounded in the number of candidates, and for k-approval and veto it is unbounded in the number of voters. Informally, this is because the optimal alternative can be preferred over the eventual winner by a relatively large fraction of the agents, and yet still lose. We now consider several social choice functions that escape this predicament, resulting in significantly better performance. We state the results for these social choice functions now, but prove them later.

Theorem 7 *For the Copeland social choice function, distortion is always ≤ 5, and this bound is tight.*

Remark: *In fact, the result for distortion being at most 5 holds whenever for any other alternative Z, the winner W either pairwise defeats Z or there exists an alternative Y whom W pairwise defeats and Y pairwise defeats Z. This precisely corresponds to the notion of W being a member of the uncovered set [19]. Thus the distortion is at most 5 for several notions of tournament winners other than Copeland such as the winner being selected from minimal covering set, bipartisan set, banks set, tournament equilibrium set, etc., as all these sets are a subset of the uncovered set [15].*

Recall that no social choice function can have distortion less than 3. Thus, Copeland is nearly optimal with a distortion of at most 5. We can show that the ranked pairs mechanism achieves the best possible distortion bound, but only in the special case when the majority graph (directed graph in which a link (X, Y) denotes that X pairwise defeats Y) has small circumference (i.e., maximum cycle size). In general, we conjecture that the worst-case distortion is ≤ 3 for ranked pairs, but our current techniques cannot obtain this bound.

Theorem 8 *The distortion of ranked pairs is ≤ 3, as long as the majority graph has circumference ≤ 4.*

Lemma 9 *Let $\vec{v} \in \mathbb{R}^m_{\geq 0}$ with $v_1 \geq v_2 \geq \cdots \geq v_m$. Let $\alpha, \beta \in \mathbb{R}^m$. If $\forall k = 1, \ldots, m, \sum_{i=1}^k \alpha_i \geq \sum_{i=1}^k \beta_i$, then $\sum_{i=1}^m \alpha_i v_i \geq \sum_{i=1}^m \beta_i v_i$*

Proof. For this proof, we simply use $\sum_{i=1}^k \alpha_i \geq \sum_{i=1}^k \beta_i$ repeatedly, once for every value of k.

$$
\begin{aligned}
\sum_{i=1}^m \alpha_i v_i &= (\alpha_1 + \beta_1 - \beta_1)\, v_1 + \sum_{i=2}^m \alpha_i v_i \\
&\geq \beta_1 v_1 + (\alpha_1 - \beta_1)\, v_2 + \sum_{i=2}^m \alpha_i v_i \\
&= \beta_1 v_1 + (\alpha_1 + \alpha_2 - \beta_1 + \beta_2 - \beta_2)\, v_2 + \sum_{i=3}^m \alpha_i v_i \\
&\geq \beta_1 v_1 + \beta_2 v_2 + (\alpha_1 + \alpha_2 - \beta_1 - \beta_2)\, v_3 + \sum_{i=3}^m \alpha_i v_i \\
&\ \ \vdots \\
&\geq \sum_{i=1}^k \beta_i v_i + \left(\sum_{i=1}^k (\alpha_i - \beta i)\right) v_k + \sum_{i=k+1}^m \alpha_i v_i \\
&\geq \sum_{i=1}^k \beta_i v_i + \left(\sum_{i=1}^k (\alpha_i - \beta i)\right) v_{k+1} + \sum_{i=k+1}^m \alpha_i v_i \\
&= \sum_{i=1}^{k+1} \beta_i v_i + \left(\sum_{i=1}^{k+1} (\alpha_i - \beta i)\right) v_{k+1} + \sum_{i=k+2}^m \alpha_i v_i \\
&\ \ \vdots \\
&\geq \sum_{i=1}^{m-1} \beta_i v_i + \left(\sum_{i=1}^{m-1} (\alpha_i - \beta i)\right) v_{m-1} + \alpha_m v_m \\
&\geq \sum_{i=1}^{m-1} \beta_i v_i + \left(\sum_{i=1}^{m-1} (\alpha_i - \beta i)\right) v_m + \alpha_m v_m \\
&\geq \sum_{i=1}^{m} \beta_i v_i.
\end{aligned}
$$

■

The next lemma allows us to instantly obtain a distortion upper bound from lower bounds of a specific form for the optimal alternative X.

Lemma 10 *If for any metric d, preferences σ induced by d, and alternatives X, W, we have that $\sum_{i \in N} d(i, X) \geq \frac{1}{\gamma} \sum_{i \in XW} \min_{Z \preceq_i W}(d(X, Z))$ for some $1 \geq \gamma$, then $\mathrm{SC}_{\sum}(W, d)/\,\mathrm{SC}_{\sum}(X, d) \leq 1 + \gamma$.*

Proof. We begin by using $d(i, W) \leq d(i, X)$ for $i \in WX$ and Inequality 3 from Lemma 5 for $i \in XW$ to upper bound the cost of the alternative W.

$$
\begin{aligned}
\frac{\mathrm{SC}_{\sum}(W, d)}{\mathrm{SC}_{\sum}(X, d)} &= \frac{\sum_{i \in N} d(i, W)}{\sum_{i \in N} d(i, X)} \\
&\leq \frac{\sum\limits_{i \in N} d(i, X) + \sum\limits_{i \in XW} (\min_{Z \preceq_i W} d(X, Z))}{\sum_{i \in N} d(i, X)} \\
&= 1 + \frac{\sum_{i \in XW} \min_{Z \preceq_i W}(d(X, Z))}{\sum_{i \in N} d(i, X)}
\end{aligned}
$$

Using the assumption that $\sum_{i\in N} d(i,X) \geq \frac{1}{\gamma}\sum_{i\in XW} \min_{Z\preceq_i W}(d(X,Z))$ and applying it to our previous equation gives us the desired result. ■

Proof of Theorem 7. Let W be the winning alternative using Copeland, and X be the optimal alternative. For any instance where W pairwise defeats X, the distortion is at most 3 by Lemma 6.

Thus, let us now consider an instance where W does not pairwise defeat X, i.e., $|XW| \geq n/2 \geq |WX|$. For Copeland, we know that this implies there exists a alternative Y such that W pairwise defeats Y and Y pairwise defeats X (Moulin 1986). Thus we have $|WY| \geq n/2 \geq |YW|$ and $|YX| \geq n/2 \geq |XY|$. We consider the cases when $d(X,Y) \geq d(X,W)$ and $d(X,Y) < d(X,W)$ separately.

Case (i). Suppose that $d(X,Y) \geq d(X,W)$. We know that at least $n/2$ agents prefer Y over X. Hence, each of these $n/2$ agents must contribute at least $d(X,Y)/2 \geq d(X,W)/2$ to the social cost of X, by Lemma 6. This observation is all we need to obtain a distortion upper bound of 5. Formally, we see that

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

We obtain the desired result by applying Lemma 10 to this $\sum_{i\in N} d(i,X)$ lower bound.

Case (ii). Suppose instead that $d(X,W) > d(X,Y)$. Unlike our previous case in which we just considered agents in YX while lower bounding $\sum_{i\in N} d(i,X)$, we must consider the sets WX, YX, and XWY. This is because the agents in these sets are the only ones guaranteed to have cost incurred by alternative X. Using bounds from Lemma 5, we begin by observing that

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

Next, we would like to apply Lemma 10, but we need our lower bound to be in terms of agents in XW rather than $WX \cup XWY \cup YXW$. Thus, we make the following observations that follow from basic set theory and properties of Copeland,

$$\begin{aligned}
|WX| + |XWY| &= |WYX| + |WXY| + |YWX| + |XWY|\\
&\geq |WY|\\
&\geq |YW|
\end{aligned}$$

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

With these two inequalities, we can now apply Lemma 9 to our previous lower bound for $\sum_{i \in N} d(i, X)$

$$
\begin{aligned}
\sum_{i \in N} d(i, X) &\geq \frac{1}{2}\left(|WX| + |XWY|\right) \cdot d(X, W) \\
&\quad + \frac{1}{2}\left(|YXW| - |XWY|\right) \cdot d(X, Y) \\
&\geq \frac{1}{4}\left(|YXW| + |XYW|\right) \cdot d(X, W) \\
&\quad + \frac{1}{4}|XWY| \cdot d(X, Y) \\
&\geq \frac{1}{4} \sum_{i \in XW} \min_{Y' \preceq_i W} (d(X, Y')).
\end{aligned}
$$

Finally, we have our lower bound for $\sum_{i \in N} d(i, X)$ in the form necessary to apply Lemma 10, which gives us the desired upper bound of 5 on distortion.

Now let us prove the tightness of this bound. Suppose there are only three alternatives W, X, Y. Let there be $\frac{n}{2} - 1$ agents corresponding to each of the preference rankings $Y \succ X \succ W$ and $X \succ W \succ Y$. Let the remaining two agents have preference ranking $W \succ Y \succ X$. Observe that W pairwise defeats Y, Y pairwise defeats X, and X pairwise defeats W. Thus every alternative pairwise defeats exactly one other alternative. In this situation, suppose Copeland chooses W as the winning alternative. Now let the underlying metric be as shown in Figure 2. (The distances not shown in the figure can be chosen to be consistent with the metric and the preference profile.) Now we have:

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

Thus as $n \to \infty$ and $\epsilon \to 0$, we get instances with distortion arbitrarily close to 5. □

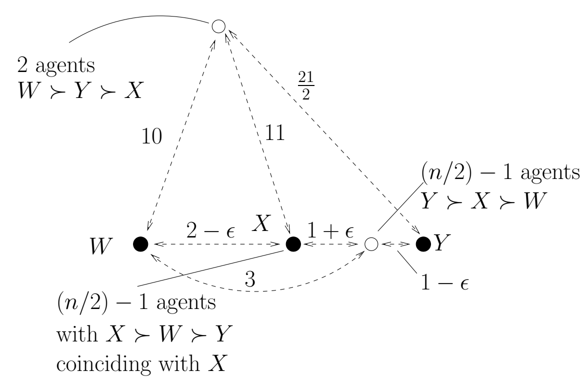

Figure 2: W being picked as the winning alternative leads to worst-case distortion arbitrarily close to 5 as $\epsilon \to 0$ and $n \to \infty$.

Proof of Theorem 8. As usual, let W be the alternative chosen by our social choice mechanism, and let X be the optimal alternative. For any instance where W pairwise defeats X then the distortion is at most 3 by Lemma 6.

Let us consider instances in which W does not pairwise defeat X. That is, $|XW| \geq n/2 \geq |WX|$. Let G be the graph generated by the ranked pairs mechanism. Since W is the source of G, then there must be at least one path in G from W to X. Furthermore, at least one of these paths has the property that the edge weights are $\geq |XW|$ for all edges in the path, because otherwise the edge (X, W) would have been added to the graph constructed by the ranked pairs mechanism, which contradicts W being the winner. Let P be this path. Then, P is also a subpath of the majority graph, and P together with (X, W) form a cycle in the majority graph. By our assumption on the circumference, 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 lengths 2 is similar and simpler. Then, if P consists of alternative W, Y, Z, X, then W defeats alternative Y pairwise, Y defeats Z pairwise, and Z defeats X pairwise with $|WY|, |YZ|, |ZX| \geq |XW|$.

Case (i). Suppose that $d(X, Z) \geq d(X, W)$. As with the first case of Copeland, we need only consider the agents $i \in ZX$. Then we know that at least $|ZX| \geq |XW|$ agents each contribute at least $d(X, Z)/2 \geq \frac{d(X,W)}{2}$ to the social cost of X, by Lemma 5. This is all we need to obtain an upper bound of 3. Formally, we see that

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

We obtain the desired result by applying Lemma 10.

Case (ii). Suppose that $d(X, W) \geq d(X, Z) \geq d(X, Y)$. We begin by making observations about $d(i, X)$ for the agents in the sets WX and ZXW, because they contribute significantly to the cost of X. Using Lemma 5, we observe that

$$\begin{aligned}
\sum_{i \in N} d(i, X) &\geq \sum_{i \in WX \cup ZXW} d(i, X) \\
&\geq \frac{1}{2} |WX| \cdot d(X, W) + \frac{1}{2} |ZXW| \cdot d(X, Z)
\end{aligned}$$

In order to apply Lemma 10, we need to lower bound $\sum_{i \in N} d(i, X)$ in terms of agents in XW. Thus, we make the following observations about how the cardinality of the sets we considered relate to $|XW|$ (recall our notation introduced at the beginning of the supplemental materials).

$$\begin{aligned}
|WX| &= n - |XW| \\
&\geq n - |WY| \\
&= |YW| \\
&\geq |*W| \\
|WX| + |ZXW| &\geq |ZX| \\
&\geq |XW| \\
&\geq |*W| + |*WZ| + |XWY|
\end{aligned}$$

Now we can combine these results with our previous lower bound for $\sum_{i \in N} d(i, X)$, and use Lemma 9 with $\alpha_1 = |WX|$, $\alpha_2 = |ZXW|$, $\beta_1 = |*W|$, and $\beta_2 = |*WZ| + |XWY|$. This allows us to derive that

$$\sum_{i \in N} d(i, X) \geq \frac{1}{2} |WX| \cdot d(X, W) + \frac{1}{2} |ZXW| \cdot d(X, Z)$$

$$
\begin{aligned}
&\geq \frac{1}{2}|*W| \cdot d(X,W) + \frac{1}{2}|*WZ| \cdot d(X,Z) \\
&\quad + \frac{1}{2}|XWY| \cdot d(X,Y) \\
&\geq \frac{1}{2}\sum_{i \in XW} \min_{Z' \preceq_i W}(d(X,Z'))
\end{aligned}
$$

We can apply Lemma 10 to obtain the desired result.

Case (iii). Suppose that $d(X,W) \geq d(X,Y) \geq d(X,Z)$. This case is similar to our previous case, except we need to consider more sets of agents. In particular, we require the following bounds

$$|WX| \geq |*W|$$

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

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

As we have done previously, we will lower bound $d(i,X)$ for agents in certain sets using Lemma 5. In this case, we use the sets WX, YXW, $XYZ \cap XW$, and $ZXW \cap ZXY$; these are all disjoint. Then we will use the inequalities we just derived above so we can apply Lemma 9.

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

By Lemma 10, we conclude that distortion ≤ 3.

Case (iv). For the final case, suppose that $d(X,Y) \geq d(X,W) \geq d(X,Z)$. This case is almost identical to our previous one. We reuse the following inequalities:

$$|WX| + |YXW| + |XYZ \cap XW| \geq |*W| + |*WY|$$

$$|WX| + |YXW| + |ZXW \cap ZXY| \geq |*W| + |*WY| + |XWZ|.$$

Using these inequalities, Lemma 5, and Lemma 9,

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

Lemma 10 gives us the desired result.

This concludes the proof that ranked pairs always yields a distortion of at most 3 when the majority graph has no large cycles. The case in which the path is length 2 is similar (and simpler). The case in which the path is length 1 is simply the case where W defeats X pairwise. Finally, we give an example to show that this bound is tight. Suppose that $d(X,W) = d(X,Y) = d(X,Z) = 1$. We construct a preference profile as follows:

There are $\frac{n}{4}$ agents i with $X \succ_i W \succ_i Y \succ_i Z$ such that $d(i,X) = 0, d(i,W) = d(i,Y) = d(i,Z) = 1$.

There are $\frac{n}{4}$ agents i with $W \succ_i Y \succ_i Z \succ_i X$ such that $d(i,W) = d(i,Y) = d(i,Z) = d(i,X) = \frac{1}{2}$.

There are $\frac{n}{4}$ agents i with $Y \succ_i Z \succ_i X \succ_i W$ such that $d(i,Y) = d(i,Z) = d(i,X) = \frac{1}{2}, d(i,W) = \frac{3}{2}$.

There are $\frac{n}{4}$ agents i with $Z \succ_i X \succ_i W \succ_i Y$ such that $d(i,Z) = d(i,X) = \frac{1}{2}, d(i,W) = d(i,Y) = \frac{3}{2}$.

We allow the remaining distances to be arbitrary, as long as the triangle inequality is obeyed to ensure that d is a metric.

We observe that W can be the winner, depending on how ties are resolved. $\mathrm{SC}_{\sum}(X) = \frac{3n}{4} \cdot \frac{1}{2} = \frac{3n}{8}$ and $\mathrm{SC}_{\sum}(W) = \frac{n}{4} \cdot \frac{1}{2} + \frac{n}{4} \cdot 1 + \frac{n}{2} \cdot \frac{3}{2} = \frac{9n}{8}$. This gives us the desired lower bound of 3. □

4 Distortion of Median Agent Cost

In this section we study the distortion of social choice functions as measured by the median agent cost. We define the median social cost of an alternative $\mathrm{SC}_{\mathrm{med}}(Y,d) = \mathrm{med}_{i \in N}\, d(i,Y)$ to be the median of the list of distances of all the agents to the alternative Y. If n is even, we define it to be the $(\frac{n}{2}+1)^{th}$ smallest value of the distances. As a shorthand, we will refer to this as $\mathrm{med}(Y)$ when the cost metric d is fixed. The distortion of a social choice function is now $\mathrm{dist}_{\mathrm{med}}$ as defined in Section 2. We begin by establishing lower bounds on the distortion achieved by any deterministic social choice function; this bound is higher than in the sum case, but first note the following trivial lemma, due to the triangle inequality.

Lemma 11 *For any two alternatives Y and Z, we have* $\mathrm{med}(Z) \leq \mathrm{med}(Y) + d(Y,Z)$.

Proof. All the agents with $d(i,Y) \leq \mathrm{med}(Y)$ have $d(i,Z) \leq d(i,Y) + d(Y,Z) \leq \mathrm{med}(Y) + d(Y,Z)$. Hence the result follows. ■

Theorem 12 *With $m = 2$ alternatives, any social choice function that picks the alternative preferred by the majority has distortion ≤ 3 for the median objective.*

Proof. Let X be an optimal alternative. Suppose that majority of the agents prefer another alternative W to X. We split the analysis into two cases: $\mathrm{med}(W) \geq 3 \cdot d(X,W)/2$ and $\mathrm{med}(W) \leq 3 \cdot d(X,W)/2$.

Let us consider the first case: let $\mathrm{med}(W) = \beta \cdot d(X,W)$ where $\beta \geq 3/2$. Applying Lemma 11, we get $\mathrm{med}(X) \geq (\beta - 1) \cdot d(X,W)$. Hence the distortion is at most $\beta/(\beta-1) \leq 3$ as we are assuming $\beta \geq 3/2$.

Now consider the other case where $\text{med}(W) \leq \frac{3}{2} \cdot d(X, W)$. If $\text{med}(X) \geq d(X, W)/2$ then we the desired distortion bound immediately follows. Hence assume that $\text{med}(X) < d(X, W)/2$. This together with Equation 5.1 implies strictly more than half of the agents prefer X over W — note that if there are even number of agents, we use $\lceil (\frac{n}{2} + 1) \rceil^{th}$ smallest value of $d(i, X)$ as $\text{med}(X)$ — thus contradicting that W is an alternative preferred by the majority of agents. ■

The proof of Theorem 12 works verbatim to give us the following corollary for any instance of a social choice function and any number of alternatives:

Corollary 13 *Any social choice function f and metric d such that the winning alternative chosen by f pairwise defeats an optimal alternative has distortion ≤ 3 for the median objective.*

When we do not restrict the number of alternatives to 2, we see that the worst-case distortion increases from 3 to 5 for all social choice functions.

Theorem 14 *No (deterministic) social choice function has worst-case distortion less than 5 for the median objective.*

Proof. Suppose there are only three alternatives W, X, Y. Let there be $n/3$ 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 social choice function picks W as the winner. Consider an underlying metric as shown in Figure 3. (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 $\epsilon \to 0$. ■

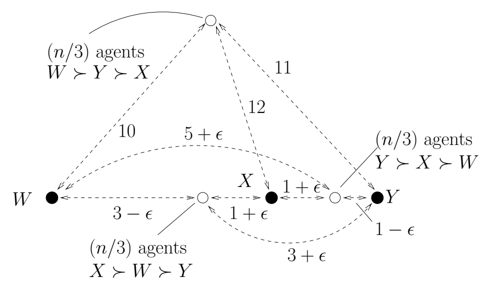

Figure 3: With median objective, W being picked as the winning alternative leads to worst-case distortion arbitrarily close to 5 as $\epsilon \to 0$.

As for the sum objective function, the distortion of the common positional scoring rules remains high for the median objective. In fact, it becomes unbounded for any $m > 2$ number of alternatives.

Theorem 15 *Plurality, Borda, k-approval, and veto have unbounded distortion for any number of alternatives $m > 2$.*

Proof. Consider the instances shown in Figure 1. In both these instances, $\text{med}(X) \leq \epsilon$ and $\text{med}(W) \geq 2 - 2\epsilon$. With plurality, W is a winning alternative in the instance on the left side, thus resulting in unbounded distortion. With Borda, veto, and k-approval (with k=2,3), W is a winning alternative in the instance on the right side, thus resulting in unbounded distortion. These instances can easily be generalized for any $m > 2$ number of alternatives. ■

Now we show that the Copeland social choice function achieves the optimal distortion bound: due to the lower bound in Theorem 14 no deterministic rule can have better median distortion than Copeland. Note that this result holds also for several other notions of tournament winners mentioned in the concluding remark in Section 3.

Theorem 16 *For the Copeland social choice function, median distortion is always ≤ 5, and this bound is tight.*

Proof. Let W denote the winning alternative and let X be an optimal alternative. Whenever W pairwise defeats X, then we know that the distortion can be at most 3 (Corollary 13). Hence let us assume that X pairwise defeats W. In this case, we know from (Moulin 1986) that there exists an intermediate candidate Y such that W pairwise defeats Y and Y pairwise defeats W. Now we split the analysis into two cases: $\text{med}(W) > \frac{5}{4} \cdot d(X, W)$ and $\text{med}(W) \leq \frac{5}{4} \cdot d(X, W)$.

Let us consider the first case: Suppose that $\text{med}(W) = \beta \cdot d(X, W)$ with $\beta > 5/4$, i.e., at least half of the agents have $d(i, W) \geq \beta \cdot d(X, W)$. Using $d(i, X) \geq d(i, W) - d(X, W)$, we conclude that all these agents must have $d(i, X) \geq (\beta - 1) \cdot d(X, W)$. Hence the distortion in this case would be at most $\frac{\beta \cdot d(X,W)}{(\beta-1) \cdot d(X,W)} \leq 5$ as $\beta > 5/4$.

Now consider the case where $\text{med}(W) \leq \frac{5}{4} \cdot d(X, W)$. We claim that the following holds:

$$\text{med}(X) \geq \max\left(\frac{d(X, Y)}{2}, \frac{d(X, W) - d(X, Y)}{2}\right) \tag{4}$$

The first term in the above inequality is due to half of the agents preferring Y over X (thus each of them having $d(i, X) \geq d(X, Y)/2$). The second term results from at least half of the agents preferring W over Y, and as shown in Lemma 5, this implies that each such agent has $d(i, X) \geq \frac{1}{2} \cdot (d(X, W) - d(X, Y))$. Hence we get

$$\frac{\text{SC}_{\text{med}}(W, d)}{\text{SC}_{\text{med}}(X, d)} \leq \frac{\frac{5 \cdot d(X,W)}{4}}{\max\left(\frac{d(X,Y)}{2}, \frac{d(X,W) - d(X,Y)}{2}\right)}$$

The above term achieves its maximum value when $d(X, W) = 2 \cdot d(X, Y)$, thus giving us the desired distortion bound. This bound is tight due to Theorem 14. ∎

4.1 Generalizing Median: Percentile Distortion

Instead of considering the happiness of the median voter or agent, it also makes sense to consider the happiness of the 25th or 75th percentile. We can generalize the median objective function $\text{med}(Y)$ above by using percentiles as follows. Let α-PC(Y) be the value from the set $\{d(i, Y) : i \in N\}$ below which lie an α fraction of the values. Thus α-PC(Y) with $\alpha = 1/2$ is the same as $\text{med}(Y)$. The distortion with α-PC is defined analogously to Section 2.

For various ranges of α, we now give lower bounds on the distortion that any social choice function must have in Theorem 17, and then give social choice functions that always achieve these bounds in Theorems 19 and 20.

Theorem 17 *For any deterministic social choice function:*

(a) For $\alpha \in [\frac{2}{3}, 1)$, worst-case α-PC distortion is at least 3.

(b) For $\alpha \in [\frac{1}{2}, \frac{2}{3})$, worst-case α-PC distortion is at least 5.

(c) For $\alpha \in [0, \frac{1}{2})$, worst-case distortion is unbounded.

Proof. Part (a): Suppose half of the agents have preference $W \succ X$ and the others have $X \succ W$. Without loss of generality, assume that the given social choice function picks W as the winner. Consider the instance shown in Figure 4. Here for $\alpha \geq \frac{1}{2}$, we have α-PC$(X) = 1 + \epsilon$ and α-PC$(W) = 3 + \epsilon$, making distortion arbitrarily close to 3 as $\epsilon \to 0$.

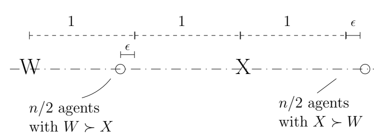

Figure 4: As $\epsilon \to 0$, the winning alternative W has distortion arbitrarily close to 3 with α-PC with $\alpha \geq \frac{1}{2}$.

Part (b): The same argument as in Theorem 14 establishes a lower bound on worst-case distortion of at least 5 for α-PC agent cost with $\frac{2}{3} > \alpha \geq \frac{1}{2}$.

Part (c): For any $\alpha < 0.5$, by choosing n suitably large, we can make α-PC$(X) = 0$ and α-PC$(W) = 1 - \epsilon$ in the instance described in the proof of Theorem 3. Hence the result follows. ∎

We now give upper bounds of the distortion of plurality and Copeland to show they are optimal for certain values of α. First, we give the following trivial lemma which will be used in the proof of these upper bounds.

Lemma 18 *For any two alternatives Y and Z, we have α-PC$(Z) \leq \alpha$-PC$(Y) + d(Y, Z)$.*

Proof. All the agents with $d(i, Y) \leq \alpha$-PC(Y) have $d(i, Z) \leq d(i, Y) + d(Y, Z) \leq \alpha$-PC$(Y) + d(Y, Z)$. Hence the result follows. ■

Theorem 19 *For the plurality social choice function, distortion is always* ≤ 3 *for* α*-PC objective with* $\alpha \geq \frac{m-1}{m}$.

Proof. This proof is almost identical to the proof of Theorem 12. Let us first consider the case when α-PC$(W) = \beta \cdot d(X, W)$ where $\beta \geq 3/2$. Applying Lemma 18, we get α-PC$(X) \geq (\beta - 1) \cdot d(X, W)$. Hence the distortion is at most $\beta/(\beta - 1) \leq 3$ as we are assuming $\beta \geq 3/2$.

Now consider the other case where α-PC$(W) \leq 3 \cdot d(X, W)/2$. Now if α-PC$(X) \geq d(X, W)/2$ then we already get the desired distortion bound. Hence assume that α-PC$(X, W) < d(X, W)/2$. This together with Equation 1 means strictly more than n/m agents prefer X over W — note that if $\alpha \cdot n$ is an integer then we use $\lceil(\alpha n + 1)\rceil^{th}$ smallest value of $d(i, X)$ as α-PC(X) — thus contradicting that W is an alternative preferred by at least n/m agents. ■

Theorem 20 *For the Copeland social choice function, distortion is always* ≤ 5 *for* α*-PC objective with* $\frac{1}{2} \leq \alpha < 1$, *and this bound is tight.*

Proof. To prove the upper bound of 5, the proof of Theorem 16 works verbatim except with median replaced by α-PC. To prove the tightness of this bound, suppose that there are 2, $k + 2$, and $k + 1$ agents corresponding to each of the preference rankings $W \succ Y \succ X$, $Y \succ X \succ W$ and $X \succ W \succ Y$. For any $k \geq 1$, W pairwise defeats Y, Y pairwise defeats X, and X pairwise defeats W. Thus for any $k \geq 1$, every alternative pairwise defeats exactly one other alternative. In this situation, suppose Copeland chooses W as the winning alternative. Now let the underlying metric be as shown in Figure 5. (The distances not shown in the figure can be chosen to be consistent with the metric and the preference profile.) For any $\alpha \in [\frac{1}{2}, 1)$, choose a large enough k so that $\frac{2k+3}{2k+5} \geq \alpha$.

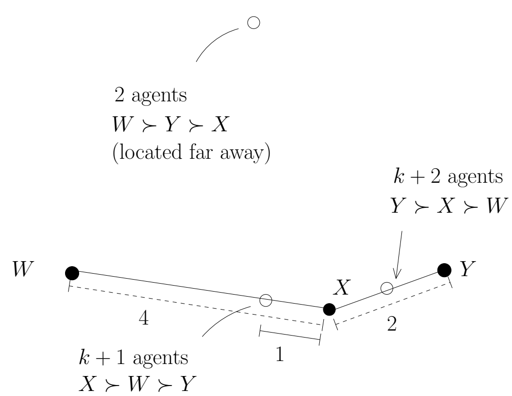

Figure 5: For Copeland and with α-PC as the objective, by choosing a suitable parameter $k \geq 1$ for any $\alpha \in [\frac{1}{2}, 1)$, we get α-PC$(W) = 5 + \epsilon$ and α-PC$(X) = 1 + \epsilon$, thus the distortion approaches 5 as $\epsilon \to 0$.

Now, consider $\text{med}(W)$. Since $\alpha \geq 1/2$, 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, at least α fraction of the agents are within distance $5 + \epsilon$ of W due to our choice of k, and so $\text{med}(W) = 5 + \epsilon$. Similarly, $\text{med}(X) = 1 + \epsilon$. Thus, the distortion approaches 5 as $\epsilon \to 0$. ■

Together with the lower bound from Theorem 17, this shows that for $\alpha \geq \frac{m-1}{m}$, no deterministic rule can have better worst-case distortion than plurality, whereas Copeland achieves the optimal worst-case distortion for $\frac{1}{2} \leq \alpha < \frac{2}{3}$.

5 Conclusion and Future Directions

We analyzed the distortion of many common social choice mechanisms 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, mechanisms like Copeland achieve a small constant-factor approximation to the optimal candidate (and in fact, for median objective function they achieve the best approximation to an optimal candidate that a deterministic mechanism can ever hope to achieve).

Nevertheless, some important open questions remain. Foremost among them is the question of a social choice rule which beats Copeland, and maybe achieves the best possible distortion of 3. While we showed some weaker results for the ranked pairs mechanism, we believe there is a good chance that it performs even better than we anticipate, and actually guarantees a distortion of 3 for all instances, not just the ones with small graph circumference. Exploring the space of randomized mechanisms could also be very fruitful. Randomized mechanisms still cannot get arbitrarily close to the optimal alternative (we can prove a lower bound of 2 on distortion instead of 3), but a small amount of randomization added to Copeland and ranked pairs has a chance to greatly improve their distortion properties.

6 Acknowledgments

This work began as a project in a class taught by Lirong Xia, and we would like to thank him greatly for many interesting discussions on the topic of social choice. This work was partially supported by NSF awards CCF-1101495 and CNS-1218374.

References

[1] John Bartholdi III, Craig A Tovey, and Michael A Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and welfare*, 6(2):157–165, 1989.

[2] Duncan Black. On the rationale of group decision-making. *The Journal of Political Economy*, pages 23–34, 1948.

[3] Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 1998.

[4] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 197–214. ACM, 2012.

[5] Simina Branzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel Procaccia. How bad is selfish voting? In *AAAI*, 2013.

[6] Ioannis Caragiannis and Ariel D Procaccia. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9):1655–1671, 2011.

[7] Deeparnab Chakrabarty and Chaitanya 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.

[8] Vincent Conitzer and Tuomas Sandholm. Vote elicitation: Complexity and strategy-proofness. In *AAAI/IAAI*, pages 392–397, 2002.

[9] Edith Elkind, Piotr Faliszewski, and Arkadii Slinko. On distance rationalizability of some voting rules. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 108–117, 2009.

[10] James M Enelow and Melvin J Hinich. *The spatial theory of voting: An introduction*. CUP Archive, 1984.

[11] Piotr Faliszewski and Ariel D Procaccia. Ai's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

[12] Joshua S Gans and Michael Smart. Majority voting with single-crossing preferences. *Journal of Public Economics*, 59(2):219–237, 1996.

[13] Sumit Ghosh, Manisha Mundhe, Karina Hernandez, and Sandip Sen. Voting for movies: the anatomy of a recommender system. In *Proceedings of the third annual conference on Autonomous Agents*, pages 434–435. ACM, 1999.

[14] John C Harsanyi. *Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility*. Springer, 1976.

[15] Gilbert Laffond, Jean Francois Laslier, and Michel Le Breton. Condorcet choice correspondences: A set-theoretical comparison. *Mathematical Social Sciences*, 30(1):23–35, 1995.

[16] Jérôme Lang and Lirong Xia. Sequential composition of voting rules in multi-issue domains. *Mathematical social sciences*, 57(3):304–324, 2009.

[17] Samuel Merrill and Bernard Grofman. *A unified theory of voting: Directional and proximity spatial models*. Cambridge University Press, 1999.

[18] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.

[19] Herve Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.

[20] Roger B Myerson. *Fundamentals of social choice theory*. 1996.

[21] Yew-Kwang Ng. A case for happiness, cardinalism, and interpersonal comparability. *The Economic Journal*, 107(445):1848–1858, 1997.

[22] Joel Oren and Brendan Lucier. Online (budgeted) social choice. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.

[23] Ariel D Procaccia and Jeffrey S Rosenschein. The distortion of cardinal preferences in voting. In *Cooperative Information Agents X*, pages 317–331. Springer, 2006.

[24] John E Roemer. *Theories of distributive justice*. Harvard University Press, 1998.

[25] Alejandro Saporiti. Strategy-proofness and single-crossing. *Theoretical Economics*, 4(2):127–163, 2009.

[26] Georgios Sigletos, Georgios Paliouras, Constantine D Spyropoulos, and Michalis Hatzopoulos. Combining information extraction systems using voting and stacked generalization. *The Journal of Machine Learning Research*, 6:1751–1782, 2005.

[27] Xin Sui, Alex Francois-Nienaber, and Craig Boutilier. Multi-dimensional single-peaked consistency and its approximations. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 375–382. AAAI Press, 2013.