

Relating Metric Distortion and Fairness of Social Choice Rules

Extended Abstract*

Ashish Goel, Reyna Hulett, and Anilesh K. Krishnaswamy

Stanford University

{ashishg,rmhulett,anilesh}@stanford.edu

ABSTRACT

One way of evaluating social choice (voting) rules is through a utilitarian distortion framework. In this model, we assume that agents submit full rankings over the alternatives, and these rankings are generated from underlying, but unknown, quantitative costs. The *distortion* of a social choice rule is then the ratio of the total social cost of the chosen alternative to the optimal social cost of any alternative; since the true costs are unknown, we consider the worst-case distortion over all possible underlying costs. Analogously, we can consider the worst-case *fairness ratio* of a social choice rule by comparing a useful notion of fairness (based on approximate majorization) for the chosen alternative to that of the optimal alternative. With an additional metric assumption – that the costs equal the agent-alternative distances in some metric space – it is known that the Copeland rule achieves both a distortion and fairness ratio of at most 5. For other rules, only bounds on the distortion are known, e.g., the popular Single Transferable Vote (STV) rule has distortion $O(\log m)$, where m is the number of alternatives. We prove that the distinct notions of distortion and fairness ratio are in fact closely linked – within an additive factor of 2 for any voting rule – and thus STV also achieves an $O(\log m)$ fairness ratio. We further extend the notions of distortion and fairness ratio to social choice rules choosing a *set* of alternatives. By relating the distortion of single-winner rules to multiple-winner rules, we establish that Recursive Copeland achieves a distortion of 5 and a fairness ratio of at most 7 for choosing a set of alternatives.

KEYWORDS

Social Choice, Metric Distortion, Fairness

ACM Reference Format:

Ashish Goel, Reyna Hulett, and Anilesh K. Krishnaswamy. 2018. Relating Metric Distortion and Fairness of Social Choice Rules: Extended Abstract. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Social choice theory studies the aggregation of agent preferences into a single collective decision via a social choice rule. Often these

preferences are expressed as total orderings over a set of possible alternatives, and a social choice rule maps any instance of preferences to one or more alternatives. The traditional approach to evaluating the quality of social choice rules has been a normative, axiomatic one. A great deal of work has been done to propose various axioms or properties that social choice rules ought to satisfy – such as strategy-proofness or the majority winner criterion – and evaluate different rules by which criteria they meet [4]. Unfortunately, even small sets of natural axioms may be impossible to satisfy simultaneously. For instance, the celebrated Gibbard-Satterthwaite Theorem rules out the existence of strategy-proof, non-dictatorial social choice rules with more than two alternatives [20, 32]. Thus we must either accept social choice rules which fail to satisfy some natural properties, or make assumptions which limit the possible agent preferences but permit strategy-proof rules [6, 29].

An alternative to the axiomatic approach, which has lately received much attention in the field of computational social choice [9, 12, 31], is to adopt a utilitarian view – agents express their ordinal preferences by ranking the alternatives, but have latent *cardinal* preferences over the alternatives. In particular, much work [2, 3, 33] has been done on the *metric distortion* problem [2]. Under this model, agents and alternatives are assumed to lie in an unknown, arbitrary metric space, and an agent's cost for an alternative is given by the distance between the two. Social choice rules are viewed as approximation algorithms trying to choose the alternative with the lowest social cost, given access to only the ordinal preferences of agents. Similar to the competitive ratio of online approximation algorithms, the quantity of interest here is the worst-case value (over all possible underlying costs) of the *distortion* – the ratio of the social cost of the chosen alternative to that of the optimal alternative, chosen omnisciently [2]. In this setting, the best known positive result for deterministic social rules is that the distortion of the Copeland rule is at most 5 [2]. Another recent result establishes that the distortion of the Single Transferable Vote (STV) rule, which is widely used in practice, is $O(\log m)$ [33], where m is the number of alternatives.

Under the metric distortion approach, another important question is to quantify how “fair” choosing a particular alternative is. Since the costs incurred could vary widely among the agents, minimizing the total cost might not be the only useful objective. For instance, imagine there are two agents and two alternatives, and the costs for the agents are $\{1, 3\}$ for the first alternative, and $\{2, 2\}$ for the second. It is natural to expect that under any reasonable notion of fairness the second is more desirable than the first. Although many notions of fairness exist, not all are applicable in this setting, and some cannot even be bounded for any social choice rule [2, 21]. We use the *fairness ratio* [21], which is a simultaneous bound, over all k , on the distortion of the social cost objective given by the sum

*For want of space, we have pushed all our proofs to the Appendix

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference'17, July 2017, Washington, DC, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

of the k largest agent costs. This definition of fairness is based on the notion of approximate majorization [23], which generalizes the various notions of fairness that have been studied in the context of routing, bandwidth allocation and load balancing problems [24–26]. For instance, it includes as special cases both total cost minimization (utilitarianism), when $k = N$ (where N is the total number of agents), and max-min fairness (egalitarianism), when $k = 1$. A bound on the fairness ratio also translates to an approximation guarantee on a wide class of convex objective functions [21]. For example, a constant factor bound on the fairness ratio implies the same bound on the distortion of social cost objectives such as the geometric mean, or any l_p norm of the agent costs [21]. Thus, this interesting notion of fairness captures a wide range of desirable properties.

It is known that Copeland achieves a fairness ratio of at most 5, thereby yielding a constant-factor approximation for a large class of convex objectives [21]. Since it also achieves a distortion of at most 5, a natural question arises as to whether the distortion and fairness ratio are inherently connected. For many popular voting rules – such as scoring rules and STV – bounds are known on the distortion but not the fairness ratio. Could the distortion bounds for these rules extend to the fairness ratio as well, as is the case with Copeland? In particular, does any bound on distortion directly imply a bound on the fairness ratio? The primary aim of this paper is to answer this question.

The majority of previous work on the metric distortion problem looks only at social choice rules that choose a single alternative and not those that choose a set of alternatives. To address this lacuna, we also study the problem of establishing bounds on the distortion and fairness ratio for the case of choosing a set of winners, given the size of the desired set. Applications of multi-winner elections are quite diverse such as proportional representation in parliament [8, 28], selecting a diverse committee [13] (e.g. locations of fire stations), offering a selection of movies to passengers on a flight [27], and shortlisting candidates for a job interview [7, 15]. We will focus on settings where a set of resources has to be chosen to serve a community of voters. Some examples of these setting are choosing overlay networks in the Internet [1], or a set of public projects to implement given a fixed budget [10, 22]. We will also assume in such settings that the preferences are additive – the total cost for an agent is just the sum of her costs over the chosen set of alternatives. It is fairly straightforward to see that a repeated application of Copeland to choose a set of winners yields a bound of 5 on the distortion. In fact, we establish that any single-winner rule can be applied recursively to obtain a multiple-winner rule with the same bound on distortion. However, it was not known if a constant-factor bound on the fairness ratio is possible in this case. We answer this question in the positive by extending the distortion-fairness relationship to rules choosing a set of winners, thus establishing that the set chosen by a repeated application of Copeland achieves a bound of at most 7 on the fairness ratio.

1.1 Our Results

Our primary focus in this paper will be to quantify the relationship between the distortion and fairness ratio for social choice rules that choose a single winner, which also leads to new results for

any rules where bounds are known on the distortion but not the fairness ratio. We then extend these results to rules that choose a set of winners, and characterize upper bounds on distortion and fairness in this setting. We also provide a lower bound on the gap between the distortion and fairness ratio.

Note: All our proofs are in the appendix due to space constraints.

1.1.1 Distortion and Fairness for Single-Winner Rules. Given the ordinal preferences of the agents, the distortion of an alternative is the worst case value (over all possible underlying metrics) of the ratio between the sum of agent costs for it, and the sum of the agent costs for the optimal alternative. For studying fairness under the metric distortion framework, we instead consider the ratio of the sum of the k largest agent costs for an alternative to the k largest costs for the optimal alternative, and take the maximum value among these ratios over all possible values of k . The overall fairness ratio is the worst case value of this quantity over all possible underlying metrics. Since for $k = N$ we get the sum of costs objective, the fairness ratio is at least as large as distortion.

Our main result is that the distortion and fairness ratio are closely related. For any given ordinal preferences, and any alternative, we show that the fairness ratio cannot be much larger than the distortion.

RESULT 1 (THEOREM 3.2). *Given the ordinal preferences of the agents, the fairness ratio of any alternative is at most 2 more than its distortion.*

We also present instances where the “obvious” winner for any reasonable social choice rule (which is also the distortion minimizing alternative) has a distortion-fairness gap approaching 2.

RESULT 2 (THEOREM 3.3). *There exist instances where the gap between the distortion and fairness ratio of the distortion-optimal alternative approaches 2.*

In fact, because the above results hold for any set of preferences and any metric space, they also imply that for any social choice rule, the worst-case fairness ratio is at most 2 more than its worst-case distortion. This implies novel bounds on the fairness ratio of several social choice rules, including STV, for which only bounds on distortion were previously known [33].

RESULT 3 (COROLLARY 3.4). *The fairness ratio of STV is $O(\log m)$ and $\Omega(\sqrt{\log m})$.*

1.1.2 Distortion and Fairness for Multiple-Winner Rules. We also consider the case where a social choice rule selects a set of winners (of a given size). Here, the cost of an agent for any set of alternatives is the sum of her costs over the alternatives in the set, and the social cost is the sum of the individual agent costs, as before.

It is straightforward to see from previous results [2] that choosing the desired set by a repeated application of the Copeland rule achieves a worst-case distortion of 5. In fact, we show that recursive application of any single-winner rule – where a set of given size ℓ is selected by first choosing the winner when the rule is applied over the entire set of m alternatives, then over the remaining $m - 1$ alternatives, and so on until we have chosen ℓ alternatives – results in an ℓ -winner rule with the same distortion bound.

RESULT 4 (THEOREM 4.1). *Recursive application of a single-winner social choice rule gives an ℓ -winner rule (for any ℓ) achieving the same bound on distortion.*

Although this result does not directly extend to the fairness ratio, we do establish that our main result relating distortion and fairness does apply to multiple-winner rules. Combining this and the above result concerning recursive social choice rules, we show that it is possible to achieve, using Recursive Copeland, a constant-factor bound on the multiple-winner fairness ratio.

RESULT 5 (COROLLARY 4.4). *Recursive Copeland achieves a worst-case distortion of at most 5, and a worst-case fairness ratio of at most 7.*

However, note that we do not know of a better lower bound than 5 for the worst-case fairness ratio of Recursive Copeland. The lower bound of 5 follows from previous work – the upper bound of 5 for both the distortion and fairness ratio of Copeland is known to be tight [2, 21].

1.2 Related Literature

Within the general model of agents with cardinal preferences who report only ordinal information, several ways of bounding distortion have been studied in the literature [30]. It is usually assumed that agents' ordinal rankings straightforwardly match the order of their cardinal utilities, with the most-preferred alternative ranked first. In this case, it is known that when the underlying utilities are unrestricted, or simply normalized, the worst-case distortion of any deterministic social choice rule is large [31]. With randomized mechanisms, it is possible to achieve an expected distortion of $O(\sqrt{m} \log^* m)$, where m is the number of alternatives [9]. However, if the mechanism has control over how agents' utilities are translated into rankings, it is possible to design randomized rules with very low distortion [12]. The preceding results are not restricted to strategy-proof voting rules, but it is possible to construct a truthful-in-expectation mechanism whose worst-case distortion is $O(m^{3/4})$ [19].

One way to improve these distortion results is to make spatial assumptions on the underlying cardinal preferences, a technique which has a long history in social choice [16, 29]. Such models, especially those using Euclidean spaces, have naturally also been studied in the approximation algorithms literature on facility location problems [5, 14]. In these models, the cost of an agent for an alternative is given by the distance between the two.

As mentioned earlier, our work follows the literature on the analysis of distortion of social choice rules under the assumption that agent costs form an unknown metric space [2, 3]. Several lower and upper bounds for the sum of costs and median objectives are known, in both deterministic and randomized settings [2, 3]. For example, in the deterministic case, Copeland achieves a distortion of at most 5 [2], and the distortion of STV is $O(\log m)$ [33]. In the special case of Euclidean metrics, it is possible to design low-distortion mechanisms, with the additional constraint of their being truthful-in-expectation [18].

Social choice rules choosing sets of alternatives rather than a single winner have been studied for some time [13, 17, 28], but

recently they have also been evaluated within the distortion framework [11]. In most of these settings, the value of a set for an agent is determined by her favorite item in the set. Such rules have been evaluated both using distortion of utilities and an additive notion of approximation, regret; the distortion in this setting remains large – $\Theta(\sqrt{m})$ – at least when the number of winners is small compared to the number of alternatives [11]. To the best of our knowledge, multiple-winner rules have not been studied under the metric costs model.

In the distortion framework, both the interpersonal comparison of utilities, and the goal of utility maximization, are implicitly assumed to be valid. While the interpersonal comparison of utilities is more meaningful in some contexts than others [9], we take it for granted.

In addition to distortion, we borrow from the various notions of fairness studied in the context of network problems such as bandwidth allocation and load balancing [25, 26], of which approximate majorization is the most general [24]. The notion of approximate majorization has been used in the metric distortion setting to study the fairness properties of social choice rules in the form of the fairness ratio [21]. Copeland is known to achieve a fairness ratio of 5 [21], and a bound on the fairness ratio translates to a bound on social cost objectives belonging to a general class of convex functions [21, 23].

2 PRELIMINARIES

2.1 Social Choice Rules

Let \mathcal{V} be the set of agents and C the set of alternatives. We will use N to denote the total number of agents, i.e., $N = |\mathcal{V}|$, and m the number of alternatives, i.e., $m = |C|$. Every agent $v \in \mathcal{V}$ has a strict (no ties) preference ordering σ_v on C . For any $c, c' \in C$, we will use $c \succ_v c'$ to denote the fact that agent $v \in \mathcal{V}$ prefers c over c' in her ordering σ_v . Let \mathcal{S} be the set of all possible preference orderings on C . We call a profile of preference orderings $\sigma \in \mathcal{S}^N$ an *instance*.

Based on the preferences of the agents, we either want to determine a single alternative as the winner, or a set of alternatives of a given size ℓ . A (deterministic) single-winner social choice rule is a function $f : \mathcal{S}^N \rightarrow C$ that maps each instance to an alternative. An ℓ -winner social choice rule is a function $f : \mathcal{S}^N \rightarrow \{S \subseteq C : |S| = \ell\}$ that maps each instance to a set of ℓ alternatives.

To define the social choice rules that we use in this paper, we need one additional definition. We say that an alternative c *pairwise beats* c' if $|\{v \in \mathcal{V} : c \succ_v c'\}| \geq \frac{N}{2}$, with ties broken arbitrarily.

- **Copeland:** Given an instance σ , define a score for each alternative $c \in C$ given by $|\{c' \in C : c \text{ pairwise beats } c'\}|$. The alternative with the highest score (the largest number of pairwise victories, with ties broken arbitrarily) is chosen as the winner.
- **Recursive rules:** For any single-winner rule f , we define the ℓ -winner “Recursive f ” (e.g., Recursive Copeland). Given an instance σ , choose ℓ winners as follows: First pick the winner c_1 under rule f among all alternatives C , then pick the winner c_2 under f among $C \setminus \{c_1\}$, and so on until c_ℓ . The set of winners is given by $\{c_1, c_2, \dots, c_\ell\}$.

- **STV**: Given an instance σ , repeatedly eliminate the alternative ranked first by the fewest agents and remove this alternative from every ranking. The last remaining alternative is the winner.

2.2 Metric costs

We assume that the agent costs over the alternatives are given by an underlying metric d on $C \cup \mathcal{V}$, such that $d(v, c)$ is the cost of an agent v for an alternative c .

Definition 2.1. A function $d : C \cup \mathcal{V} \times C \cup \mathcal{V} \rightarrow \mathcal{R}_{\geq 0}$ is a metric if and only if $\forall x, y, z \in C \cup \mathcal{V}$, we have the following:

- (1) $d(x, y) \geq 0$
- (2) $d(x, x) = 0$
- (3) $d(x, y) = d(y, x)$
- (4) $d(x, z) \leq d(x, y) + d(y, z)$

We can do with a much simpler yet equivalent assumption on the agents' costs. We need to first define a q-metric ("q" for quadrilateral) by replacing the triangle inequalities by "quadrilateral" inequalities (Condition 2 in the definition below).

Definition 2.2. A function $d : \mathcal{V} \times C \rightarrow \mathcal{R}_{\geq 0}$ is a q-metric if and only if $\forall v, v' \in \mathcal{V}$, and $\forall c, c' \in C$, we have the following:

- (1) $d(v, c) \geq 0$
- (2) $d(v, c) \leq d(v, c') + d(v', c') + d(v', c)$

The following equivalence has been shown in earlier work [21].

LEMMA 2.3 (GOEL ET AL. [21]). *If d is a q-metric, then there exists a metric d' such that $d(v, c) = d'(v, c)$ for all $v \in \mathcal{V}$ and $c \in C$.*

Henceforth, we will use the terms *metric* and *q-metric* interchangeably.

2.3 Distortion

We say that a metric d is *consistent* with an instance σ , if whenever any agent v prefers c over c' , then her cost for c must be at most her cost for c' , i.e., $c \succ_v c' \implies d(v, c) \leq d(v, c')$. We denote by $\rho(\sigma)$ the set of all metrics d that are consistent with σ .

The social cost, ϕ , of an alternative is defined as the sum of the agent costs for it. For any metric d , and any alternative $c \in C$, we define $\phi(c, d) = \sum_{v \in \mathcal{V}} d(v, c)$. The social cost of a set of alternatives is the sum of social costs of the constituent alternatives. For any set of alternatives $S \subseteq C$, we define

$$\phi(S, d) = \sum_{v \in \mathcal{V}} \sum_{c \in S} d(v, c).$$

Below, we will define distortion in terms of choosing sets of alternatives. The corresponding definitions for single-winner social choice rules are obtained by using singleton sets instead.

We view social choice rules as trying to approximate the optimal set of alternatives of a given size ℓ , with knowledge of only the rankings σ , but not the underlying metric cost d that induces σ . To measure how close a social choice rule gets to the optimal set of size ℓ in terms of social cost, we define the *distortion* of any given set S of size ℓ to be the ratio of the social cost of S to the cost of the optimal set according to d :

$$\text{distortion}(S, d) = \frac{\phi(S, d)}{\min_{T \subseteq C: |T|=\ell} \phi(T, d)}.$$

The worst-case distortion of a set alternatives $S \subseteq C$ for a given instance σ is given by

$$\text{distortion}(S, \sigma) = \sup_{d \in \rho(\sigma)} \text{distortion}(S, d).$$

The distortion of a social choice rule f is said to be at most β , if for any instance σ , we have $\text{distortion}(f(\sigma), \sigma) \leq \beta$. In other words, the worst-case distortion of the set of alternatives chosen by f , over all possible instances σ , and all metrics $d \in \rho(\sigma)$ that are consistent with it, is at most β .

2.4 Fairness

Given an underlying metric, based on the alternative chosen, the costs incurred might vary widely among the agents. We want to formally quantify how "fair" choosing a particular alternative is. For this purpose, we look at social cost defined as the sum of k largest agent costs, for all $1 \leq k \leq N$. For any metric d and $S \subseteq C$, we define

$$\phi_k(S, d) = \max_{V \subseteq \mathcal{V}: |V|=k} \sum_{v \in V} \sum_{c \in S} d(v, c).$$

We measure fairness by a worst-case bound on how well a social choice rule approximates (simultaneously for all k) the optimal set of alternatives in terms of the social cost given by ϕ_k , with knowledge of only the rankings σ , but not the underlying metric d . To this end, we define the *fairness ratio* of a given set of alternatives S of size ℓ as

$$\text{fairness}(S, d) = \max_{1 \leq k \leq N} \frac{\phi_k(S, d)}{\min_{T \subseteq C: |T|=\ell} \phi_k(T, d)}.$$

The worst-case fairness ratio of a set of alternatives S of size ℓ , for a given instance σ then becomes

$$\text{fairness}(S, \sigma) = \sup_{d \in \rho(\sigma)} \text{fairness}(S, d).$$

Furthermore, the fairness ratio of a social choice rule f is said to be at most β , if for any instance σ , we have $\text{fairness}(f(\sigma), \sigma) \leq \beta$.

Another reason for studying the fairness ratio is that for deterministic social choice rules, a bound on the fairness ratio translates to an approximation result with respect to a broad class of convex cost functions (see Goel et al. [21]).

3 RELATING DISTORTION AND FAIRNESS

In this section, we establish our main result, which ties the distortion and fairness ratio closely together (within an additive factor of 2). This result additionally implies bounds on the fairness ratio for several popular social choice rules for which only bounds on the distortion were previously known, including k -Approval, Veto, Plurality, Borda, and STV. Until Sect. 4, we consider only single-winner social choice rules $f : \mathcal{S}^N \rightarrow C$.

We will prove that the distortion-fairness gap is at most 2 for any fixed metric, which in turn implies the gap is at most 2 for any instance (taking the supremum over all metrics) and for any rule (taking the supremum over all instances). Perhaps surprisingly, although this relationship is established on the level of a specific metric, we will show that it is tight on the level of instances, i.e., there exist instances for which the distortion-fairness gap of a given alternative approaches 2.




Figure 1: A graph metric inducing a series of instances σ_N for any $N = |\mathcal{V}|$, where $c_2 \succ_v c_1$ and $c_1 \succ_{v'} c_2$ for all $v' \neq v$. This metric shows that $\text{fairness}(c_1, \sigma_N) \geq 3$ (letting $\delta \rightarrow 0$); however $\text{distortion}(c_1, \sigma_N) \rightarrow 1$ as $N \rightarrow \infty$.

First, we note the following trivial relationship between the distortion and fairness ratio.

THEOREM 3.1. *For any instance σ and alternative c ,*

$$\text{distortion}(c, \sigma) \leq \text{fairness}(c, \sigma).$$

Theorem 3.1 also implies that if the fairness ratio of a single-winner rule f is at most β , then the distortion of f is at most β .

Next, we establish a corresponding lower bound on distortion, giving the desired gap of 2. Again, we establish the bound for fixed instances by proving it for individual metrics, and this also implies that the same bound applies to any single-winner rule f .

THEOREM 3.2. *For any instance σ and any alternative c ,*

$$\text{fairness}(c, \sigma) - 2 < \text{distortion}(c, \sigma).$$

Theorem 3.2 also implies that if the distortion of a single-winner rule f is at most β , then the fairness ratio of f is at most $\beta + 2$.

Additionally, the bound of Theorem 3.1 is tight, in the sense that there exists a series of instances and an alternative c for which the distortion-fairness gap approaches 2 as the number of agents approaches ∞ . One such series of instances is illustrated in Fig. 1.

THEOREM 3.3. *For the series of instances σ_N illustrated in Fig. 1, $\text{fairness}(c_1, \sigma_N) - \text{distortion}(c_1, \sigma_N) \rightarrow 2$ as $N \rightarrow \infty$.*

This example also implies that choosing the alternative with lowest distortion (in this case, c_1) can result in a distortion-fairness gap approaching 2. Indeed, any reasonable social choice rule should choose c_1 , which beats c_2 by an overwhelming majority, and thus would have a distortion-fairness gap approaching 2 for this instance.

In fact, this gap makes sense when we note that $\text{fairness}(c, \sigma)$ is always at least 3 unless c is ranked first by every agent. Otherwise, let v be an agent that prefers some alternative c_{adv} above c , i.e., $c_{\text{adv}} \succ_v c$, and consider a metric d such that

$$d(v', c') = \begin{cases} 3 & \text{if } v' = v, c_{\text{adv}} \succ_v c' \\ 1 & \text{otherwise} \end{cases}. \quad (1)$$

It is not hard to see that this metric satisfies the quadrilateral inequality and is consistent with σ . For this metric and $k = 1$, we have

$$\frac{\phi_1(c, d)}{\min_{c' \in C} \phi_1(c', d)} = \frac{3}{1} = 3.$$

Since this consistent metric achieves a fairness ratio of 3, we conclude that $\text{fairness}(c, \sigma) \geq 3$ unless every agent ranks c first. Since distortion can approach 1 without this strict requirement, a gap of 2 naturally arises. (However, note that this flooring effect is not the only reason for a distortion-fairness gap; such a gap may still exist when the fairness ratio is larger than 3.)

3.1 New Bounds on the Fairness of Specific Rules

Now that we have established a close relationship between the distortion and fairness ratio – within a constant additive factor – we immediately get new results for any single-winner rule where distortion is known but fairness is not. These include many rules in common use: both simple rules such as k -Approval, Veto, Plurality, and Borda [2]; and more complex ones such as STV and the general class of scoring rules, i.e., rules where each position on an agent’s ballot is worth a certain number of points and the alternative with the most points is the winner [33]. The following corollary arises directly from these known bounds and Theorems 3.1 and 3.2 establishing the relationship between the distortion and fairness ratio.

COROLLARY 3.4. *For an instance σ with N agents and m alternatives,*

- (1) *The fairness ratio of k -Approval and Veto is $\Theta(N)$.*
- (2) *The fairness ratio of Plurality and Borda is $\Theta(m)$.*
- (3) *The fairness ratio of the harmonic scoring rule¹ is $O(\frac{m}{\sqrt{\log m}})$ and $\Omega(\frac{m}{\log m})$.*
- (4) *The fairness ratio of any scoring rule is $\Omega(\sqrt{\log m})$.*
- (5) *The fairness ratio of STV is $O(\log m)$ and $\Omega(\sqrt{\log m})$.*

3.2 Calculating Fairness

For completeness, we can give a straightforward way to calculate the fairness ratio exactly, similar to the program given by Goel et al. [21] for calculating distortion. Although we do not have a polynomial time algorithm for calculating the fairness ratio, it can be computed by means of a binary linear program. We provide details in the Section 6.2 in the Appendix.

4 DISTORTION AND FAIRNESS FOR MULTIPLE-WINNER RULES

The notions of distortion and fairness ratio extend naturally to choosing a *set* of winners rather than a single alternative. In this section, we show that the same distortion bounds for single-winner rules can also be obtained for multiple-winner rules (under the sum of costs objective), and thus we can use Recursive Copeland to achieve a distortion of 5 even for multiple-winner rules. Additionally, we note that the close relationship between the distortion and fairness ratio extends to multiple-winner rules, and thus Recursive Copeland has a fairness ratio of at most 7.

First, we demonstrate a simple way to turn a single-winner social choice rule f into an ℓ -winner rule f' for any ℓ and which has the same distortion bound.

THEOREM 4.1. *Let f be a single-winner social choice rule having distortion at most β . Let $\ell > 1$. Then the ℓ -winner social choice rule $f' = \text{“Recursive } f \text{”}$ has distortion at most β .*

Thus any distortion results obtainable for single-winner social choice rules can also be obtained for ℓ -winner rules when we care about the sum of distances from an agent to each chosen alternative. Unfortunately, the results for fairness cannot be easily extended in

¹see Skowron and Elkind [33] for a definition.

the same way. In particular, the last few steps of the preceding proof relied on the ability to decompose the numerator and denominator in the calculation of distortion by separating out the costs for the individual alternatives. That is, we used the fact that

$$\phi(c_1, d) + \phi(c_2, d) = \phi(\{c_1, c_2\}, d).$$

The analogous statement for fairness does *not* hold, i.e.,

$$\phi_k(c_1, d) + \phi_k(c_2, d) \neq \phi_k(\{c_1, c_2\}, d),$$

because the specific set of agents used to calculate the social cost ϕ_k can differ in these three cases unless $k = N$. Thus the left hand side of this equation could conceivably be larger than the right hand side, e.g., if the sets of agents farthest from c_1 and those farthest from c_2 have little overlap. Intuitively, this suggests that the iterative rule f' might perform worse than the original rule f in terms of fairness ratio, because even if each chosen alternative individually has a low fairness ratio, combining the adversarial alternatives into one set could decrease the denominator.

Nevertheless, it turns out that the fairness ratio of the iterative social choice rule f' can't be too much worse than the fairness ratio of f , because the relationship between the distortion and fairness ratio extends to multiple-winner rules.

THEOREM 4.2. *For any instance σ any set of alternatives S ,*

$$\text{fairness}(S, \sigma) - 2 < \text{distortion}(S, \sigma) \leq \text{fairness}(S, \sigma).$$

COROLLARY 4.3. *Let f be a single-winner social choice rule which has fairness ratio at most β . Let $\ell > 1$. Then ℓ -winner social choice rule $f' = \text{"Recursive } f\text{"}$ has fairness ratio at most $\beta + 2$.*

This result leads to a low constant-factor bound for the fairness ratio of Recursive Copeland (though we conjecture that the true fairness ratio is lower).

COROLLARY 4.4. *The ℓ -winner rule Recursive Copeland has distortion at most 5 and fairness ratio at most 7.*

5 CONCLUSION

We demonstrated that the distinct notions of distortion and fairness ratio are in fact closely linked – within an additive factor of 2. This further lead to new bounds on the fairness ratio for several common social choice rules for which only bounds on the distortion were previously known, including STV and various scoring rules. Additionally, we showed that the distortion of any single-winner rule can also be obtained by a recursive multiple-winner rule. Together with the relationship between the distortion and fairness ratio for multiple-winner rules, this implied that Recursive Copeland achieves distortion at most 5 and fairness ratio at most 7.

ACKNOWLEDGMENTS

Ashish Goel and Anilesh K. Krishnaswamy were supported in part by NSF grant no. CCF-1637418, ONR grant no. N00014-15-1-2786 and ARO grant no. W911NF-14-1-0526. Reyna Hulett's research was supported in part by NSF GRFP grant DGE-1656518.

REFERENCES

- [1] David Andersen, Hari Balakrishnan, Frans Kaashoek, and Robert Morris. 2001. *Resilient overlay networks*. Vol. 35. ACM.
- [2] Elliot Anshelevich, Onkar Bhardwaj, and John Postl. 2015. Approximating Optimal Social Choice under Metric Preferences. *Association for the Advancement of Artificial Intelligence, 15th Conference of the* (2015).
- [3] Elliot Anshelevich and John Postl. 2016. Randomized Social Choice Functions Under Metric Preferences. *25th International Joint Conference on Artificial Intelligence* (2016).
- [4] Kenneth J Arrow. 2012. *Social choice and individual values*. Vol. 12. Yale university press.
- [5] Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. 2004. Local search heuristics for k-median and facility location problems. *SIAM Journal on computing* (2004).
- [6] Salvador Barbera. 2001. An introduction to strategy-proof social choice functions. *Social Choice and Welfare* (2001).
- [7] Salvador Barberà and Danilo Coelho. 2008. How to choose a non-controversial list with k names. *Social Choice and Welfare* 31, 1 (2008), 79–96.
- [8] Nadja Betzler, Arkadii Slinko, and Johannes Uhlmann. 2013. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research* 47 (2013), 475–519.
- [9] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D Procaccia, and Or Sheffet. 2015. Optimal social choice functions: A utilitarian view. *Artificial Intelligence* (2015).
- [10] Yves Cabannes. 2004. Participatory budgeting: a significant contribution to participatory democracy. *Environment and urbanization* 16, 1 (2004), 27–46.
- [11] Ioannis Caragiannis, Swaprava Nath, Ariel D Procaccia, and Nisarg Shah. 2017. Subset Selection Via Implicit Utilitarian Voting. *Journal of Artificial Intelligence Research* (2017).
- [12] Ioannis Caragiannis and Ariel D Procaccia. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* (2011).
- [13] John R Chamberlin and Paul N Courant. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *American Political Science Review* (1983).
- [14] Zvi Drezner and Horst W Hamacher. 1995. *Facility location*. Springer-Verlag New York, NY.
- [15] Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. 2017. Properties of multiwinner voting rules. *Social Choice and Welfare* 48, 3 (2017), 599–632.
- [16] James M Enelow and Melvin J Hinich. 1984. *The spatial theory of voting: An introduction*. CUP Archive.
- [17] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner voting: A new challenge for social choice theory. *Trends in Computational Social Choice* (2017), 27.
- [18] Michal Feldman, Amos Fiat, and Iddan Golomb. 2016. On voting and facility location. *Proceedings of the 2016 ACM Conference on Economics and Computation* (2016).
- [19] Aris Filos-Ratsikas and Peter Bro Miltersen. 2014. Truthful approximations to range voting. *International Conference on Web and Internet Economics* (2014).
- [20] Allan Gibbard. 1973. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society* (1973).
- [21] Ashish Goel, Anilesh K Krishnaswamy, and Kamesh Munagala. 2017. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 287–304.
- [22] Ashish Goel, Anilesh K Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. 2015. Knapsack voting. *Collective Intelligence* (2015).
- [23] Ashish Goel and Adam Meyerson. 2006. Simultaneous optimization via approximate majorization for concave profits or convex costs. *Algorithmica* (2006).
- [24] Ashish Goel, Adam Meyerson, and Serge Plotkin. 2001. Approximate majorization and fair online load balancing. *Symposium on Discrete Algorithms: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms* (2001).
- [25] Jon Kleinberg, Yuval Rabani, and Éva Tardos. 1999. Fairness in routing and load balancing. *Foundations of Computer Science, 1999. 40th Annual Symposium on* (1999).
- [26] Amit Kumar and Jon Kleinberg. 2000. Fairness measures for resource allocation. *Foundations of Computer Science, 41st Annual Symposium on* (2000).
- [27] Tyler Lu and Craig Boutilier. 2011. Budgeted social choice: From consensus to personalized decision making. In *IJCAI*, Vol. 11. 280–286.
- [28] Burt L Monroe. 1995. Fully Proportional Representation. *American Political Science Review* (1995).
- [29] Hervé Moulin. 1980. On strategy-proofness and single peakedness. *Public Choice* (1980).
- [30] Hervé Moulin, Felix Brandt, Vincent Conitzer, Ulle Endriss, Ariel D Procaccia, and Jérôme Lang. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- [31] Ariel D Procaccia and Jeffrey S Rosenschein. 2006. The distortion of cardinal preferences in voting. *International Workshop on Cooperative Information Agents* (2006).

- [32] Mark Allen Satterthwaite. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory* (1975).
- [33] Piotr Krzysztof Skowron and Edith Elkind. 2017. Social Choice Under Metric Preferences: Scoring Rules and STV. *AAAI* (2017).

6 APPENDIX

6.1 All proofs

PROOF OF THEOREM 3.1. Observe that for any fixed metric d , $\phi_N(c, d) = \phi(c, d)$, and therefore

$$\begin{aligned} \text{distortion}(c, d) &= \frac{\phi(c, d)}{\min_{c' \in C} \phi(c', d)} \leq \max_{1 \leq k \leq N} \frac{\phi_k(c, d)}{\min_{c' \in C} \phi_k(c', d)} \\ &= \text{fairness}(c, d) \leq \text{fairness}(c, \sigma), \end{aligned}$$

since $\text{fairness}(c, \sigma)$ is the worst-case fairness ratio over all metrics. This implies $\text{fairness}(c, \sigma)$ upper bounds $\text{distortion}(c, d)$ for every d , and thus by the definition of supremum,

$$\text{distortion}(c, \sigma) = \sup_{d \in \rho(\sigma)} \text{distortion}(c, d) \leq \text{fairness}(c, \sigma),$$

as desired. \square

PROOF OF THEOREM 3.2. Denote the N agents of σ by \mathcal{V} and the m alternatives by C . We will establish that for any fixed k and metric d ,

$$\frac{\phi_k(c, d)}{\min_{c' \in C} \phi_k(c', d)} - 2 \frac{N-1}{N} \leq \text{distortion}(c, d), \quad (2)$$

and thus taking the maximum over all k ,

$$\begin{aligned} \text{fairness}(c, d) - 2 \frac{N-1}{N} &= \max_{1 \leq k \leq N} \frac{\phi_k(c, d)}{\min_{c' \in C} \phi_k(c', d)} - 2 \frac{N-1}{N} \\ &\leq \text{distortion}(c, d) \leq \text{distortion}(c, \sigma), \end{aligned}$$

since $\text{distortion}(c, \sigma)$ is the worst-case distortion over all metrics. Thus by the definition of supremum,

$$\text{fairness}(c, \sigma) - 2 = \sup_{d \in \rho(\sigma)} \text{fairness}(c, d) - 2 < \text{distortion}(c, \sigma),$$

as desired.

For convenience, we denote the adversarial alternative for fairness by $c_{\text{opt}} = \text{argmin}_{c' \in C} \phi_k(c', d)$, and the set of k agents farthest from c by $V_k = \text{argmax}_{V \subseteq \mathcal{V}: |V|=k} \sum_{v \in V} d(v, c)$. Furthermore, let the numerator and denominator of the fairness ratio be denoted respectively as $f_k = \phi_k(c, d)$ and $g_k = \phi_k(c_{\text{opt}}, d)$. Thus we can rewrite the desired result 2 as

$$\frac{f_k}{g_k} - 2 \frac{N-1}{N} = \frac{\sum_{v \in V_k} d(v, c)}{\phi_k(c_{\text{opt}}, d)} - 2 \frac{N-1}{N} \leq \frac{\phi(c, d)}{\min_{c' \in C} \phi(c', d)}.$$

We proceed by calculating a lower bound on $\frac{\phi(c, d)}{\phi(c_{\text{opt}}, d)}$, which will immediately also apply to $\frac{\phi(c, d)}{\min_{c' \in C} \phi(c', d)}$. We divide the agents into two groups, V_k and $\mathcal{V} \setminus V_k$. Note that by definition, $\sum_{v \in V_k} d(v, c) = f_k$, and furthermore that

$$\sum_{v \in V_k} d(v, c_{\text{opt}}) \leq \max_{V \subseteq \mathcal{V}: |V|=k} \sum_{v \in V} d(v, c_{\text{opt}}) = g_k,$$

as shown in Fig. 2.




Figure 2: A partial representation of the metric d highlighting the set of k agents farthest from c , V_k , and an arbitrary agent $v \notin V_k$, with distances. Note that the top two distances refer to total distance to V_k .

Now, consider any $v \notin V_k$. For any $w \in V_k$, we have $d(w, c) \leq d(w, c_{\text{opt}}) + d(v, c) + d(v, c_{\text{opt}})$ by the quadrilateral inequality. Solving for $d(v, c)$ and averaging over all k nodes in V_k , we get

$$\begin{aligned} d(v, c) &\geq \frac{\sum_{w \in V_k} (d(w, c) - d(w, c_{\text{opt}}) - d(v, c_{\text{opt}}))}{k} \\ &\geq \frac{f_k - g_k}{k} - d(v, c_{\text{opt}}). \end{aligned} \quad (3)$$

We can now calculate a lower bound on the distortion of c relative to c_{opt} .

$$\begin{aligned} \frac{\phi(c, d)}{\phi(c_{\text{opt}}, d)} &= \frac{f_k + \sum_{v \notin V_k} d(v, c)}{\sum_{v \in \mathcal{V}} d(v, c_{\text{opt}})} \\ &\geq \frac{f_k + \sum_{v \notin V_k} d(v, c)}{N \frac{g_k}{k}} \end{aligned} \quad (4)$$

$$\begin{aligned} &\geq \frac{f_k + (N-k) \frac{f_k - g_k}{k} - \sum_{v \notin V_k} d(v, c_{\text{opt}})}{N \frac{g_k}{k}} \quad (5) \\ &= \frac{f_k}{g_k} - \frac{N-k}{N} - \frac{\sum_{v \notin V_k} d(v, c_{\text{opt}})}{N \frac{g_k}{k}}, \end{aligned}$$

where (4) follows because the average distance from an agent to c_{opt} can't be more than $\frac{1}{k}$ times the sum of the largest k distances, and (5) follows from (3). We now consider two cases. In the first, let $k \leq \frac{N}{2}$. Then $|\mathcal{V} \setminus V_k| \geq k$ so we can repeat the argument that the average distance from an agent in $\mathcal{V} \setminus V_k$ can't be more than $\frac{1}{k}$ times the sum of the largest k distances for any agents. Thus

$$\begin{aligned} \frac{\phi(c, d)}{\phi(c_{\text{opt}}, d)} &\geq \frac{f_k}{g_k} - \frac{N-k}{N} - \frac{\sum_{v \notin V_k} d(v, c_{\text{opt}})}{N \frac{g_k}{k}} \\ &\geq \frac{f_k}{g_k} - \frac{N-k}{N} - \frac{(N-k) \frac{g_k}{k}}{N \frac{g_k}{k}} \\ &= \frac{f_k}{g_k} - 2 \frac{N-k}{N} \geq \frac{f_k}{g_k} - 2 \frac{N-1}{N}, \end{aligned}$$

as desired. (Note that the last inequality is tight when $k = 1$.) On the other hand, if $k \geq \frac{N}{2}$, then $|\mathcal{V} \setminus V_k| \leq k$ so we can only say

$\sum_{v \notin V_k} d(v, c_{\text{opt}}) \leq g_k$. Thus

$$\begin{aligned} \frac{\phi(c, d)}{\phi(c_{\text{opt}}, d)} &\geq \frac{f_k}{g_k} - \frac{N-k}{N} - \frac{\sum_{v \notin V_k} d(v, c_{\text{opt}})}{N \frac{g_k}{k}} \\ &\geq \frac{f_k}{g_k} - \frac{N-k}{N} - \frac{g_k}{N \frac{g_k}{k}} \\ &= \frac{f_k}{g_k} - 1 \geq \frac{f_k}{g_k} - 2 \frac{N-1}{N}, \end{aligned}$$

as desired, where the last inequality holds provided $N \geq 2$. (When $N = 1$, distortion and fairness are identical so the desired result holds regardless.) \square

PROOF OF THEOREM 3.3. As shown in Fig. 1, for σ_N , we let $C = \{c_1, c_2\}$ where one agent v votes $c_2 \succ_v c_1$ and the remaining $N-1$ agents vote $c_1 \succ_{v'} c_2$ for $v' \in \mathcal{V} \setminus \{v\}$. As $N \rightarrow \infty$, the distortion of c_1 approaches 1, but as the example metric shows, the fairness ratio of c_1 is 3. Thus $\text{fairness}(c_1, \sigma_N) - \text{distortion}(c_1, \sigma_N)$ approaches 2 as N approaches ∞ . \square

PROOF OF THEOREM 4.1. We wish to show that for any instance σ with N agents \mathcal{V} and $m \geq \ell$ alternatives C , $\text{distortion}(f', \sigma) \leq \beta$. Equivalently, for any fixed metric d and adversarial set T of size ℓ , we wish to show

$$\frac{\phi(f'(\sigma), d)}{\phi(T, d)} \leq \beta.$$

Let $f'(\sigma) = \{c_1, \dots, c_\ell\}$ in the order in which they are selected by Recursive f . Furthermore, order T as $\{c'_1, \dots, c'_\ell\}$ such that the elements $f'(\sigma) \cap T$ (if any) appear first and in the same order as in $f'(\sigma)$. Note that this implies if any $c_j = c'_i \in T$, then $i \leq j$. Let $\sigma_i = \sigma \setminus \{c_1, \dots, c_{i-1}\}$ be the instance obtained from σ by deleting candidates $\{c_1, \dots, c_{i-1}\}$ from every ranking, so $c_i = f(\sigma_i)$, and analogously let C_i be the set of alternatives remaining in σ_i , $C_i = C \setminus \{c_1, \dots, c_{i-1}\}$. Then we know that

$$\sup_{d \in \rho(\sigma_i)} \frac{\phi(f(\sigma_i), d)}{\min_{c \in C_i} \phi(c, d)} \leq \beta,$$

simply because the distortion of f is at most β .

However, observe that the specific metric d we are considering satisfies $d \in \rho(\sigma) \subseteq \rho(\sigma_i)$ (because any consistent metric for σ is also consistent when considering only the alternatives C_i remaining in σ_i). Additionally, we must have $c'_i \in C_i = C \setminus \{c_1, \dots, c_{i-1}\}$ because of the way we ordered the alternatives of T – otherwise, we would have $c'_i = c_j$ for some $j < i$, which is a contradiction. Thus, we can look specifically at the metric d and adversarial alternative c'_i to get

$$\frac{\phi(f(\sigma_i), d)}{\phi(c'_i, d)} \leq \beta.$$

Since this holds for every i , we can combine these equations to obtain

$$\begin{aligned} \frac{\phi(f(\sigma_1), d) + \dots + \phi(f(\sigma_\ell), d)}{\phi(c'_1, d) + \dots + \phi(c'_\ell, d)} &\leq \beta \\ \frac{\sum_{c \in f'(\sigma)} \phi(c, d)}{\sum_{c \in T} \phi(c, d)} &\leq \beta \\ \frac{\phi(f'(\sigma), d)}{\phi(T, d)} &\leq \beta, \end{aligned}$$

as desired. \square

PROOF OF THEOREM 4.2. As with single-winner rules,

$$\text{distortion}(S, \sigma) \leq \text{fairness}(S, \sigma)$$

holds because fairness bounds the largest k agent costs simultaneously for all k , including $k = N$ which is equivalent to distortion.

To show $\text{fairness}(S, \sigma) - 2 < \text{distortion}(S, \sigma)$, we first observe that the quadrilateral inequality holds for fixed-size sets of alternatives. That is, if we define $d(v, S) = \sum_{c \in S} d(v, c)$ then for any $v, w \in \mathcal{V}$, $S, T \subseteq C$, $|S| = |T| = \ell$, we have $d(v, S) \leq d(v, T) + d(w, S) + d(w, T)$. This can be established from the quadrilateral inequality $d(v, c_1) \leq d(v, c_2) + d(w, c_1) + d(w, c_2)$ applied to every pair $c_1 \in S, c_2 \in T$. Summing these inequalities over the ℓ^2 such pairs and dividing by ℓ gives the desired statement for sets of ℓ elements. The remainder of the proof is identical to that of Theorem 3.2 where we replace the single alternatives c, c_{opt} with sets S, S_{opt} , so we omit the details. \square

PROOF OF COROLLARY 4.4. By Theorem 3.1, if the fairness ratio of f is at most β then the distortion of f is also at most β . Theorem 4.1 then implies the distortion of f' is at most β . Thus using Theorem 4.2, for any instance σ , we have

$$\text{fairness}(f'(\sigma), \sigma) < \text{distortion}(f'(\sigma), \sigma) + 2 \leq \beta + 2.$$

By definition, this means the fairness ratio of f' is at most $\beta + 2$, as desired. \square

6.2 Calculating Fairness

In what follows, we outline a straightforward way to calculate the fairness ratio exactly, similar to the program given by Goel et al. [21] for calculating distortion. Although we do not have a polynomial time algorithm for calculating the fairness ratio, it can be computed with the following binary linear program on an instance σ and alternative c for a fixed k and adversarial alternative c_{opt} . (Note that as with the program of Goel et al. [21], this program could be used on any specific instance to choose the alternative with lowest worst-case fairness ratio – though in this case the resulting social choice rule would not necessarily run in polynomial time.) Here we use $\{v_i : i \in [N]\}$ to represent the agents.

In the above program, M is any sufficiently large number, e.g., 3^m where m is the number of alternatives. Recall that $\rho(\sigma)$ is the set of all valid metrics consistent with σ ; this requirement can also be encoded with a polynomial number of inequalities on the distances $d(v_i, c_j)$ for $i \in [N], j \in [m]$ [21].

We will say a word about the choice of M . The program above essentially works by allowing the d_i 's to represent the k largest distances $d(v_i, c)$, and 0 for any other distances. (Meanwhile, the k largest distances $d(v_i, c_{\text{opt}})$ are represented by $d_i^{(\text{opt})}$, normalized so that their sum is at most 1; thus maximizing $\sum_{i=1}^N d_i$ is equivalent to maximizing the worst-case fairness ratio.) Thus when $b_i = 0$, we have $d_i \leq M b_i = 0$, so v_i is not included in the largest k distances, and when $b_i = 1$ we want $d_i \leq M b_i = M$ to be unrestrictive. That is, we would like M to be sufficiently large that, subject to the normalization of the distances to c_{opt} , $d(v_i, c) \leq M$ for any i and metric d .

Strictly speaking, no M can guarantee this, because for some instances σ the fairness ratio can be unbounded. However, the following lemma gives a simple way to tell whether the fairness

Algorithm 1 A binary linear program for calculating fairness.

$$\begin{aligned}
 \max \quad & \sum_{i=1}^N d_i \\
 \text{s.t.} \quad & d_i \leq d(v_i, c) & \forall i \\
 & d_i \leq M b_i & \forall i \\
 & \sum_{i=1}^N b_i \leq k \\
 & \text{binary } b_i \\
 & kt + \sum_{i=1}^N d_i^{(\text{opt})} \leq 1 \\
 & d_i^{(\text{opt})} \geq 0 & \forall i \\
 & d_i^{(\text{opt})} \geq d(v_i, c_{\text{opt}}) - t \\
 & d \in \rho(\sigma)
 \end{aligned}$$

ratio of one alternative over another is unbounded, which will also allow us to give an upper bound on M when it is bounded.

LEMMA 6.1. *Consider an instance σ with $|C| = m$ alternatives c_i and $|\mathcal{V}| = N$ agents v_i . Furthermore, consider a directed graph on C , $D = (C, E)$, where $(c_i, c_j) \in E$ whenever at least one agent prefers c_i over c_j , i.e., whenever $\exists v c_i \succ_v c_j$. Then the fairness ratio (and therefore the distortion) of an alternative c_{alg} over c_{opt} is bounded if and only if there exists a directed path $c_{\text{alg}} \rightarrow c_{\text{opt}}$ in D .*

PROOF. (\implies) We prove the contrapositive, namely, that if no directed path $c_{\text{alg}} \rightarrow c_{\text{opt}}$ exists, then the fairness ratio is unbounded. Let S be the set of vertices (alternatives) which have a directed path to c_{opt} ; by assumption $c_{\text{alg}} \notin S$. Consider the following metric: S and \mathcal{V} are co-located at one point, and the remaining alternatives $C \setminus S$, are co-located at another; the distance between the two points is 1. (Alternatively, we can say the distance between any pair of “co-located” points is $\epsilon \ll 1$.) It is not hard to see that this satisfies the triangle inequality. Furthermore, it is consistent. This holds because every agent must rank every alternative in S above every alternative in $C \setminus S$; otherwise there would exist $v, c \in C \setminus S, c' \in S$ such that $c \succ_v c'$, implying $(c, c') \in E$. But this, along with the existence of a directed path $c' \rightarrow c_{\text{opt}}$, implies the existence of a path $c \rightarrow c_{\text{opt}}$, which contradicts $c \notin S$. Thus we have defined a consistent metric, and in this metric, the fairness ratio is $\frac{1}{\epsilon}$ which approaches ∞ as $\epsilon \rightarrow 0$.

(\impliedby) Consider an arbitrary edge $(c_i, c_j) \in E$, and let d_j be any upper bound $d_j \geq d(v, c_j) \forall v$. Since $(c_i, c_j) \in E$, there must exist some agent v' such that $c_i \succ_{v'} c_j$. Thus using the quadrilateral inequality, for any agent v , we have

$$d(v, c_i) \leq d(v, c_j) + d(v', c_i) + d(v', c_j) \leq d(v, c_j) + 2d(v', c_j) \leq 3d_j$$

Thus inductively, if there exists a path $c_{\text{alg}} \rightarrow c_{\text{opt}}$, which must have length at most $|C| = m$, then $d(v, c_{\text{alg}}) \leq 3^{m-1} d_{\text{opt}}$ where d_{opt} is any upper bound on $d(v, c_{\text{opt}})$. For instance, we can set

$d_{\text{opt}} = \max_{v \in \mathcal{V}} d(v, c_{\text{opt}})$. Then the fairness ratio of c_{alg} over c_{opt} for any metric is at most

$$\max_{1 \leq k \leq N} \frac{\phi_k(c_{\text{alg}}, \sigma)}{\phi_k(c_{\text{opt}}, \sigma)} \leq \frac{N 3^{m-1} d_{\text{opt}}}{d_{\text{opt}}} = N 3^{m-1},$$

and thus is bounded. \square \square

The above proof incidentally suggests a way of setting M to be unrestrictive when the distortion and fairness ratio are bounded. Specifically, since our binary linear program normalizes

$$\max_{V \subseteq \mathcal{V}: |V|=k} \sum_V d(v, c_{\text{opt}}) \leq 1,$$

we can upper bound $d(v, c_{\text{opt}})$ with $d_{\text{opt}} = 1$ which implies $d(v, c) \leq 3^{m-1} d_{\text{opt}} = 3^{m-1}$ if the fairness ratio is bounded. Thus we can set $M = 3^m$, and we will get the correct answer whenever the fairness ratio is bounded (and the output will be $k 3^m$ if and only if the fairness ratio is unbounded).