Making Right Decisions Based on Wrong Opinions

GERDUS BENADE, ANSON KAHNG, and ARIEL D. PROCACCIA, Carnegie Mellon University

We revisit the classic problem of designing voting rules that aggregate objective opinions, in a setting where voters have noisy estimates of a true ranking of the alternatives. Previous work has replaced structural assumptions on the noise with a worst-case approach that aims to choose an outcome that minimizes the maximum error with respect to any feasible true ranking. This approach underlies algorithms that have recently been deployed on the social choice website RoboVote.org. We take a less conservative viewpoint by minimizing the *average* error with respect to the set of feasible ground truth rankings. We derive (mostly sharp) analytical bounds on the expected error and establish the practical benefits of our approach through experiments.

1 INTRODUCTION

The field of *computational social choice* [Brandt et al., 2016] has been undergoing a transformation, as rigorous approaches to voting and resource allocation, previously thought to be purely theoretical, are being applied to group decision making and social computing in practice [Chen et al., 2016]. From our biased viewpoint, RoboVote.org, a not-for-profit social choice website launched in November 2016, gives a compelling (and unquestionably recent) example. Its short-term goal is to facilitate effective group decision making by providing free access to optimization-based voting rules. In the long term, one of us has argued [Procaccia, 2016] that RoboVote and similar applications of computational social choice can change the public's perception of democracy.[1]

RoboVote distinguishes between two types of social choice tasks: aggregation of *subjective preferences*, and aggregation of *objective opinions*. Examples of the former task include a group of friends deciding where to go to dinner or which movie to watch; family members selecting a vacation spot; and faculty members choosing between faculty candidates. In all of these cases, there is no single correct choice — the goal is to choose an outcome that makes the participants as happy as possible overall.

By contrast, the latter task involves situations where some alternatives are objectively better than others, i.e., there is a true *ranking* of the alternatives by quality, but voters can only obtain noisy estimates thereof. The goal is, therefore, to aggregate these noisy opinions, which are themselves rankings of the alternatives, and uncover the true ranking. For example, consider a group of engineers deciding which product prototype to develop based on an objective metric, such as projected market share. Each prototype, if selected for development (and, ultimately, production), would achieve a particular market share, so a true ranking of the alternatives certainly exists. Other examples include a group of investors deciding which company to invest in, based on projected revenue; and employees of a movie studio selecting a movie script for production, based on projected box office earnings.

In this paper, we focus on the second setting — aggregating objective opinions. This is a problem that boasts centuries of research: it dates back to the work of the Marquis de Condorcet, published in 1785, in which he proposed a random noise model that governs how voters make mistakes when estimating the true ranking. He further suggested — albeit in a way that took 203 years to decipher [Young, 1988] — that a voting rule should be a *maximum likelihood estimator (MLE)*, that is, it should select an outcome that is most likely to coincide with the true ranking, given the observed

[1] Of course, many people disagree. We were especially amused by a reader comment on a sensationalist story about RoboVote in the British tabloid Daily Mail: "Bloody robots coming over telling us how to vote! Take our country back!"

votes and the *known* structure of the random noise model. Condorcet's approach is the foundation of a significant body of modern work [Azari Soufiani et al., 2012, 2013, 2014, Caragiannis et al., 2014, 2016, Conitzer et al., 2009, Conitzer and Sandholm, 2005, Elkind et al., 2010, Elkind and Shah, 2014, Lu and Boutilier, 2011b, Mao et al., 2013, Procaccia et al., 2012, Xia, 2016, Xia and Conitzer, 2011, Xia et al., 2010].

While the MLE approach is conceptually appealing, it is also fragile. Indeed, it advocates rules that are tailor-made for one specific noise model, which is unlikely to accurately represent real-world errors [Mao et al., 2013]. Recent work [Caragiannis et al., 2014, 2016] circumvents this problem by designing voting rules that are robust to large families of noise models, at the price of theoretical guarantees that only kick in when the number of voters is large — a reasonable assumption in crowdsourcing settings. However, here we are most interested in helping small groups of people make decisions — on RoboVote, typical instances have 4–10 voters — so this approach is a nonstarter.

1.1 The Worst-Case Approach

In recent work, Procaccia et al. [2016] have taken another step towards robustness (we will argue shortly that it is perhaps a step too far). Instead of positing a random noise model, they essentially remove all assumptions about the errors made by voters. To be specific, first fix a distance metric d on the space of rankings. For example, the *Kendall tau (KT)* distance between two rankings is the number of pairs of alternatives on which they disagree. We are given a vote profile and an upper bound t on the average distance between the input votes and the true ranking. This induces a set of *feasible* true rankings — those that are within average distance t from the votes. The *worst-case optimal voting rule* returns the ranking that minimizes the maximum distance (according to d) to any feasible true ranking. If this minimax distance is k, then we can guarantee that our output ranking is within distance k from the true ranking. The most pertinent theoretical results of Procaccia et al. are that for any distance metric d, one can always recover a ranking that is at distance at most $2t$ from the true ranking, i.e., $k \leq 2t$; and that for the four most popular distance metrics used in the social choice literature (including the KT distance), there is a tight lower bound of (roughly) $k \geq 2t$.

Arguably the more compelling results of Procaccia et al., though, are empirical. In the case of objective opinions, the measure used to evaluate a voting rule is almost indisputable: the distance (according to the distance metric of interest, say KT) between the output ranking and the *actual* true ranking. And, indeed, according to this measure, the worst-case approach significantly outperforms previous approaches — including those based on random noise models — on real data [Mao et al., 2013]; we elaborate on this dataset later.

Based on the foregoing empirical results, the algorithms deployed on RoboVote for aggregating objective opinions implement the worst-case approach. Specifically, given an upper bound t on the average KT distance between the input votes and the true ranking,[2] the algorithm computes the set of feasible true rankings (by enumerating the solutions to an integer program), and selects a ranking that minimizes the KT distance to any ranking in that set (by solving another integer program).

RoboVote also supports two additional output types: single winning alternative, and a subset of alternatives. When the user requests a single alternative as the output, the algorithm computes the set of feasible true rankings as before, and returns the alternative that minimizes the maximum position in any feasible true ranking, that is, the alternative that is guaranteed to be as close to

[2]This value is set by minimizing the average distance between any input vote and the remaining votes. This choice guarantees a nonempty set of feasible true rankings, and performs extremely well in experiments.

the top as possible. Computing a subset is similar, with the exception that the *loss* of a subset with respect to a specific feasible true ranking is determined based on the *top-ranked* alternative in the subset; the algorithm selects the subset that minimizes the maximum loss over all feasible true rankings. In other words, if this loss is s then any feasible true ranking has an alternative in the subset among its top s alternatives.

1.2 Our Approach and Results

To recap, the worst-case approach to aggregating objective opinions has proven quite successful. Nevertheless, it is very conservative, and it seems likely that better results can be achieved in practice by modifying it. We therefore take a more "optimistic" angle by carefully injecting some randomness into the worst-case approach.

In more detail, we refer to the worst-case approach as "worst case" because the errors made by voters are arbitrary, but there is actually another crucial aspect that makes it conservative: the optimization objective — minimizing the *maximum* distance to any feasible true ranking when the output is a ranking, and minimizing the maximum position or loss in any feasible true ranking when the output is a single alternative or a subset of alternatives, respectively. We propose to modify these objective functions, by replacing (in both cases) the word "maximum" with the word "average". Equivalently, we assume a uniform prior over the set of all rankings, which induces a uniform posterior over the set of feasible true rankings, and replace the word "maximum" with the word "expected".[3] Note that this model is fundamentally different from assuming that the *votes* are random: as we mentioned earlier, it is arguable whether real-world votes can be captured by any particular random noise model, not to mention a uniform distribution.[4] By contrast, we make no structural assumptions about the noise, and, in fact, we do not make any new assumptions about the world; we merely modify the optimization objective with respect to the same set of feasible true rankings.

In Section 3, we study the case where the output is a ranking. We find that for any distance metric, if the average distance between the vote profile and the true ranking is at most t, then we can recover a ranking whose average distance to the set of feasible true rankings is also t. We also establish essentially matching lower bounds for the four distance metrics studied by Procaccia et al. [2016]. Note that our relaxed goal allows us to improve their bound from $2t$ to t, which, in our view, is a *qualitative* improvement, as now we can guarantee performance that is at least as good as the average voter. While we we would like to outperform the average voter, this is a worst-case (over noisy votes) guarantee, and, as we shall see, in practice we indeed achieve excellent performance.

In Section 4, we explore the case where the output is a subset of alternatives (including the all-important case of a single winning alternative). This problem was not studied by Procaccia et al. [2016], in part because their model does not admit nontrivial analytical solutions (as we explain in detail later) — but it is just as important in practice, if not even more (see Section 1.1). We find significant gaps between the guarantees achievable under different distance metrics. Our main technical result concerns the practically significant KT distance and the closely related *footrule* distance: If the average distance between the vote profile and the true ranking is at most t, we can pinpoint a subset of alternatives of size z, whose average loss — that is, the average position of the subset's top-ranked alternative in the set of feasible true rankings (smaller position is closer to the top) — is $O(\sqrt{t/z})$. We also prove a lower bound of $\Omega(\sqrt{t}/z)$, which is tight for a constant subset size z (note that z is now outside of the square root). For the *maximum displacement* distance, we

[3]Our positive results actually work for any distribution; see Section 6.

[4]That said, some social choice papers do analyze uniformly random vote profiles [Pritchard and Wilson, 2009, Tsetlin et al., 2003] — a model known as *impartial culture*.

have asymptotically matching upper and lower bounds of $\Theta(t/z)$. Interestingly, for the *Cayley* distance and $z = 1$, we prove a lower bound of $\Omega(\sqrt{m})$, showing that there is no hope of obtaining positive results that depend only on t.

In Section 5, we present empirical results from real data. Our key finding is that our methods are robust to overestimates of the true average level of noise in the vote profile — significantly more so than the methods of Procaccia et al. [2016], which are currently deployed on RoboVote. We believe that this conclusion is meaningful for real-world implementation.

2 PRELIMINARIES

Let A be a set of *alternatives* with $|A| = m$. Let $\mathcal{L}(A)$ be the set of possible rankings of A, which we think of as permutations $\sigma : A \to [m]$, where $[m] = \{1, \ldots, m\}$. That is, $\sigma(a)$ gives the position of $a \in A$ in σ, with $\sigma^{-1}(1)$ being the highest-ranked alternative, and $\sigma^{-1}(m)$ being the lowest-ranked alternative. A ranking σ induces a strict total order $\succ_\sigma$, such that $a \succ_\sigma b$ if and only if $\sigma(a) < \sigma(b)$. A *vote profile* $\pi = (\sigma_1, \ldots, \sigma_n) \in \mathcal{L}(A)^n$ consists of n votes, where σ_i is the vote of voter i.

We next introduce notations that will simplify the creation of vote profiles. For a subset of alternatives $A_1 \subseteq A$, let σ^{A_1} be an arbitrary ranking of A_1. For a partition A_1, A_2 of A, $A_1 \succ A_2$ is a *partial order* of A which specifies that every alternative in A_1 is preferred to any alternative in A_2. Similarly, $A_1 \succ \sigma^{A_2}$ is a partial ordering where the alternatives in A_1 are preferred to those in A_2 and the order of the alternatives in A_2 is specified to coincide with σ^{A_2}. An *extension* of a partial order $\mathcal{P}$ is any ranking $\sigma \in \mathcal{L}(A)$ satisfying the partial order. Denote by $\mathcal{F}(\mathcal{P})$ the set of possible extensions of $\mathcal{P}$. For example, $|\mathcal{F}(A_1 \succ A_2)| = |A_1|! \cdot |A_2|!$ and $|\mathcal{F}(A_1 \succ \sigma^{A_2})| = |A_1|!$.

Distance metrics on permutations play an important role in the paper. We pay special attention to the following well-known distance metrics:

- The *Kendall tau (KT)* distance, denoted d_{KT}, measures the number of pairs of alternatives on which the two rankings disagree:
$$d_{KT}(\sigma, \sigma') \triangleq |\{(a, b) \in A^2 |\ a \succ_\sigma b \text{ and } b \succ_{\sigma'} a\}|.$$
Equivalently, the KT distance between σ and σ' is the number of swaps between *adjacent* alternatives required to transform one ranking into the other. Some like to think of it as the "bubble sort" distance.
- The *footrule* distance, denoted d_{FR}, measures the total displacement of alternatives between two rankings:
$$d_{FR}(\sigma, \sigma') \triangleq \sum_{a \in A} |\sigma(a) - \sigma'(a)|.$$
- The *maximum displacement* distance, denoted d_{MD}, measures the largest absolute displacement of any alternative between two rankings:
$$d_{MD}(\sigma, \sigma') \triangleq \max_{a \in A} |\sigma(a) - \sigma'(a)|.$$
- The *Cayley* distance, denoted d_{CY}, measures the number of pairwise swaps required to transform one ranking into the other. In contrast to the KT distance, the swapped alternatives need not be adjacent.

We also require the following definitions that apply to any distance metric d. For a ranking $\sigma \in \mathcal{L}(A)$ and a set of rankings $S \subseteq \mathcal{L}(A)$, define the average distance between σ and S in the obvious way,
$$d(\sigma, S) \triangleq \frac{1}{|S|} \sum_{\sigma' \in S} d(\sigma, \sigma').$$

Similarly, define the average distance between two sets of rankings $S, T \subseteq \mathcal{L}(A)$ as

$$d(S,T) \triangleq \frac{1}{|S| \cdot |T|} \sum_{\sigma \in S} \sum_{\sigma' \in T} d(\sigma, \sigma').$$

Finally, let $d^{\downarrow}(k)$ be the largest distance allowed under the distance metric d which is at most k, i.e.,

$$d^{\downarrow}(k) \triangleq \max\{s \leq k : \exists \sigma, \sigma' \in \mathcal{L}(A) \text{ s.t. } d(\sigma, \sigma') = s\}.$$

3 RETURNING THE RIGHT RANKING, IN THEORY

We first tackle the setting where our goal is to return an accurate ranking. We assume that there is an objective ground truth ranking σ^*, and that n voters submit a *vote profile* π of noisy estimates of this true ranking. As in the work of Procaccia et al. [2016], an individual vote is allowed to deviate from the ground truth in any way, but we expect that the average error is bounded, that is, the average distance between the vote profile and the ground truth is no more than some parameter t. Formally, for a distance metric d on $\mathcal{L}(A)$, we are guaranteed that

$$d(\pi, \sigma^*) = \frac{1}{n} \sum_{\sigma \in \pi} d(\sigma, \sigma^*) \leq t.$$

There are several approaches for obtaining good estimates for this upper bound t; we return to this point later.

A combinatorial structure that plays a central role in our analysis is the "ball" of feasible ground truth rankings,

$$\mathcal{B}_t(\pi) \triangleq \{\sigma \in \mathcal{L}(A) : d(\pi, \sigma) \leq t\}.$$

If this ball were a singleton (or empty), our task would be easy. But it typically contains multiple feasible ground truths, as the following example shows.

Example 3.1. Suppose that $A = \{a, b, c\}$ and the vote profile consists of 5 votes, $\pi = \{(a \succ b \succ c), (a \succ b \succ c), (b \succ c \succ a), (c \succ a \succ b), (a \succ c \succ b)\}$. For each distance metric, let the bound on average error equal half of the maximum distance allowed by the distance metric; in other words, $t_{KT} = 1.5, t_{FR} = 2, t_{MD} = 1$ and $t_{CY} = 1$. The set of feasible ground truths for the vote profile π under the respective distance metrics may be found in Table 1.

Table 1. The set of feasible ground truths in Example 3.1 for various distance metrics.

d	t	$\mathcal{B}_t(\pi)$
KT	1.5	$\{(a \succ b \succ c), (c \succ a \succ b), (a \succ c \succ b)\}$
FR	2	$\{(a \succ b \succ c), (a \succ c \succ b)\}$
MD	1	$\{(a \succ b \succ c), (a \succ c \succ b)\}$
CY	1	$\{(a \succ b \succ c), (a \succ c \succ b)\}$

Procaccia et al. [2016] advocate a conservative approach — they choose a ranking that minimizes the *maximum* distance to any feasible ground truth. By contrast, we are concerned with the *average* distance to the set of feasible ground truths. In other words, we assume that each of the feasible ground truths is equally likely, and our goal is to find a ranking that has a small expected distance to the set of feasible ground truths $\mathcal{B}_t(\pi)$.

Our first result is that is it always possible to find a ranking $\sigma \in \pi$ that is close to $\mathcal{B}_t(\pi)$.

Theorem 3.2. *Given a profile π of n noisy rankings with average distance at most t from the ground truth according to some distance metric d, there always exists a ranking within average distance t from the set of feasible ground truths $\mathcal{B}_t(\pi)$ according to the same metric.*

Proof. For any $\sigma \in \mathcal{B}_t(\pi)$, $d(\sigma, \pi) \leq t$. It follows from the definitions that

$$d(\pi, \mathcal{B}_t(\pi)) = \frac{1}{n \cdot |\mathcal{B}_t(\pi)|} \sum_{\sigma' \in \pi} \sum_{\sigma \in \mathcal{B}_t(\pi)} d(\sigma, \sigma') = \frac{1}{|\mathcal{B}_t(\pi)|} \sum_{\sigma \in \mathcal{B}_t(\pi)} \frac{1}{n} \sum_{\sigma' \in \pi} d(\sigma, \sigma')$$

$$= \frac{1}{|\mathcal{B}_t(\pi)|} \sum_{\sigma \in \mathcal{B}_t(\pi)} d(\sigma, \pi) \leq t.$$

To conclude the proof, observe that if the average distance from π to $\mathcal{B}_t(\pi)$ is no more than t, then there certainly exists $\sigma'' \in \pi$ with $d(\sigma'', \mathcal{B}_t(\pi)) \leq t$. □

This result holds for any distance metric. Interestingly, it also generalizes to any probability distribution over $\mathcal{B}_t(\pi)$, not just the uniform distribution (see Section 6 for additional discussion of this point).

We next derive essentially matching lower bounds for the four common distance metrics introduced in Section 2.

Theorem 3.3. *For $d \in \{d_{KT}, d_{FR}, d_{MD}, d_{CY}\}$, there exists a profile π of n noisy rankings with average distance at most t from the ground truth, such that for any ranking, its average distance (according to d) from $\mathcal{B}_t(\pi)$ is at least $d^{\downarrow}(2t)/2$.*

The proof of this theorem relies heavily on constructions by Procaccia et al. [2016]; it is relegated to Appendix A.

4 RETURNING THE RIGHT ALTERNATIVES, IN THEORY

In the previous section, we derived bounds on the expected distance of the ranking closest to the set of feasible ground truth rankings. In practice, we may not be interested in eliciting a complete ranking of alternatives, but rather in selecting a subset of the alternatives (often a single alternative) on which to focus attention, time, or effort.

In this section, we bound the average position of the *best* alternative in a subset of alternatives, where the average is taken over the set of feasible ground truths as before. This type of utility function, where the utility of a set is defined by its highest utility member, is consistent with quite a few previous papers that deal with selecting subsets of alternatives in different social choice settings [Caragiannis et al., 2017, Chamberlin and Courant, 1983, Lu and Boutilier, 2011a, Monroe, 1995, Procaccia et al., 2012, 2008]. For example, when selecting a menu of movies to show on a three hour flight, the utility of passengers depends on their most preferred alternative. From a technical viewpoint, this choice has the advantage of giving bounds that improve as the subset size increases, which matches our intuition. Of course, in the important special case where the subset is a singleton, all reasonable definitions coincide.

Formally, let $Z \subseteq A$ be a subset of alternatives; the *loss of Z in σ* is $\ell(Z, \sigma) \triangleq \min_{a \in Z} \sigma(a)$, and therefore the *average loss of Z in $\mathcal{B}_t(\pi)$* is

$$\ell(Z, \mathcal{B}_t(\pi)) \triangleq \frac{1}{|\mathcal{B}_t(\pi)|} \sum_{\sigma \in \mathcal{B}_t(\pi)} \ell(Z, \sigma).$$

For given average error t and subset size z, we are interested in bounding

$$\max_{\pi \in \mathcal{L}(A)^n} \min_{Z \subseteq A \text{ s.t. } |Z| = z} \ell(Z, \mathcal{B}_t(\pi)).$$

In words, we wish to bound the the average loss of the best Z (of size z) in $\mathcal{B}_t(\pi)$, in the worst case over vote profiles.

Let us return to Example 3.1. For the footrule, maximum displacement, and Cayley distance metrics, it is clear from Table 1 that selecting $\{a\}$ when $z = 1$ guarantees average loss 1, as $\mathcal{B}_t(\pi)$ only contains rankings that place a first. For the KT distance, the set $\{a\}$ has average loss 4/3, and the set $\{a, c\}$ has average loss 1.

We now turn to the technical results, starting with some lemmas that are independent of specific distance metrics. Throughout this section, we will rely on the following lemma, which is the discrete analogue of selecting a set of z numbers uniformly at random in an interval and studying their order statistics. No doubt someone has proved it in the past, but we include our (cute, if we may say so ourselves) proof, as we will need to reuse specific equations.

Lemma 4.1. *When choosing z elements $Y_1, \ldots, Y_z$ uniformly at random without replacement from the set $[k]$, $\mathbb{E}[\min_{i\in[z]} Y_i] = \frac{k+1}{z+1}$.*

Proof. Let $Y_{min} = \min_{i\in[z]} Y_i$ be the minimum value of the z numbers chosen uniformly at random from $[k]$ without replacement. It holds that

$$\Pr[Y_{min} = y] = \frac{\binom{k-y}{z-1}}{\binom{k}{z}},$$

and therefore

$$\mathbb{E}[Y_{min}] = \sum_{y=1}^{k} y\frac{\binom{k-y}{z-1}}{\binom{k}{z}} = \frac{1}{\binom{k}{z}}\sum_{y=1}^{k} y\binom{k-y}{z-1} = \frac{1}{\binom{k}{z}}\sum_{y=1}^{k-z+1} y\binom{k-y}{z-1}. \tag{1}$$

We claim that

$$\sum_{y=1}^{k-z+1} y\binom{k-y}{z-1} = \binom{k+1}{z+1}. \tag{2}$$

Indeed, the left hand side can be interpreted as follows: for each choice of $y \in [k-z+1]$, elements $\{1, \ldots, y\}$ form a committee of size y. We have y possibilities for choosing the head of the committee. Then we choose $z-1$ clerks among the elements $\{y+1, \ldots, k\}$. We can interpret the right hand side of Equation (2) in the same way. To see how, choose $z+1$ elements from $[k+1]$, and sort them in increasing order to obtain $s_1, \ldots, s_{z+1}$. Now s_1 is the head of the committee, $y = s_2 - 1$ is the number of committee members, and $s_3 - 1, \ldots, s_{z+1} - 1$ are the clerks.

Plugging Equation (2) into Equation (1), we get

$$\mathbb{E}[Y_{min}] = \frac{\binom{k+1}{z+1}}{\binom{k}{z}} = \frac{k+1}{z+1}.$$

□

Our strategy for proving upper bounds also relies on the following lemma, which relates the performance of randomized rules on the worst ranking in $\mathcal{B}_t(\pi)$, to the performance of deterministic rules on average, and is reminiscent of Yao's Minimax Principle [Yao, 1977]. This lemma actually holds for any distribution over ground truth rankings, as we discuss in Section 6.

Lemma 4.2. *Suppose that for a given $\mathcal{B}_t(\pi)$, there exists a distribution D over subsets of A of size z such that*

$$\max_{\sigma\in\mathcal{B}_t(\pi)} \mathbb{E}_{Z\sim D}[\ell(Z, \sigma)] = k.$$

Then there exists $Z^* \subseteq A$ of size z whose average loss in $\mathcal{B}_t(\pi)$ is at most k.

Proof. Let U be the uniform distribution over rankings in $\mathcal{B}_t(\pi)$. Then clearly

$$\mathbb{E}_{Z\sim D,\sigma\sim U}\left[\ell(Z,\sigma)\right] \leq k,$$

as this inequality holds pointwise for all $\sigma \in \mathcal{B}_t(\pi)$. It follows there must exist at least one Z^* such that

$$\ell(Z^*, \mathcal{B}_t(\pi)) = \mathbb{E}_{\sigma\sim U}\left[\ell(Z^*,\sigma)\right] \leq k,$$

that is, the average loss of Z^* in $\mathcal{B}_t(\pi)$ is at most k. □

Finally, we require a simple lemma of Procaccia et al. [2016].

Lemma 4.3. *Given a profile π of n noisy rankings with average distance at most t from the ground truth according to a distance metric d, there exists $\sigma \in \mathcal{L}(A)$ such that for all $\tau \in \mathcal{B}_t(\pi)$, $d(\sigma,\tau) \leq 2t$.*

4.1 The KT and Footrule Distances

We first focus on the KT distance and the footrule distance. The KT distance is by far the most important distance metric over permutations, both in theory, and in practice (see Section 1.1). We study it together with the footrule distance because the two distances are closely related, as the following lemma, due to Diaconis and Graham [1977], shows.

Lemma 4.4. *For all $\sigma, \sigma' \in \mathcal{L}(A)$, $d_{KT}(\sigma,\sigma') \leq d_{FR}(\sigma,\sigma') \leq 2d_{KT}(\sigma,\sigma')$.*

Despite this close connection between the two metrics, it is important to note that it does not allow us to automatically transform a bound on the loss for one into a bound for the other.

The next upper bound is, in our view, our most significant theoretical result. It is formulated for the footrule distance, but, as we show shortly, also holds for the KT distance.

Theorem 4.5. *For $d = d_{FR}$, given a profile π of n noisy rankings with average distance at most t from the ground truth, and a number $z \in [m]$, there always exists a subset of size z whose average loss in the set of feasible ground truths $\mathcal{B}_t(\pi)$ is at most $O(\sqrt{t/z})$.*

At some point in the proof, we will rely on the following (almost trivial) lemma.

Lemma 4.6. *Given two positive sequences of k real numbers, P, and Q, such that P is non-decreasing, Q is strictly decreasing and $\sum_{i=1}^{k} P_i = C$, the sequence P that maximizes $S = \sum_{i=1}^{n} P_i Q_i$ is constant, i.e., $P_i = C/k$ for all $i \in [k]$.*

Proof. Assume for contradiction that P maximizes S and contains consecutive elements such that $P_j < P_{j+1}$. Now moving mass from P_{j+1} and distributing it to all lower positions in the sequence will strictly increase S. Concretely, if $P_{j+1} = P_j + \varepsilon$, we can subtract $j\varepsilon/(j+1)$ from P_{j+1} and add $\varepsilon/(j+1)$ to P_i for all $i \in [j]$. Because Q is strictly decreasing, this increases S by

$$\left(\sum_{i=1}^{j} \frac{Q_i \varepsilon}{j+1}\right) - \frac{Q_{j+1} j\varepsilon}{j+1} > \left(\sum_{i=1}^{j} \frac{Q_j \varepsilon}{j+1}\right) - \frac{Q_{j+1} j\varepsilon}{j+1} = \frac{j\varepsilon}{j+1}(Q_j - Q_{j+1}) > 0,$$

contradicting the assumption that P maximizes S. We may conclude that P is constant. □

Proof of Theorem 4.5. By Lemma 4.2, it is sufficient to construct a randomized rule that has expected loss at most $O(\sqrt{t/z})$ in *any* ranking in $\mathcal{B}_t(\pi)$. To this end, let $\sigma \in \mathcal{L}(A)$ such that $d(\sigma,\tau) \leq 2t$ for any $\tau \in \mathcal{B}_t(\pi)$; its existence is guaranteed by Lemma 4.3. Let $k = \sqrt{tz}$, and assume for ease of exposition that k is an integer. For $y = 1, \ldots, k$, let $a_y = \sigma^{-1}(y)$. Our randomized rule simply selects z alternatives uniformly at random from the top k alternatives in σ, that is, from the set $T \triangleq \{a_1, \ldots, a_k\}$. So, fixing some $\tau \in \mathcal{B}_t(\pi)$, we need to show that choosing z elements uniformly at random from the worst-case positions occupied by T in τ has expected minimum position at most $O(\sqrt{t/z})$.

Let Y^{σ}_{min} be the minimum position in σ of a random subset of size z from T. By Lemma 4.1 and Equation (1), we have

$$\mathbb{E}[Y^{\sigma}_{min}] = \sum_{y=1}^{k} y \frac{\binom{k-y}{z-1}}{\binom{k}{z}} = \frac{k+1}{z+1}.$$

However, we are interested in the positions of these elements in $\tau \in \mathcal{B}_t(\pi)$, not σ. Instead of appearing in position y, alternative a_y appears in position $p_y \triangleq \tau(a_y)$. Therefore, the expected minimum position in τ is

$$\mathbb{E}[Y^{\tau}_{min}] = \sum_{y=1}^{k} p_y \frac{\binom{k-y}{z-1}}{\binom{k}{z}}.$$

We wish to upper bound $\mathbb{E}[Y^{\tau}_{min}]$. Equivalently, because $\mathbb{E}[Y^{\sigma}_{min}]$ is fixed and independent of τ, it is sufficient to maximize the expression

$$\begin{aligned}\mathbb{E}[Y^{\tau}_{min}] - \mathbb{E}[Y^{\sigma}_{min}] &= \sum_{y=1}^{k} p_y \frac{\binom{k-y}{z-1}}{\binom{k}{z}} - \sum_{y=1}^{k} y \frac{\binom{k-y}{z-1}}{\binom{k}{z}} \\ &= \sum_{y=1}^{k} (p_y - y) \frac{\binom{k-y}{z-1}}{\binom{k}{z}}.\end{aligned} \tag{3}$$

Let us now assume that $p_y < p_{y+1}$ for all $y \in [k-1]$, that is, τ and σ agree on the order of the alternatives in T; we will remove this assumption later. Since the original positions of the alternatives in T were $\{1, \ldots, k\}$ it follows that $p_y \geq y$ for all $y \in [k]$. Moreover, because

$$\frac{\binom{k-y}{z-1}}{\binom{k}{z}} > \frac{\binom{k-(y+1)}{z-1}}{\binom{k}{z}},$$

the sequence of probabilities

$$Q = \left\{ \frac{\binom{k-y}{z-1}}{\binom{k}{z}} \right\}_{y \in [k]}$$

is strictly decreasing in y. Additionally, the sequence $P = \{p_y - y\}_{y \in [k]}$ is non-decreasing, because $p_{y+1} > p_y$, coupled with the fact that both values are integers, implies that $p_{y+1} \geq p_y + 1$.

In light of these facts, let us return to Equation (3). We wish to maximize

$$\mathbb{E}[Y^{\tau}_{min}] - \mathbb{E}[Y^{\sigma}_{min}] = \sum_{y=1}^{k} (p_y - y) \frac{\binom{k-y}{z-1}}{\binom{k}{z}} = \sum_{y=1}^{k} P_y Q_y.$$

By Lemma 4.6, $p_y - y$ is the same for all $y \in [k]$, that is, all alternatives in T are shifted by the same amount from σ to form τ. Moreover, we have that

$$\sum_{y=1}^{k} (p_y - y) \leq d(\sigma, \tau) \leq 2t.$$

Using $k = |T| = \sqrt{zt}$, we conclude that $p_y - y \leq 2\sqrt{t/z}$ for all $y \in [k]$. Therefore, in the worst $\tau \in \mathcal{B}_t(\pi)$, we have that the alternatives in T occupy positions $2\sqrt{t/z} + 1$ to $2\sqrt{t/z} + \sqrt{tz}$ in τ. By

Lemma 4.1, the expected minimum position of T in τ is

$$2\sqrt{\frac{t}{z}} + \frac{\sqrt{tz}+1}{z+1} = O\left(\sqrt{\frac{t}{z}}\right).$$

To complete the proof, it remains to show that our assumption that $p_y < p_{y+1}$ for all $y \in [k-1]$ is without loss of generality. To see this, note that since we are selecting uniformly at random from T, Y^{τ}_{min} only depends on the positions occupied by T in τ. Moreover, if τ does not preserve the order over T, we can find a ranking τ' that has the following properties:

(1) $d(\sigma, \tau') \leq 2t$.
(2) T occupies the same positions: $\{\tau(a_1), \ldots, \tau(a_k)\} = \{\tau'(a_1), \ldots, \tau'(a_k)\}$.
(3) τ' preserves the order over T: $\tau'(a_y) < \tau'(a_{y+1})$ for all $y \in [k-1]$.

Now all our arguments would apply to τ', and $\mathbb{E}[Y^{\tau}_{min}] = \mathbb{E}[Y^{\tau'}_{min}]$.

In order to construct τ', suppose that $\tau(a_y) > \tau(a_{y+1})$, and consider τ'' that is identical to τ except for swapping a_y and a_{y+1}. Then

$$\begin{aligned} d(\tau'', \sigma) &= d(\tau, \sigma) + \Big(|\tau''(a_y) - y| + |\tau''(a_{y+1}) - (y+1)| - |\tau(a_y) - y| - |\tau(a_{y+1} - (y+1)|\Big) \\ &\leq d(\tau, \sigma) \leq 2t. \end{aligned}$$

By iteratively swapping alternatives we can easily obtain the desired τ'. □

We next formulate the same result for the KT distance. The proof is very similar, so instead of repeating it, we just give a proof sketch that highlights the differences.

Theorem 4.7. *For $d = d_{KT}$, given a profile π of n noisy rankings with average distance at most t from the ground truth, and a number $z \in [m]$, there always exists a subset of size z whose average loss in the set of feasible ground truths $\mathcal{B}_t(\pi)$ is at most $O(\sqrt{t/z})$.*

Proof sketch. The proof only differs from the proof of Theorem 4.7 in two places.

First, the footrule proof had the inequality

$$\sum_{y=1}^{k} (p_y - y) \leq d_{FR}(\sigma, \tau) \leq 2t.$$

In our case,

$$\sum_{y=1}^{k} (p_y - y) \leq d_{FR}(\sigma, \tau) \leq 2 \cdot d_{KT}(\sigma, \tau) \leq 4t,$$

where the second inequality follows from Lemma 4.4.

Second, if τ does not preserve the order over T, we needed to find a ranking τ' that has the following properties:

(1) $d(\sigma, \tau') \leq 2t$.
(2) T occupies the same positions: $\{\tau(a_1), \ldots, \tau(a_k)\} = \{\tau'(a_1), \ldots, \tau'(a_k)\}$.
(3) τ' preserves the order over T: $\tau'(a_y) < \tau'(a_{y+1})$ for all $y \in [k-1]$.

To construct τ' under $d = d_{KT}$, we use the same strategy as before: Suppose that $\tau(a_y) > \tau(a_{y+1})$, and consider τ'' that is identical to τ except for swapping a_y and a_{y+1}. We claim that $d(\tau', \sigma) \leq d(\tau, \sigma) \leq 2t$. Indeed, notice that all $a \in T$ precede all $b \in A \setminus T$ in σ. Therefore, holding all else equal, switching the relative order of alternatives in T will not change the number of pairwise disagreements on alternatives $b \in T$, $b' \in A \setminus T$, nor will it change the number of pairwise disagreements on alternatives $b, b' \in A \setminus T$. It will only (strictly) decrease the number of disagreements on alternatives in T. □

Our next result is a lower bound of $\Omega(\sqrt{t}/z)$ for both distance metrics. Note that here z is outside the square root, i.e., there is a gap of $\sqrt{z}$ between the upper bounds given in Theorems 4.5 and 4.7, and the lower bound. That said, the lower bound is tight for a constant z, including the important case of $z = 1$.

Theorem 4.8. *For $d \in \{d_{FR}, d_{KT}\}$, $z \in [m]$, and an even n, there exist $t = O(m^2)$ and a profile π of n noisy rankings with average distance at most t from the ground truth, such that for any subset of size z, its average loss in the set of feasible ground truths $\mathcal{B}_t(\pi)$ is at least $\Omega(\sqrt{t}/z)$.*

Proof. We first prove the theorem for the KT distance, that is, $d = d_{KT}$. For any $k \geq 1$, let $t = \binom{k}{2}/2$; equivalently, let

$$k = \frac{1 + \sqrt{1 + 16t}}{2} = \Theta\left(\sqrt{t}\right).$$

Let $\sigma = (a_1 \succ \cdots \succ a_m)$, and let $\sigma_{R(k)} = (a_k, a_{k-1}, \ldots, a_1, a_{k+1}, \ldots, a_m)$ be the ranking that reverses the first k alternatives of σ. Consider the vote profile π with $n/2$ copies of each ranking σ and $\sigma_{R(k)}$.

Let $A_k = \{a_1, \ldots, a_k\}$ and denote by σ_{-k} the ranking of $A \setminus A_k$ ordered as in σ. We claim that $\mathcal{B}_t(\pi) = \mathcal{F}(A_k \succ \sigma_{-k})$, i.e., exactly the rankings that have some permutation of A_k in the first k positions, and coincide with σ in all the other positions. Indeed, consider any $\tau \in \mathcal{L}(A)$. This ranking will disagree with exactly one of σ and $\sigma_{R(k)}$ on every pair of alternatives in A_k, so

$$d(\tau, \pi) \geq \frac{\binom{k}{2}}{2} = t.$$

It follows that if $\tau \in \mathcal{B}_t(\pi)$ then τ must agree with σ_{-k} on the remaining alternatives.

Now let Z be a subset of z alternatives. Note that for every $a \in A \setminus A_k$ and $\tau \in B$, $\tau(a) > k$, so it is best to choose $Z \subset A_k$. We are interested in the expected loss of Z under the uniform distribution on $\mathcal{B}_t(\pi)$, which amounts to a random permutation of A_k. This is the same as choosing z positions at random from $[k]$. By Lemma 4.1, the expected minimum position of a randomly chosen subset of size z is $\frac{k+1}{z+1}$. Since $k = \frac{1+\sqrt{1+16t}}{2}$, it holds that

$$\mathbb{E}[Y_{min}] = \frac{\frac{1+\sqrt{1+16t}}{2} + 1}{z + 1} = \Omega\left(\frac{\sqrt{t}}{z}\right).$$

For $d = d_{FR}$, the construction is analogous to above, with one minor modification. For any $k \geq 1$, we let $t = \lfloor k^2/2 \rfloor/2$, because the footrule distance between σ and $\sigma_{R(k)}$ is $\lfloor k^2/2 \rfloor$, instead of $\binom{k}{2}$ as in the KT case. Now, the proof proceeds as before. □

An important remark is in order. Suppose that instead of measuring the average loss of the subset Z in $\mathcal{B}_t(\pi)$, we measured the *maximum* loss in any ranking in $\mathcal{B}_t(\pi)$, in the spirit of the model of Procaccia et al. [2016]. Then the results would be *qualitatively* different. To see why on an intuitive level, consider the KT distance, and suppose that the vote profile π consists of n copies of the same ranking σ. Then for any $a \in A$, $\mathcal{B}_t(\pi)$ includes a ranking σ' such that $\sigma'(a) \geq t$ (by using our "budget" of t to move a downwards in the ranking). Therefore, for $z = 1$, it is impossible to choose an alternative whose maximum position (i.e., loss) in $\mathcal{B}_t(\pi)$ is smaller than t. In contrast, Theorem 4.5 gives us an upper bound of $O(\sqrt{t})$ in our model.

4.2 The Maximum Displacement Distance

We now turn to the maximum displacement distance. Here the bounds are significantly worse than in the KT and footrule settings. On an intuitive level, the reason is that two rankings that are at maximum displacement distance t from each other can be drastically different, because *every* alternative can move by up to t positions. Therefore, $\mathcal{B}_t(\pi)$ under maximum displacement would typically be larger than under the distance metrics we previously considered. Indeed, this is the case in Example 3.1 if one sets $t_{MD} \geq 1.5$.

Theorem 4.9. *For $d = d_{MD}$, given a profile π of n noisy rankings with average distance at most t from the ground truth, and a number $z \in [m]$, there always exists a subset of size z whose average loss in the set of feasible ground truths $\mathcal{B}_t(\pi)$ is at most $O(t/z)$.*

Proof. By Lemma 4.2, it is sufficient to construct a randomized rule that has expected loss at most $O(t/z)$ in *any* ranking in $\mathcal{B}_t(\pi)$. To this end, let $\sigma \in \mathcal{L}(A)$ such that $d(\sigma, \tau) \leq 2t$ for any $\tau \in \mathcal{B}_t(\pi)$; its existence is guaranteed by Lemma 4.3. For $y = 1, \ldots, 3t$, let $a_y = \sigma^{-1}(y)$. Our randomized rule selects z alternatives uniformly at random from the top $3t$ alternatives in σ, that is, from the set $T \triangleq \{a_1, \ldots, a_{3t}\}$.

Let T' be the top t alternatives in a ranking $\tau \in \mathcal{B}_t(\pi)$. Since $d(\sigma, \tau) \leq 2t$, we know that $T' \subset T$. Moreover, for any $a_y \in T$, we have that $p_y \triangleq \tau(a_y) \leq 5t$. Assume without loss of generality that $p_y \leq p_{y+1}$ for all $y \in [3t-1]$; then we have that the vector of positions $(p_1, \ldots, p_{3t})$ is pointwise at least as small as the vector $(1, 2, \ldots, t, 5t, 5t, \ldots, 5t)$. Using Lemma 4.1 and Equation (1), we conclude that the minimum position in τ when selecting z alternatives uniformly at random from T, denoted Y^{τ}_{min}, satisfies

$$\mathbb{E}[Y^{\tau}_{min}] = \sum_{y=1}^{3t} p_y \frac{\binom{3t-y}{z-1}}{\binom{3t}{z}} = \sum_{y=1}^{t-1} p_y \frac{\binom{3t-y}{z-1}}{\binom{3t}{z}} + \sum_{y=t}^{3t} p_y \frac{\binom{3t-y}{z-1}}{\binom{3t}{z}} \leq \sum_{y=1}^{t-1} y \frac{\binom{3t-y}{z-1}}{\binom{3t}{z}} + \sum_{y=t}^{3t} 5t \frac{\binom{3t-y}{z-1}}{\binom{3t}{z}}$$

$$\leq 5 \cdot \sum_{y=1}^{3t} y \frac{\binom{3t-y}{z-1}}{\binom{3t}{z}} = 5 \cdot \frac{3t+1}{z+1} = \Theta\left(\frac{t}{z}\right).$$

□

We next establish a lower bound of $\Omega(t/z)$ on the average loss achievable under the maximum displacement distance. Note that this lower bound matches the upper bound of Theorem 4.9.

Theorem 4.10. *For $d = d_{MD}$, given $k \in \mathbb{N}$ and $z \in [m]$, there exist $t = \Theta(k)$ and a vote profile π of $k!$ noisy votes at average distance at most t from the ground truth, such that for any subset of size z, its average loss in the set of feasible ground truths $\mathcal{B}_t(\pi)$ is at least $\Omega(t/z)$.*

Proof. Let $\pi = \mathcal{F}(A_k \succ \sigma^{A \setminus A_k})$, where $|A_k| = k$. For some $\tau \in \pi$, let $t = d(\tau, \pi)$. By symmetry, $\tau' \in \mathcal{B}_t(\pi)$ for all $\tau' \in \pi$.

We first claim that $t = \Omega(k)$. Indeed, t is the average distance between τ and π. Letting U be the uniform distribution over π, we have that $t = \mathbb{E}_{\tau' \sim U}[d(\tau, \tau')]$. Now consider the top-ranked alternative in τ, $a \triangleq \tau^{-1}(1)$. Because U amounts to a random permutation over A_k, it clearly holds that $\mathbb{E}_{\tau' \sim U}[\tau'(a)] = (k+1)/2$, and therefore

$$t = \mathbb{E}_{\tau' \sim U}[d(\tau, \tau')] = \mathbb{E}_{\tau' \sim U}\left[\max_{b \in A} |\tau'(b) - \tau(b)|\right] \geq \mathbb{E}_{\tau' \sim U}[\tau'(a) - \tau(a)] = \frac{k+1}{2} - 1 = \Omega(k).$$

Now, suppose that we have shown that $\mathcal{B}_t(\pi) = \pi$; we argue that the theorem follows. Let $Z \subseteq A$ be a subset of alternatives of size z. We can assume without loss of generality that $Z \subseteq A_k$, as A_k is ranked at the top of every $\tau \in \mathcal{B}_t(\pi)$. But because $\mathcal{B}_t(\pi)$ consists of all permutations of A_k, $\ell(Z, \mathcal{B}_t(\pi))$ is equal to the expected minimum position when z elements are selected uniformly at random from the positions occupied by A_k, namely $[k]$. That is, we have that

$$\ell(Z, \mathcal{B}_t(\pi)) = \frac{k+1}{z+1} = \Omega\left(\frac{t}{z}\right).$$

Therefore, it only remains to show that $\mathcal{B}_t(\pi) = \pi$. Indeed, let $\tau \notin \pi$, then there exists $a \in A_k$ such that $\tau(a) > k$. Without loss of generality assume a is unique and let $\tau(a) = k + 1$. There must then be some $b \in A \setminus A_k$ with $\tau(b) \leq k$. Recall that the alternatives in $A \setminus A_k$ remain in fixed positions in π, and, again without loss of generality, suppose that $\sigma(b) = k + 1$ for all $\sigma \in \pi$. We wish to show that $d(\tau, \pi) > d(\sigma, \pi)$ for all $\sigma \in \pi$.

Let τ' be τ except that a and b are swapped, so $\tau'(a) = \tau(b)$ and $\tau'(b) = \tau(a)$. Observe that $\tau' \in \pi$ since a is unique. By definition, $d(\tau', \pi) = d(\sigma, \pi)$ for all $\sigma \in \pi$. It is therefore sufficient to show that $d(\tau', \pi) < d(\tau, \pi)$.

To this end, we partition the rankings $\sigma \in \pi \setminus \{\tau'\}$ into two sets, analyze them separately and in both cases show that $d(\tau', \sigma) \leq d(\tau, \sigma)$.

(1) $\sigma(a) \leq \tau'(a)$ (see Figure 1): In this case, we have that $|\sigma(a) - \tau(a)| \geq |\sigma(a) - \tau'(a)|$. Also, because σ and τ' agree on the position of $b \in A \setminus A_k$, $0 = |\sigma(b) - \tau'(b)| \leq |\sigma(b) - \tau(b)|$. We conclude that $d(\tau', \sigma) \leq d(\tau, \sigma)$.

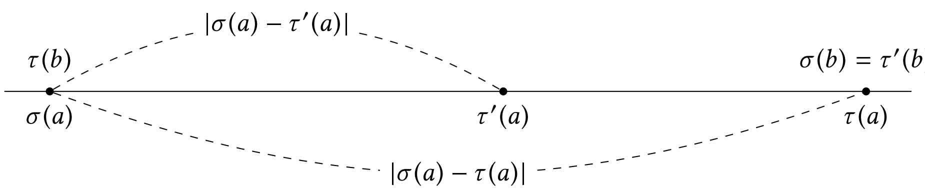

Fig. 1. Illustration of Case 1 of the proof of Theorem 4.10.

(2) $\sigma(a) > \tau'(a)$ (see Figure 2): It again holds that $0 = |\sigma(b) - \tau'(b)| \leq |\sigma(b) - \tau(b)|$, so if $d(\tau', \sigma) > d(\tau, \sigma)$ then $d(\tau', \sigma)$ is determined by a (i.e., a has the maximum displacement). Assume for contradiction that $d(\tau', \sigma) > d(\tau, \sigma)$. Then it holds that

$$d(\tau', \sigma) = |\sigma(a) - \tau'(a)| \leq |\sigma(a) - \tau'(a)| + |\sigma(a) - \tau(a)| = |\sigma(b) - \tau(b)| \leq d(\tau, \sigma),$$

a contradiction. We may conclude that $d(\tau', \sigma) \leq d(\tau, \sigma)$.

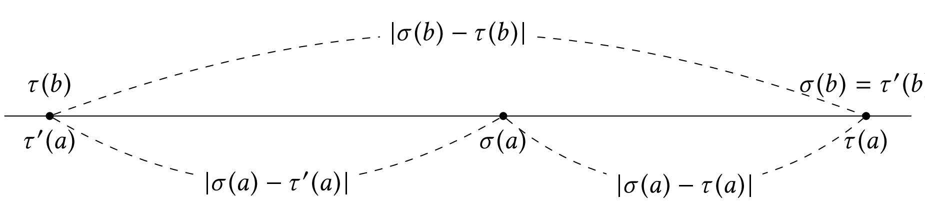

Fig. 2. Illustration of Case 2 of the proof of Theorem 4.10.

Since $d(\tau', \sigma) \leq d(\tau, \sigma)$ for all $\sigma \in \pi \setminus \{\tau'\}$, and $d(\tau', \tau') = 0 < d(\tau, \tau')$ we may conclude that $d(\tau, \pi) > d(\tau', \pi) = t$. It follows that $\mathcal{B}_t(\pi) = \pi$, thereby completing the proof. □

4.3 The Cayley Distance

In the previous sections, we have seen that our bounds are very different for different distance metrics. Still, all those bounds depended on t. By contrast, we establish a lower bound of $\Omega(\sqrt{m})$ on the average loss of any subset with $z = 1$ (i.e., the average position of any alternative) under the

Cayley distance. We view this as a striking negative result: Even if the votes are extremely accurate, i.e., t is very small, the ball $\mathcal{B}_t(\pi)$ could be such that the average position of any alternative is as large as $\Omega(\sqrt{m})$.

Theorem 4.11. *For $d = d_{CY}$ and every $k \in [\sqrt{m}/3]$, there exists $t = \Theta(k)$ and a vote profile π with*

$$n = k!\binom{\sqrt{m}}{k}^2$$

noisy rankings at average distance at most t from the ground truth, such that for any single alternative, its average position in the set of feasible ground truths $\mathcal{B}_t(\pi)$ is at least $\Omega(\sqrt{m})$.

The theorem's proof appears in Appendix B. Note that the delicate construction is specific to the case of $z = 1$. It remains open whether the theorem still holds when, say, $z = 2$, and, more generally, how the bound decreases as a function of z.

5 MAKING THE RIGHT DECISIONS, IN PRACTICE

We have two related goals in practice, to recover a ranking that is close to the ground truth, and identify a subset of alternatives with small loss in the ground truth. We compare the optimal rules that minimize the *average* distance or loss on $\mathcal{B}_t(\pi)$, denoted AVG^d, which we developed, to those that minimize the *maximum* distance or loss, denoted MAX^d, which were developed by Procaccia et al. [2016]. Importantly, at least for the case where the output is a ranking, Procaccia et al. [2016] have compared their methods against a slew of previously studied methods — including MLE rules for famous random noise models like the one due to Mallows [1957] — and found theirs to be superior. In addition, their methods are the ones currently used in practice, on RoboVote. Therefore we focus on comparing our methods to theirs.

Datasets. Like Procaccia et al. [2016], we make use of two real-world datasets collected by Mao et al. [2013]. In both of these datasets — *dots* and *puzzle* — the ground truth rankings are known, and data was collected via Amazon Mechanical Turk. Dataset *dots* was obtained by asking workers to rank four images containing different numbers of dots in increasing order. Dataset *puzzle* was obtained by asking workers to rank four different states of a puzzle according to the minimal number of moves necessary to reach the goal state. Each dataset consists of four different noise levels, corresponding to levels of difficulty, represented using a single noise parameter. In *dots*, higher noise corresponds to smaller differences between the number of dots in the images, whereas in *puzzle*, higher noise entails ranking states that are all a constant number of steps further from the goal state. Overall the two datasets contain thousands of votes — 6367, to be precise.

Experimental design. When recovering complete rankings, the evaluation metric is the distance of the returned ranking to the *actual* (known) ground truth. We reiterate that, although MAX^d is designed to minimize the maximum distance to any feasible ground truth given an input profile π and an estimate of the average noise t, that is, it is designed for the worst case, it is known to work well in practice [Procaccia et al., 2016]. Similarly, AVG^d is designed to optimize the average distance to the set of feasible ground truths; our experiments will determine whether this is a useful proxy for minimizing the distance to an unknown ground truth.

When selecting a subset of alternatives, the evaluation metric is the loss of that subset in the *actual* ground truth. As discussed above, the current implementation of RoboVote uses the rule MAX^d that returns the set of alternatives that minimizes the maximum loss in any feasible true ranking. As in the complete ranking setting, the rule AVG^d returns the set of alternatives that minimizes the average loss over the feasible true rankings.

It is important to emphasize that in both these settings, MAX^d and AVG^d optimize an objective over the set of feasible ground truths, but are evaluated on the actual known ground truth. It is therefore impossible to predict in advance which of the methods will perform best.

Our theoretical results assume that an upper bound t on the average error is given to us, and our guarantees depend on this bound. In practice, though, t has to be estimated. For example, the current RoboVote implementation uses $t_{RV} = \min_{\sigma \in \pi} d(\sigma, \pi)/|\pi|$, or the minimum average distance from one ranking in π to all other rankings in π.

In our experiments, we wish to study the impact of the choice of t on the performance of AVG^d and MAX^d. A natural choice is $t^* \triangleq d(\pi, \sigma^*)$, where π is the vote profile and σ^* is the actual ground truth. That is, t^* is the average distance between the vote profile and the actual ground truth. In principle it is an especially good choice because it induces the smallest ball $\mathcal{B}_t(\pi)$ that contains the actual ground truth. However, it is also an impractical choice, because one cannot compute this value without knowing the ground truth. We also consider

$$t_{KEM} \triangleq \min_{\sigma \in \mathcal{L}(A)} d(\sigma, \pi)$$

(named after the *Kemeny* rule) — the minimum possible distance between the vote profile and any ranking.

In order to synchronize results across different profiles, we let $\hat{t}$ be the estimate of t that we feed into the methods, and define

$$r = \frac{\hat{t} - t_{KEM}}{t^* - t_{KEM}}.$$

Note that because t_{KEM} is the minimum value that allows for a nonempty set of feasible ground truths, we know that $t^* - t_{KEM} \geq 0$. For any profile, $r = 0$ implies that $\hat{t} = t_{KEM}$, $r < 1$ implies that $\hat{t} < t^*$, $r = 1$ implies that $\hat{t} = t^*$, and $r > 1$ implies that $\hat{t} > t^*$. In our experiments, as in the work of Procaccia et al. [2016], we use $r \in [0, 2]$.

Results and their interpretation. Our results for three output types — ranking, subset with $z = 1$ (single winner), and subset with $z = 2$ — can be found in Figures 3, 4, and 5, respectively. Each has two subfigures, for the KT distance, and the Cayley distance. All Figures show r on the x axis. In Figure 3, the y axis shows the distance between the output ranking and the actual ground truth.

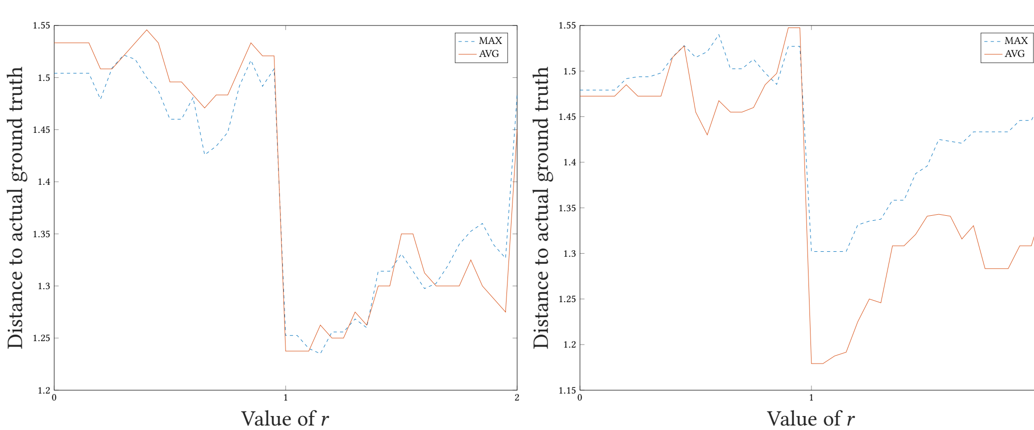

(a) Kendall tau (b) Cayley

Fig. 3. Dots dataset (noise level 4), ranking output.

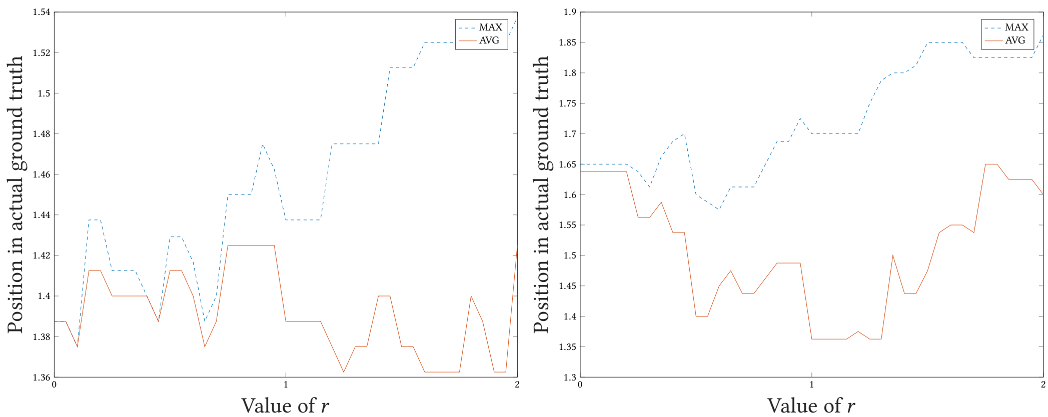

(a) Kendall tau (b) Cayley

Fig. 4. Dots dataset (noise level 4), subset output with $z = 1$.

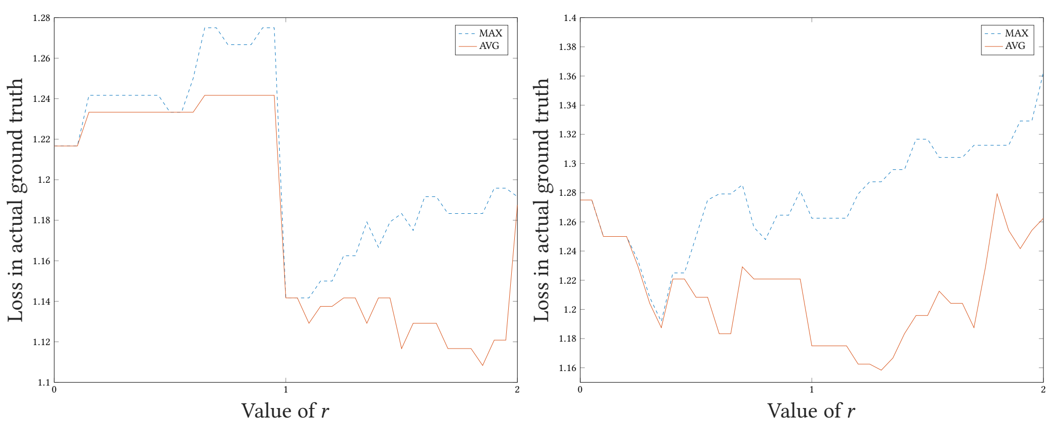

(a) Kendall tau (b) Cayley

Fig. 5. Dots dataset (noise level 4), subset output with $z = 2$.

In Figures 4 and 5, the y axis shows the loss of the selected subset on the actual ground truth. All figures are based on the dots dataset with the highest noise level (4). The results for the puzzle dataset are similar (albeit not as crisp), and the results for different noise levels are quite similar. The results differ across distance functions, but the conclusions below apply to all four, not just the two that are shown here. Additional figures can be found in appendix C.

It is interesting to note that, while in Figure 3 the accuracy of each distance metric is measured using that metric (i.e., KT is measured with KT and Cayley with Cayley), in the other two figures the two distances are measured in the exact same way: based on position or loss in the ground truth. Despite the dismal theoretical results for Cayley (Theorem 4.11), its performance in practice is comparable to KT.

More importantly, we see that although MAX^d and AVG^d perform similarly on low values of r, AVG^d significantly outperforms MAX^d on medium and high values of r, and especially when $r > 1$, that is, $\hat{t} > t^*$. This is true in all cases (including the two distance metrics that are not shown), except for the ranking output type under the KT distance (Figure 3a) and the footrule distance (Appendix C), where the performance of the two methods is almost identical across the board (values of r, datasets, and noise levels).

These results match our intuition. As r increases so does $\hat{t}$, and the set $\mathcal{B}_{\hat{t}}(\pi)$ grows larger. When this set is large, the conservatism of MAX^d becomes a liability, as it minimizes the maximum distance with respect to rankings that are unlikely to coincide with the actual ground truth. By contrast, AVG^d is more robust: It takes the new rankings into account, but does not allow them to dictate its output.

The practical implication is clear. Because we do not have a way of divining t^*, which is often the most effective choice in practice, we resort to relatively crude estimates, such as the deployed choice of t_{RV} discussed above. Moreover, underestimating t^* is often risky, as the results show, because the ball $\mathcal{B}_{\hat{t}}(\pi)$ does not contain the actual ground truth when $\hat{t} < t^*$. Therefore in practice we try to aim for estimates such that $\hat{t} > t^*$, and robustness to the value of $\hat{t}$ is crucial. In this sense AVG^d is a better choice than MAX^d.

6 DISCUSSION

We wrap up with a brief discussion of several key points.

Non-uniform distributions. All of our upper bound results, namely Theorems 3.2, 4.5, 4.7, and 4.9, apply to any distribution over $\mathcal{B}_t(\pi)$, not just the uniform distribution (when replacing "average" distance/loss with "expected" distance/loss). To see why this is true for the latter three theorems, note that their proofs construct a randomized rule and leverage Lemma 4.2, which easily extends to any distribution. While this is a nice point to make, we do not believe that non-uniform distributions are especially well motivated — where would such a distribution come from? By contrast, the uniform distribution represents an agnostic viewpoint.

Computational complexity. We have not paid much attention to computational complexity. In our experiments there are only four alternatives, so we can easily compute $\mathcal{B}_t(\pi)$ by enumeration. For real-world instances, integer programming is used, as we briefly discussed in Section 1.1. While those implementations are for rules that minimize the *maximum* distance or loss over $\mathcal{B}_t(\pi)$ [Procaccia et al., 2016], they can be easily modified to minimize the average distance or loss. Therefore, at least for the purposes of applications like RoboVote, computational complexity is *not* an obstacle.

Real-world implications. As noted in Section 5, our empirical results suggest that minimizing the average distance or loss has a significant advantage in practice over minimizing the maximum distance or loss. We are therefore planning to continue refining our methods, and ultimately deploy them on RoboVote, where they will influence the way thousands of people around the world make group decisions.

ACKNOWLEDGMENTS

Kahng and Procaccia were partially supported by the National Science Foundation under grants IIS-1350598 and CCF-1525932, by the Office of Naval Research, and by a Sloan Research Fellowship.

REFERENCES

H. Azari Soufiani, D. C. Parkes, and L. Xia. 2012. Random Utility Theory for Social Choice. In *Proceedings of the 26th Annual Conference on Neural Information Processing Systems (NIPS)*. 126–134.

H. Azari Soufiani, D. C. Parkes, and L. Xia. 2013. Preference Elicitation For General Random Utility Models. In *Proceedings of the 29th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 596–605.

H. Azari Soufiani, D. C. Parkes, and L. Xia. 2014. Computing Parametric Ranking Models via Rank-Breaking. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*. 360–368.

F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. 2017. Subset Selection Via Implicit Utilitarian Voting. *Journal of Artificial Intelligence Research* 58 (2017), 123–152.

I. Caragiannis, A. D. Procaccia, and N. Shah. 2014. Modal Ranking: A Uniquely Robust Voting Rule. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*. 616–622.

I. Caragiannis, A. D. Procaccia, and N. Shah. 2016. When Do Noisy Votes Reveal the Truth? *ACM Transactions on Economics and Computation* 4, 3 (2016), article 15.

J. R. Chamberlin and P. N. Courant. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *American Political Science Review* 77, 3 (1983), 718–733.

Y. Chen, A. Ghosh, M. Kearns, T. Roughgarden, and J. Wortman Vaughan. 2016. Mathematical Foundations for Social Computing. *Communications of the ACM* 59, 12 (2016), 102–108.

V. Conitzer, M. Rognlie, and L. Xia. 2009. Preference Functions That Score Rankings and Maximum Likelihood Estimation. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*. 109–115.

V. Conitzer and T. Sandholm. 2005. Common Voting Rules as Maximum Likelihood Estimators. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 145–152.

P. Diaconis and R. L. Graham. 1977. Spearman's footrule as a measure of disarray. *Journal of the Royal Statistical Society: Series B* 39, 2 (1977), 262–268.

E. Elkind, P. Faliszewski, and A. Slinko. 2010. Good Rationalizations of Voting Rules. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*. 774–779.

E. Elkind and N. Shah. 2014. Electing the Most Probable Without Eliminating the Irrational: Voting Over Intransitive Domains. In *Proceedings of the 30th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 182–191.

T. Lu and C. Boutilier. 2011a. Budgeted Social Choice: From Consensus to Personalized Decision Making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*. 280–286.

T. Lu and C. Boutilier. 2011b. Robust Approximation and Incremental Elicitation in Voting Protocols. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*. 287–293.

C. L. Mallows. 1957. Non-null ranking models. *Biometrika* 44 (1957), 114–130.

A. Mao, A. D. Procaccia, and Y. Chen. 2013. Better Human Computation Through Principled Voting. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*. 1142–1148.

B. L. Monroe. 1995. Fully Proportional Representation. *American Political Science Review* 89, 4 (1995), 925–940.

G. Pritchard and M. Wilson. 2009. Asymptotics of the minimum manipulating coalition size for positional voting rules under impartial culture behaviour. *Mathematical Social Sciences* 58, 1 (2009), 35–57.

A. D. Procaccia. 2016. Science Can Restore America's Faith in Democracy. (2016). https://www.wired.com/2016/12/science-can-restore-americas-faith-democracy/

A. D. Procaccia, S. J. Reddi, and N. Shah. 2012. A Maximum Likelihood Approach For Selecting Sets of Alternatives. In *Proceedings of the 28th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 695–704.

A. D. Procaccia, J. S. Rosenschein, and A. Zohar. 2008. On the Complexity of Achieveing Proportional Representation. *Social Choice and Welfare* 30, 3 (2008), 353–362.

A. D. Procaccia, N. Shah, and Y. Zick. 2016. Voting rules as error-correcting codes. *Artificial Intelligence* 231 (2016), 1–16.

I. Tsetlin, M. Regenwetter, and B. Grofman. 2003. The impartial culture maximizes the probability of majority cycles. *Social Choice and Welfare* 21 (2003), 387–398.

L. Xia. 2016. Bayesian Estimators As Voting Rules. In *Proceedings of the 32nd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*.

L. Xia and V. Conitzer. 2011. A maximum likelihood approach towards aggregating partial orders. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*. 446–451.

L. Xia, V. Conitzer, and J. Lang. 2010. Aggregating preferences in multi-issue domains by using maximum likelihood estimators. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 399–408.

A. C. Yao. 1977. Probabilistic Computations: Towards a Unified Measure of Complexity. In *Proceedings of the 17th Symposium on Foundations of Computer Science (FOCS)*. 222–227.

H. P. Young. 1988. Condorcet's theory of voting. *The American Political Science Review* 82, 4 (1988), 1231–1244.

A PROOF OF THEOREM 3.3

To prove the lower bounds we will make use of several technical lemmas. The next three lemmas were established by Procaccia et al. [2016, Theorem 5].

Lemma A.1. *For $d = d_{KT}$ and $t \le (m/12)^2$, there exists a partition of A into A_1, A_2, A_3, A_4, and a vote profile consisting of $n/2$ copies of each of the rankings*

$$\sigma = \sigma^{A_1} \succ \sigma^{A_2} \succ \sigma^{A_3} \succ \sigma^{A_4}$$
$$\sigma' = \sigma^{A_1}_{rev} \succ \sigma^{A_2}_{rev} \succ \sigma^{A_3}_{rev} \succ \sigma^{A_4},$$

for which $\mathcal{B}_t(\pi) = \mathcal{F}(A_1 \succ A_2 \succ A_3 \succ \sigma^{A_4})$ and $\lfloor 2t \rfloor = \sum_{i=1}^{3} \binom{m_i}{2}$, where $m_i \triangleq |A_i|$ for $i \in [4]$.

Lemma A.2. *For $d = d_{FR}$ and $t \le (m/8)^2$, there exists a partition of A into A_1, A_2, A_3, A_4, and A_5, and a vote profile $\pi \in \mathcal{L}(A)^n$ consisting of $n/2$ copies of each of the following rankings,*

$$\sigma = \sigma^{A_1} \succ \sigma^{A_2} \succ \sigma^{A_3} \succ \sigma^{A_4} \succ \sigma^{A_5}$$
$$\sigma' = \sigma^{A_1}_{rev} \succ \sigma^{A_2}_{rev} \succ \sigma^{A_3}_{rev} \succ \sigma^{A_4}_{rev} \succ \sigma^{A_5},$$

for which

$$\mathcal{B}_t(\pi) = \Big\{\rho \in \mathcal{L}(A) \mid \{\rho(a_i^j), \rho(a_i^{2m_i+1-j})\} = \{\sigma(a_i^j), \sigma(a_i^{2m_i+1-j})\} \text{ for } i \in [4], j \in [2m_i], \text{ and}$$
$$\rho(a_5^j) = \sigma(a_5^j) = \sigma'(a_5^j) \text{ for } j \in [m_5]\Big\},$$

where $2m_i = |A_i|$ for $i \in [4]$, $m_5 = |A_5|$, and

$$d^{\downarrow}_{FR}(2t) = \sum_{i=1}^{4} \left\lfloor \frac{(2m_i)^2}{2} \right\rfloor.$$

Lemma A.3. *For $d = d_{CY}$ and t such that $2\lfloor 2t \rfloor \le m$, there exists a vote profile π consisting of $n/2$ copies of each of the following rankings,*

$$\sigma = (a_1 \succ \cdots \succ a_{2\lfloor 2t \rfloor} \succ a_{2\lfloor 2t \rfloor+1} \succ \cdots \succ a_m)$$
$$\sigma' = (a_{2\lfloor 2t \rfloor} \succ \cdots \succ a_1 \succ a_{2\lfloor 2t \rfloor+1} \succ \cdots \succ a_m),$$

for which

$$\mathcal{B}_t(\pi) = \{\rho \in \mathcal{L}(A) | \{\rho(a_i), \rho(a_{2\lfloor 2t \rfloor+1-i})\} = \{i, 2\lfloor 2t \rfloor + 1 - i\} \text{ for } i \in [\lfloor 2t \rfloor], \text{ and}$$
$$\rho(a_i) = i \text{ for } i > 2\lfloor 2t \rfloor\}.$$

We will need a similar result for maximum displacement.

Lemma A.4. *For $d = d_{MD}$ and t such that $2\lfloor 2t \rfloor \le m$, there exists a vote profile π consisting of $n/2$ copies of each of the following rankings,*

$$\sigma = (a_1 \succ \cdots \succ a_{\lfloor 2t \rfloor}) \succ (a_{\lfloor 2t \rfloor+1} \succ \cdots \succ a_{2\lfloor 2t \rfloor}) \succ \sigma^{A'}$$
$$\sigma' = (a_{\lfloor 2t \rfloor+1} \succ \cdots \succ a_{2\lfloor 2t \rfloor}) \succ (a_1 \succ \cdots \succ a_{\lfloor 2t \rfloor}) \succ \sigma^{A'},$$

where $A' = A \setminus \{a_1, \ldots, a_{2\lfloor 2t \rfloor}\}$, for which $\mathcal{B}_t(\pi) = \{\sigma, \sigma'\}$.

Proof. It is easy to see that $\sigma \in \mathcal{B}_t(\pi)$ and $\sigma' \in \mathcal{B}_t(\pi)$, as $d(\sigma, \sigma') = \lfloor 2t \rfloor$. We therefore need to show that $\mathcal{B}_t(\pi)$ does not contain any other rankings.

Let $\rho \in \mathcal{B}_t(\pi)$, and consider its first-ranked alternative, $a = \rho^{-1}(1)$. It holds that $\sigma(a) \ge \lfloor 2t \rfloor + 1$ or $\sigma'(a) \ge \lfloor 2t \rfloor + 1$, because the two rankings place disjoint subsets of alternatives in the first $\lfloor 2t \rfloor$ positions. Suppose first that the former inequality holds; then

$$d(\rho, \sigma) \ge \sigma(a) - \rho(a) \ge \lfloor 2t \rfloor.$$

If $\rho \neq \sigma'$ then $d(\rho, \sigma') \geq 1$, and therefore

$$d(\rho, \mathcal{B}_t(\pi)) = \frac{d(\rho, \sigma) + d(\rho, \sigma')}{2} \geq \frac{\lfloor 2t \rfloor + 1}{2} > t.$$

It follows that $\rho = \sigma'$. Similarly, if the latter inequality holds, then $\rho = \sigma$. □

We are now in a position to prove Theorem 3.3.

Proof of Theorem 3.3. We address each of the four distance metrics separately.

The Kendall tau distance. Let π and $\mathcal{B}_t(\pi)$ have the structure specified in Lemma A.1. For all $\rho \in \mathcal{L}(A)$ and $i \in [3]$, and every pair of alternatives $a \in A_i$, $b \in A_i \setminus \{a\}$, we can divide the rankings in $\mathcal{B}_t(\pi)$ into pairs that are identical except for swapping a and b. Note that for each pair, one ranking agrees with ρ on a and b, and one does not. Therefore,

$$d(\rho, \mathcal{B}_t(\pi)) \geq \frac{\sum_{i=1}^{3} \binom{m_i}{2}}{2} = \frac{\lfloor 2t \rfloor}{2} \geq \frac{d^{\downarrow}(2t)}{2}.$$

The footrule distance. Let π and $\mathcal{B}_t(\pi)$ have the structure specified in Lemma A.2. For all $\rho \in \mathcal{L}(A)$ and $i \in [4]$, and for every alternative $a_i^j \in A_i$, we can divide the rankings in $\mathcal{B}_t(\pi)$ into pairs that are identical except for swapping a_i^j and $a_i^{2m_i+1-j}$. Note that for each such pair σ and σ', $|\sigma(a_i^j) - \sigma'(a_i^j)| = 2m_i + 1 - 2j$, and using the triangle inequality,

$$|\rho(a_i^j) - \sigma(a_i^j)| + |\rho(a_i^j) - \sigma'(a_i^j)| \geq 2m_i + 1 - 2j.$$

Furthermore, by the structure of $\mathcal{B}_t(\pi)$, we know that

$$\sum_{j=1}^{2m_i} 2m_i + 1 - 2j = \left\lfloor \frac{(2m_i)^2}{2} \right\rfloor.$$

By summing over all $j \in [2m_i]$ and $i \in [4]$, we get

$$d(\rho, \mathcal{B}_t(\pi)) \geq \frac{\sum_{i=1}^{4} \sum_{j=1}^{2m_i} 2m_i + 1 - 2j}{2} = \frac{\sum_{i=1}^{4} \left\lfloor \frac{(2m_i)^2}{2} \right\rfloor}{2} = \frac{d^{\downarrow}(2t)}{2}.$$

The Cayley distance. Let π and $\mathcal{B}_t(\pi)$ have the structure specified in Lemma A.3. For all $\rho \in \mathcal{L}(A)$, and every pair of alternatives $\{a_i, a_{2\lfloor 2t \rfloor+1-i}\}$ for $i \in [\lfloor 2t \rfloor]$, we can divide the rankings in $\mathcal{B}_t(\pi)$ into pairs τ_i and τ_i' that are identical except for swapping a and b. Note that for each pair, one ranking agrees with ρ on a and b, and one does not. Since each swap places at most two alternatives in their correct positions, each of the $\lfloor 2t \rfloor$ pairs adds at least $1/2$ to $d(\rho, \mathcal{B}_t(\pi))$ because $d(\rho, \tau_i) + d(\rho, \tau_i') \geq 1$. Overall we have

$$d(\rho, \mathcal{B}_t(\pi)) \geq \frac{\lfloor 2t \rfloor}{2} \geq \frac{d^{\downarrow}(2t)}{2}.$$

The maximum displacement distance. Let π and $\mathcal{B}_t(\pi)$ have the structure specified in Lemma A.4. Consider any ranking $\rho \in \mathcal{L}(A)$. Let $a \in A$ be the alternative ranked first in ρ, i.e., $a = \rho^{-1}(1)$. If $a \in \{a_1, \ldots, a_{\lfloor 2t \rfloor}\}$, then $d(\rho, \sigma') \geq \lfloor 2t \rfloor$. Similarly, if $a \in \{a_{2t+1}, \ldots, a_{2\lfloor 2t \rfloor}\}$ then $d(\rho, \sigma) \geq \lfloor 2t \rfloor$. Therefore,

$$d(\rho, \mathcal{B}_t(\pi)) = \frac{d(\rho, \sigma) + d(\rho, \sigma')}{2} \geq \frac{\lfloor 2t \rfloor}{2} \geq \frac{d^{\downarrow}(2t)}{2}.$$

□

B PROOF OF THEOREM 4.11

Suppose for ease of exposition that $\sqrt{m} \in \mathbb{Z}$. Let $\sigma = (a_1 \succ a_2 \succ \ldots \succ a_m)$ be a ranking and let $L = \{1, 2, \ldots, \sqrt{m}\}$, $M = \{\sqrt{m}+1, \ldots, m-\sqrt{m}\}$ and $R = \{m-\sqrt{m}+1, \ldots, m\}$. Define the ranking σ_{ij} for $i \in L, j \in R$ to have $\sigma_{ij}(a_i) = \sigma(a_j)$ and $\sigma_{ij}(a_j) = \sigma(a_i)$ while $\sigma_{ij}(a_c) = \sigma(a_c)$ for all $c \in [m] \setminus \{i, j\}$. In other words, σ_{ij} is exactly σ with element $i \in L$ and element $j \in R$ swapped.

Construct a vote in π by selecting $S \subseteq L$, $T \subseteq R$ with $|S| = |T| = k$, then selecting a perfect matching $M : S \to T$, and finally swapping each a_i for $i \in S$ with a_j for $j = M(i)$. We have such a vote for every choice of S and T, and every perfect matching between them. This results in a vote profile of cardinality

$$n = |\pi| = k!\binom{\sqrt{m}}{k}^2.$$

Let $t = k + 1 - \frac{2k}{m}$. By construction $d(\tau, \sigma) = k$ for all $\tau \in \pi$. It follows that $d(\pi, \sigma) = k \le t$, and therefore $\sigma \in \mathcal{B}_t(\pi)$.

We next claim that

$$d(\sigma_{ij}, \pi) \le k + 1 - \frac{2k}{m} = t.$$

It suffices to consider two classes of rankings $\tau \in \pi$. First, if $\tau(a_i) = j = \sigma_{ij}(a_i)$ and $\tau(a_j) = i = \sigma_{ij}(a_j)$, then $d(\sigma_{ij}, \tau) \le k - 1$, since reversing the other $k - 1$ pairwise swaps changes τ into σ_{ij}. There are

$$\hat{n} = \binom{\sqrt{m}-1}{k-1}^2 \cdot (k-1)!$$

such rankings in π. Second, for all other $\tau \in \pi$, we have $d(\sigma_{ij}, \tau) \le k + 1$, since it is always possible to reverse the k pairwise exchanges that changed σ into $\tau \in \pi$, and then perform one additional exchange to put a_i and a_j into the correct positions. It follows that for all $i \in L, j \in R$,

$$\begin{aligned} d_{CY}(\sigma_{ij}, \pi) &\le \frac{1}{|\pi|}(k-1)\hat{n} + \frac{1}{|\pi|}(k+1)(|\pi| - \hat{n}) \\ &= (k+1) + \frac{(k-1)\hat{n} - (k+1)\hat{n}}{|\pi|} = (k+1) - \frac{2\hat{n}}{|\pi|} \\ &= (k+1) - 2 \cdot \frac{\binom{\sqrt{m}-1}{k-1}^2 \cdot (k-1)!}{|\pi|} = k + 1 - \frac{2k}{m}. \end{aligned}$$

We conclude that $\{\sigma\} \cup \{\sigma_{ij} : i \in L, j \in R\} \subseteq \mathcal{B}_t(\pi)$.

We next show that this, in fact, fully describes $\mathcal{B}_t(\pi)$. To show this, we must use the Hamming distance, denoted d_{HM}, which is defined as the number of positions at which two rankings of the same length differ. In particular, we use the relationship $d_{CY}(\tau, \tau') \ge \frac{1}{2} d_{HM}(\tau, \tau')$ between the Cayley and Hamming distance metrics for all $\tau, \tau' \in \mathcal{L}(A)$. This is a direct result of the fact that a single swap can place at most two alternatives in their correct positions.

For an arbitrary $\tau' \in \mathcal{L}(A)$ we can decompose the Hamming distance metric as

$$\begin{aligned} d_{HM}(\tau', \pi) &= \frac{1}{|\pi|}\sum_{\tau \in \pi} d_{HM}(\tau', \tau) = \frac{1}{|\pi|}\sum_{\tau \in \pi}\sum_{i \in [m]} \mathbb{I}[\tau(a_i) \ne \tau'(a_i)] \\ &= \sum_{i \in [m]} \frac{1}{|\pi|}\sum_{\tau \in \pi} \mathbb{I}[\tau(a_i) \ne \tau'(a_i)] = \sum_{i \in [m]} q_i(\pi, \tau'), \end{aligned} \tag{4}$$

where

$$q_i(\pi, \tau') \triangleq \frac{1}{|\pi|}\sum_{\tau \in \pi} \mathbb{I}[\tau(a_i) \ne \tau'(a_i)]$$

is the average *penalty* that a_i incurs in τ' with respect to π under the Hamming distance metric.

Consider $q_i(\pi, \tau')$ for $i \in L$. If $\tau'(a_i) = i$, then $q_i(\pi, \tau') = k/\sqrt{m}$ since a_i is swapped with an alternative in the right endpoint in a $k/\sqrt{m}$ fraction of the rankings in π. If $\tau'(a_i) \in (L \setminus \{i\}) \cup M$, then a penalty is incurred in every $\tau \in \pi$, so $q_i(\pi, \tau') = 1$. If $\tau'(a_i) \in R$, then $q_i(\pi, \tau') = 1 - (k/\sqrt{m})(1/\sqrt{m}) = 1 - k/m$. The analysis for $q_i(\pi, \tau')$, $i \in R$ is identical. For $q_i(\pi, \tau')$, $i \in M$, observe that $\tau(a_i) = i$ for all $\tau \in \pi$, so $q_i(\pi, \tau') = 0$ if $\tau'(a_i) = i$ and 1 otherwise.

It is clear from the decomposition and above discussion that $\tau' = \sigma$ is the unique ranking minimizing $d_{HM}(\tau', \pi)$. We partition the rankings $\tau' \in \mathcal{L}(A)$ according to their Hamming distance from σ and analyze which rankings can appear in $\mathcal{B}_t(\pi)$.

(1) $d_{HM}(\tau', \sigma) = 1$: The Hamming distance metric does not allow rankings at distance 1 from each other.
(2) $d_{HM}(\tau', \sigma) = 2$: We have shown that $\sigma_{ij} \in \mathcal{B}_t(\pi)$. If $\tau' \notin \{\sigma_{ij} : i \in L, j \in R\}$, then $d(\tau', \tau) = k + 1$ for all $\tau \in \pi$ and thus $\tau' \notin \mathcal{B}_t(\pi)$. This is because the Cayley distance between σ and any $\tau \in \pi$ is exactly k due to the k pairwise disjoint swaps described above, and τ' involves an additional swap that is not allowed when transforming σ into $\tau \in \pi$.
(3) $d_{HM}(\tau', \sigma) \geq 3$: For every ranking $\tau' \in \mathcal{L}(A)$ at Hamming distance at least 3 from σ, it holds that $\tau'(a_i) \neq i$ for at least three values of i, and therefore at least three of the penalties in Equation (4) are not minimal, meaning that they are at least $1 - k/m$. Moreover, the minimal penalty for $i \in L \cup R$ is $k/\sqrt{m}$. It follows that

$$\begin{aligned} d_{CY}(\tau', \pi) &\geq \frac{1}{2} d_{HM}(\tau', \pi) \\ &\geq \frac{1}{2}\left[\frac{k}{\sqrt{m}}(2\sqrt{m} - 3) + 3\left(1 - \frac{k}{m}\right)\right] \\ &= k + \frac{3}{2} - \frac{3k}{2m} - \frac{3k}{2\sqrt{m}} \\ &= k + 1 - \frac{2k}{m} + \left(\frac{1}{2} + \frac{k}{2m} - \frac{3k}{2\sqrt{m}}\right) \\ &\geq k + 1 - \frac{2k}{m} + \left(\frac{1}{2} + \frac{k}{2m} - \frac{1}{2}\right) \\ &= k + 1 - \frac{2k}{m} + \frac{k}{2m} > k + 1 - \frac{2k}{m}, \end{aligned}$$

where the fifth transition follows from the assumption that $k \leq \sqrt{m}/3$.

We conclude that $\mathcal{B}_t(\pi) = \{\sigma\} \cup \{\sigma_{ij} : i \in L, j \in R\}$ and thus that $|\mathcal{B}_t(\pi)| = m + 1$.

To complete the proof, we show that every alternative has average position at least $\Omega(\sqrt{m})$ in $\mathcal{B}_t(\pi)$. For every a_i with $i \in L$, a_i appears in position $j \in R$ in $\sqrt{m}$ of the $m + 1$ rankings in $\mathcal{B}_t(\pi)$. Therefore the average loss of a_i over $\mathcal{B}_t(\pi)$ is at least

$$\frac{m + 1 - \sqrt{m}}{m + 1} \cdot 1 + \frac{\sqrt{m}}{m + 1} \cdot \frac{m}{2} = \Omega(\sqrt{m}).$$

For $i \in M$, alternative a_i never appears in position smaller than $\sqrt{m} + 1$ in $\mathcal{B}_t(\pi)$ and clearly has average position $\Omega(\sqrt{m})$. Finally, for $j \in R$, alternative a_j appears in position j in at least $m + 1 - \sqrt{m}$ of the rankings in $\mathcal{B}_t(\pi)$, and also has average position at least $\Omega(\sqrt{m})$. □

C ADDITIONAL EXPERIMENTAL RESULTS

We provide additional evidence that our experimental results of Section 5 do not depend on any particular distance metric, dataset, or noise level. Specifically, the results for the footrule and maximum displacement distance metrics (instead of KT and Cayley) under noise level 3 (instead of 4) of the puzzle dataset (instead of dots) when returning a complete ranking are presented in Figure 6, and the results for returning a subset of size 1 and 2 in Figures 7 and 8, respectively.

Although the results obtained from the puzzle dataset are somewhat noisier in general, it does still hold that AVG^d is more robust than MAX^d to overestimates of t^*, as we concluded in Section 5 (with the exception of Figure 6a, as noted there).

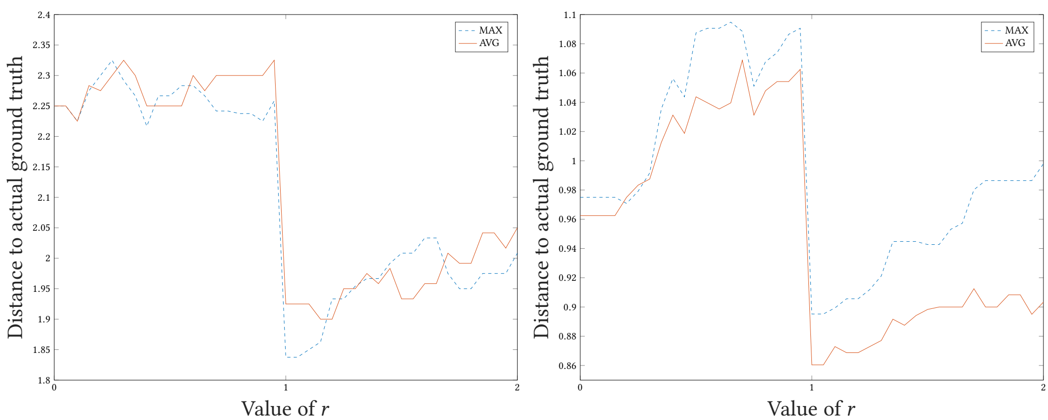

(a) Footrule (b) Maximum displacement

Fig. 6. Puzzle dataset (noise level 3), ranking output.

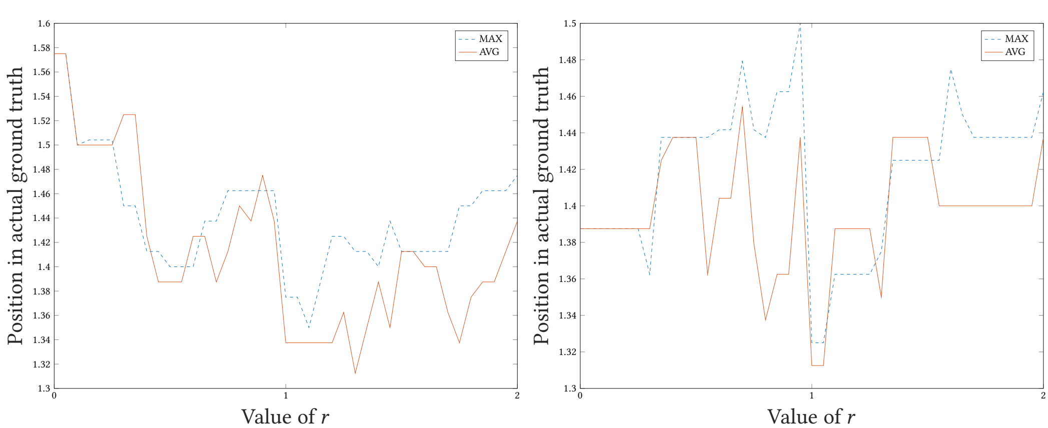

(a) Footrule (b) Maximum displacement

Fig. 7. Puzzle dataset (noise level 3), subset output with $z = 1$.

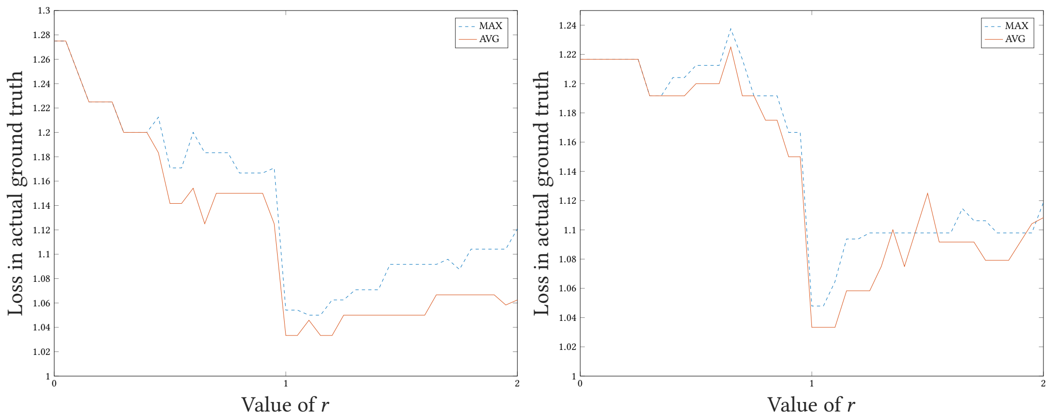

(a) Footrule

(b) Maximum displacement

Fig. 8. Puzzle dataset (noise level 3), subset output with $z = 2$.