Metric-Distortion Bounds under Limited Information

Ioannis Anagnostides

National Technical University of Athens

`ioannis.anagnostides@gmail.com`

Dimitris Fotakis

National Technical University of Athens

`fotakis@cs.ntua.gr`

Panagiotis Patsilinakos

National Technical University of Athens

`patsilinak@corelab.ntua.gr`

July 7, 2021

Abstract

In this work we study the metric distortion problem in voting theory under a limited amount of ordinal information. Our primary contribution is threefold. First, we consider mechanisms which perform a sequence of pairwise comparisons between candidates. We show that a widely-popular deterministic mechanism employed in most knockout phases yields distortion $\mathcal{O}(\log m)$ while eliciting only $m-1$ out of $\Theta(m^2)$ possible pairwise comparisons, where m represents the number of candidates. Our analysis for this mechanism leverages a powerful technical lemma recently developed by Kempe [Kem20a]. We also provide a matching lower bound on its distortion. In contrast, we prove that any mechanism which performs fewer than $m-1$ pairwise comparisons is destined to have unbounded distortion. Moreover, we study the power of deterministic mechanisms under incomplete rankings. Most notably, when every agent provides her k-top preferences we show an upper bound of $6m/k + 1$ on the distortion, for any $k \in \{1, 2, \ldots, m\}$. Thus, we substantially improve over the previous bound of $12m/k$ recently established by Kempe [Kem20a, Kem20b], and we come closer to matching the best-known lower bound. Finally, we are concerned with the sample complexity required to ensure near-optimal distortion with high probability. Our main contribution is to show that a random sample of $\Theta(m/\epsilon^2)$ voters suffices to guarantee distortion $3 + \epsilon$ with high probability, for any sufficiently small $\epsilon > 0$. This result is based on analyzing the sensitivity of the deterministic mechanism introduced by Gkatzelis, Halpern, and Shah [GHS20]. Importantly, all of our sample-complexity bounds are distribution-independent.

From an experimental standpoint we present several empirical findings on real-life voting applications, comparing the scoring systems employed in practice with a mechanism explicitly minimizing (metric) distortion. Interestingly, for our case studies we find that the winner in the actual competition is typically the candidate who minimizes the distortion.

1   Introduction

Aggregating the preferences of individuals into a collective decision lies at the heart of social choice, and has recently found numerous applications in areas such as information retrieval, recommender systems, and machine learning [VZ14, VLZ12, BCJ11, FH10, LN05]. According to the classic theory of Von Neumann and Morgenstern [vNM44] individual preferences are captured through a *utility function*, which assigns numerical (or *cardinal*) values to each alternative. Yet, in voting theory, as well as in most practical applications, mechanisms typically elicit only *ordinal* information from the voters, indicating an order of preferences over the candidates. Although this might seem at odds with a utilitarian framework, it has been recognized that it might be hard for a voter to specify a precise numerical value for an alternative, and providing only ordinal information substantially reduces the cognitive burden. This begs the question: What is the loss in efficiency of a mechanism extracting only ordinal information with respect to the *utilitarian social welfare*, i.e. the sum of individual utilities over a chosen candidate? The framework of *distortion* introduced by Procaccia and Rosenschein [PR06] measures exactly this loss from an approximation-algorithms standpoint, and has received considerable attention in recent years.

As it turns out, the approximation-guarantees we can hope for crucially depend on the assumptions we make on the utility functions. For example, in the absence of any structure Procaccia and Rosenschein observed that *no* ordinal mechanism can obtain bounded distortion [PR06]. In this work we focus on the *metric distortion* framework, introduced by Anshelevich et al. [ABP15], wherein voters and candidates are thought of as points in some arbitrary metric space; this is akin to some models in spatial voting theory [CE03]. In this context, the voters' preferences are measured by means of their "proximity" from each candidate, and the goal is to output a candidate who (approximately) minimizes the *social cost*, i.e. the cumulative distances from the voters. A rather simplistic view of this framework manifests itself when agents and candidates are embedded into a one-dimensional line, and their locations indicate whether they are "left" or "right" on the political spectrum. However, the metric distortion framework has a far greater reach since no assumptions whatsoever are made for the underlying metric space. Indeed, the dimensionality of the space is potentially unbounded, while we are not even restricted in Euclidean spaces.

Importantly, this paradigm offers a compelling way to quantitatively compare different voting rules commonly employed in practice [SE17, Kem20a, GKM17, ABP15], while it also serves as a benchmark for designing new mechanisms in search of better distortion bounds [GHS20, MW19]. A common assumption made in this line of work is that the algorithm has access to the entire total rankings of the voters. However, there are many practical scenarios in which it might be desirable to truncate the ordinal information elicited by the mechanism. For example, requesting only the top preferences could further relieve the cognitive burden since it might be hard for a voter to compare alternatives which lie on the bottom of her list (for additional motivation for considering incomplete or partial orderings see [FKS21, CLWA13, BLP04], and references therein), while any truncation in the elicited information would also translate to more efficient communication. These reasons have driven several authors to study the decay of distortion under missing information [Kem20b, AP16, GAX17, FGMP19], potentially allowing some randomization (see our related work subsection). In this work we follow this line of research, offering several new insights and improved bounds over prior results.

1.1   Contributions & Techniques

First, we study voting rules which perform a sequence of pairwise comparisons between two candidates, with the result of each comparison being determined by the majority rule over the entire population of voters. This class includes many common mechanisms such as Copeland's rule [SM96] or the minimax scheme of Simpson and Kramer [LN95], and has received consider-

able attention in the literature of social choice; cf., see [LPR$^+$07], and references therein. Within the framework of (metric) distortion the following fundamental question arises:

> *How many pairwise comparisons between two candidates are needed to guarantee non-trivial bounds on the distortion?*

For example, Copeland's rule (and most of the common voting rules within this class) elicits all possible pairwise comparisons, i.e. $\binom{m}{2} = \Theta(m^2)$, and guarantees distortion at most 5 [ABP15]. Thus, it is natural to ask whether we can substantially truncate the number of elicited pairwise comparisons without sacrificing too much the efficiency of the mechanism. We stress that we allow the queries to be dynamically adapted during the execution of the algorithm. In this context, we provide the following strong positive result:

Theorem 1.1. *There exists a deterministic mechanism which elicits only $m-1$ pairwise comparisons and guarantees distortion $\mathcal{O}(\log m)$.*

Our mechanism is particularly simple and natural: In every round we arbitrarily pair the remaining candidates and we only extract the corresponding comparisons. Next, we eliminate all the candidates who lost and we continue recursively until a single candidate emerges victorious. Interestingly, this mechanism is widely employed in practical applications, for example in the knockout phases of most competitions, with the difference that typically some "prior" is used in order to construct the pairings. The main technical ingredient of the analysis is a powerful lemma developed by Kempe via an LP duality argument [Kem20a]. Specifically, Kempe characterized the social cost ratio between two candidates when there exists a sequence of intermediate alternatives such that every candidate in the chain pairwise-defeats the next one. We also supplement our analysis for this mechanism with a matching lower bound on a carefully constructed instance. Moreover, we show that any mechanism which performs (strictly) fewer than $m-1$ pairwise comparisons has *unbounded* distortion. This limitation applies even if we allow randomization either during the elicitation or the winner-determination phase. Indeed, there are instances for which only a single alternative can yield bounded distortion, but the mechanism simply does not have enough information to identify the "right" candidate.

Next, we study deterministic mechanisms which only receive an incomplete order of preferences from every voter, instead of the entire rankings as it is usually assumed. This setting has already received attention in the literature, most notably by Kempe [Kem20b], and has numerous applications in real-life voting systems. Arguably the most important such consideration arises when every voter provides her k-top preferences, for some parameter $k \in [m]$. Kempe showed [Kem20b] that there exists a deterministic mechanism which elicits only the k-top preferences and whose distortion is upper-bounded by $79m/k$; using a powerful tool developed in [Kem20a] this bound can be improved all the way down to $12m/k$. However, this still leaves a substantial gap with respect to the best-known lower bound, which is $2m/k$ if we ignore some additive constant factors. Thus, Kempe [Kem20b] left as an open question whether the aforementioned upper bound can be improved. In our work we make substantial progress towards bridging this gap, proving the following:

Theorem 1.2. *There exists a deterministic mechanism which only elicits the k-top preferences and yields distortion at most $6m/k + 1$.*

We should stress that the constant factors are of particular importance in the framework of metric distortion; indeed, closing the gap even for the special case of $k = m$ has received intense scrutiny in recent years [ABP15, MW19, Kem20a, GHS20]. From a technical standpoint the main technique for proving such upper bounds consists of identifying a candidate for which there exists a path to any other node such that every candidate in the path pairwise-defeats the next one by a sufficiently large margin (which depends on k). Importantly, the derived upper

bound crucially depends on the length of the path. Our main technical contribution is to show that there always exists a path of length 2 with the aforedescribed property, while the previous best result by Kempe established the claim only for paths of length 3.

Although our approach can potentially bring further improvements, closing the gap inevitably requires different techniques. In particular, a promising direction appears to stem from extending some of the claims established by Gkatzelis et al. [GHS20]. Indeed, we observe that a natural generalization of the main technical ingredient in [GHS20] would lower the upper bound to $4m/k - 1$ (Proposition 4.7), which appears to be optimal when k is close to m. More precisely, the authors in [GHS20] proved that a certain graph always has a perfect matching when the entire rankings are available; we conjecture that under k-top preferences there always exists a perfect matching for a subset of a k/m fraction of the voters (see Conjecture 4.6 for a more precise statement).

We also provide some other interesting bounds for deterministic mechanisms under missing information. Most notably, if the voting rule performs well on an arbitrary (potentially adversarially selected) subset of the voters can we quantify its distortion over the entire population? We answer this question with a sharp upper bound in Theorem 4.3. In fact, we use this result as a tool for some of our other proofs, but nonetheless we consider it to be of independent interest. It should be noted that even in the realm of partial or incomplete rankings there exists an *instance-optimal* mechanism via linear programming; this was first observed by Goel et al. [GKM17] when the total orders are available, but it directly extends in more general settings. Interestingly, we show that the recently introduced mechanism of Gkatzelis et al. [GHS20] which always obtains distortion at most 3 can be substantially outperformed by the LP mechanism. Namely, for some instances the mechanism of Gkatzelis et al. [GHS20] yields distortion almost 3 while the instance-optimal mechanism yields distortion close to 1.

Finally, we consider mechanisms which receive information from only a "small" *random sample* of voters; that is, we are concerned with the *sample complexity* required to ensure efficiency, which boils down to the following fundamental question:

> *How large should the size of the sample be in order to guarantee near-optimal distortion with high probability?*

More precisely, we are interested in deriving sample-complexity bounds which are *independent* on the number of voters n. This endeavor is particularly motivated given that in most applications $n \gg m$. Naturally, sampling approximations are particularly standard in the literature of social choice. Indeed, in many scenarios one wishes to predict the outcome of an election based on small sample (e.g. in polls or exit polls), while in many other applications it is considered even infeasible to elicit the entire input (e.g. in online surveys). In this context, we will be content with obtaining near-optimal distortion *with high probability* (e.g. 99%). This immediately deviates from the line of research studying randomized mechanisms (cf. see [AP16]) wherein it suffices to obtain a guarantee in expectation. We point out that it has been well-understood that a guarantee only in expectation might be insufficient in many cases; for example, Fain et al. [FGMP19] considered as the objective the *squared distortion* as a proxy in order to limit as much as possible the variance in the distortion. In fact, the authors in [FGMP19] are also concerned with sample complexity issues, but from a very different standpoint.

We stress that we only allow randomization during the preference elicitation phase; for a given random sample, which corresponds to the *entire* rankings of the voters, the mechanisms we consider act deterministically. Specifically, we analyze two main voting rules along this vein.

Theorem 1.3 (Approximate Copeland). *For any sufficiently small $\epsilon > 0$ there exists a mechanism which takes a sample of size $\Theta(\log(m)/\epsilon^2)$ voters and yields distortion at most $5 + \epsilon$ with probability 0.99.*

The techniques required for the proof of this theorem are fairly standard. More importantly, we analyze the sample complexity of PLURALITYMATCHING, the mechanism of Gkatzelis et al. [GHS20] which recovers the optimal distortion bound of 3 (among deterministic mechanisms). In this context, we establish the following result:

Theorem 1.4 (Approximate PLURALITYMATCHING). *For any sufficiently small $\epsilon > 0$ there exists a mechanism which takes a sample of size $\Theta(m/\epsilon^2)$ voters and yields distortion at most $3 + \epsilon$ with probability 0.99.*

More precisely, the main ingredient of PLURALITYMATCHING is a maximum-matching subroutine for a certain bipartite graph. Our first observation is that the size of the maximum matching can be determined through a much smaller graph which satisfies a "proportionality" condition with respect to a maximum-matching decomposition. Although this condition cannot be explicitly met since the algorithm is agnostic to the decomposition, our observation is that sampling (with sufficiently many samples) will approximately satisfy this requirement, eventually leading to the desired conclusion.

We stress that we do *not* guarantee that the winner in our sample will coincide with that over the entire population. In fact, the sample complexity bounds for the *winner determination* problem—for virtually every reasonable voting rule—depend on the *margin* of victory (see [DB15]); however, we argue that this feature is undesirable. For one thing the algorithm does not have any prior information on the margin, and hence it is unclear how to tune this parameter in practice. More importantly, in many scenarios the margin might be very small, leading to a substantial overhead in the sample-complexity requirements of the mechanism. One of our conceptual contributions is to show that we can circumvent such limitations once we espouse a utilitarian framework. Indeed, observe that all of our bounds are *distribution-independent* (and instance-oblivious).

We should also point out that although we are emphasizing on sample-complexity considerations, we believe that our results have another very clear motivation. Namely, given that in most applications $n \gg m$, it is important to provide *sublinear* algorithms whose running time does not depend on n. In this context, we provide a Monte Carlo implementation of PLURALITYMATCHING whose time complexity scales independently of n.

To conclude, we provide several experimental findings in real-life voting applications from the standpoint of the (metric) distortion framework. We are mostly concerned with comparing the results of the scoring systems used in practice against a mechanism which explicitly attempts to minimize the distortion; the latter is realized with the linear programming mechanism of Goel et al. [GKM17]. Specifically, we analyze the efficiency of the scoring rule used in the Eurovision song contest and the Formula One world championship. Interestingly, on both occasions we find that the winner in the actual competition is usually the candidate who minimizes the distortion. Our implementation is publicly available at `https://github.com/ioannisAnagno/Voting-MetricDistortion`.

1.2 Related Work

Research in the metric distortion framework was initiated by Anshelevich et al. [ABP15]. Specifically, they analyzed the distortion of several common voting rules, most notably establishing that Copeland's rule has distortion at most 5, with the bound being tight for certain instances. They also conjectured that the *ranked pairs* mechanism always achieves distortion at most 3, which is also the lower bound for any deterministic mechanism. This conjecture was disproved by Goel et al. [GKM17],[1] while they also studied *fairness* properties of certain voting rules. Moreover, Skowron and Elkind [SE17] established that a popular rule named *single transferable vote* (STV) has distortion $\mathcal{O}(\log m)$, along with a nearly-matching lower bound. The barrier

---

[1]A tight bound of $\Theta(\sqrt{m})$ for the ranked pairs mechanism was subsequently given by Kempe [Kem20a].

of 5 set out by Copeland was broken by Munagala and Wang [MW19], presenting a novel deterministic rule with distortion $2 + \sqrt{5}$. The same bound was obtained by Kempe [Kem20a] through an LP duality framework, who also articulated sufficient conditions for proving the existence of a deterministic mechanism with distortion 3. This conjecture was only recently confirmed by Gkatzelis et al. [GHS20], introducing the *plurality matching* mechanism. Closely related to our study is also the work of Gross et al. [GAX17], wherein the authors provide a near-optimal mechanism which only asks $m + 1$ voters for their top-ranked alternatives. The main difference with our setting is that we require an efficiency-guarantee with high probability, and not in expectation.

Broader Context. Beyond the metric case most focus has been on analyzing distortion under a *unit-sum* assumption on the utility function, ensuring that agents have equal "weights". In particular, Boutilier et al. [BCH+15] provide several upper and lower bounds, while they also study learning-theoretic aspects under the premise that every agent's utility is drawn from a distribution (cf., see [PZPR09]). Moreover, several multi-winner extensions have been studied in the literature. Caragiannis et al. [CNPS16] studied the *committee selection* problem, which consists of selecting k alternatives that maximize the social welfare, assuming that the value of each agent is defined as the maximum value derived from the committee's members. We also refer to [BNPS17] for the *participatory budgeting* problem, and to [BPQ19] when the output of the mechanism should be a total order over alternatives (instead of a single winner).

   More special cases were considered in [FFG16, FGMS17], strengthening some of the results we previously described. The trade-off between efficiency and communication has been addressed in [MPSW19, MSW20], while Amanatidis et al. [ABFV20] investigated the decay of distortion under a limited amount of cardinal queries—in addition to the ordinal information. We should also note a series of works analyzing the power of ordinal preferences for some fundamental graph-theoretic problems [FFZ14, AS16a, AS16b, AZ18]. Finally, we point out that strategic issues are typically ignored within this line of work. We will also posit that agents provide truthfully their preferences, but we refer to [BDG18, CFF+16, CFNV18] for rigorous considerations on the strategic issues that arise. We refer the interested reader to the survey of Anshelevich et al. [AFRSV21], as we have certainly not exhausted the literature.

2   Preliminaries

A *metric space* is a pair $(\mathcal{M}, d)$, where $d : \mathcal{M} \times \mathcal{M} \mapsto \mathbb{R}$ constitutes a *metric* on $\mathcal{M}$, i.e., (i) $\forall x, y \in \mathcal{M}, d(x, y) = 0 \iff x = y$ (identity of indiscernibles), (ii) $\forall x, y \in d(x, y) = d(y, x)$ (symmetry), and (iii) $\forall x, y, z \in \mathcal{M}, d(x, y) \leq d(x, z) + d(z, y)$ (triangle inequality). Consider a set of n voters $V = \{1, 2, \ldots, n\}$ and a set of m candidates $C = \{a, b, \ldots, \}$; candidates will be typically represented with lowercase letters such as a, b, w, x, but it will be sometimes convenient to use numerical values as well. We assume that every voter $i \in V$ is associated with a point $v_i \in \mathcal{M}$, and every candidate $a \in C$ to a point $c_a \in \mathcal{M}$. Our goal is to select a candidate x who minimizes the *social cost*: $\mathrm{SC}(x) = \sum_{i=1}^{n} d(v_i, c_x)$. This task would be trivial if we had access to the agents' distances from all the candidates. However, in the standard *metric distortion* framework every agent i provides only a *ranking* (a total order) σ_i over the points in C according to the *order* of i's distances from the candidates. We assume that ties are broken arbitrarily, subject to transitivity, but we will not abuse the tie-breaking assumption.

   In this work we are considering a substantially more general setting, wherein every agent provides a subset of σ_i. More precisely, we assume that agent i provides as input a set $\mathcal{P}_i$ of ordered pairs of distinct candidates, such that $(a, b) \in \mathcal{P}_i \implies a \succ_i b$, where $a, b \in C$; it will always be assumed that $\mathcal{P}_i$ corresponds to the *transitive closure* of the input. We will allow $\mathcal{P}_i$ to be the empty set, in which case i does not provide any information to the mechanism; with a slight abuse of notation we will let $\mathcal{P}_i \equiv \sigma_i$ when i provides the entire order of preferences.

We will say that the input $\mathcal{P} = (\mathcal{P}_1, \dots, \mathcal{P}_n)$ is *consistent* with the metric d if $(a, b) \in \mathcal{P}_i \implies d(v_i, c_a) \leq d(v_i, c_b), \forall i \in V$, and this will be denoted with $d \triangleright \mathcal{P}$. We will represent with $\mathrm{top}(i)$ and $\mathrm{sec}(i)$ i's first and second most preferred candidates respectively. We may also sometimes use the notation $ab = \{i \in V : a \succ_i b\}$.

A *deterministic social choice rule* is a function that maps an *election* in the form of a 3-tuple $\mathcal{E} = (V, C, \mathcal{P})$ to a single candidate $a \in C$. We will measure the performance of f for a given input of preferences $\mathcal{P}$ in terms of its *distortion*, namely, the worst-case approximation ratio it provides with respect to the social cost:

$$\mathrm{distortion}(f; \mathcal{P}) = \sup \frac{\mathrm{SC}(f(\mathcal{P}))}{\min_{a \in C} \mathrm{SC}(a)}, \tag{1}$$

where the supremum is taken over all metrics such that $d \triangleright \mathcal{P}$. The distortion of a social choice rule f is the maximum of $\mathrm{distortion}(f; \mathcal{P})$ over all possible input preferences $\mathcal{P}$. In other words, once the mechanism selects a candidate (or a distribution over candidates if the social choice rule is *randomized*) an adversary can select any metric space subject to being consistent with the input preferences. Similarly, in Section 3 where we study mechanisms which perform pairwise comparisons, the adversary can select any metric space consistent with the elicited comparisons. We should point out the following:

Proposition 2.1. *Under any given preferences $\mathcal{P}$, there exists a metric space consistent with $\mathcal{P}$.*

This proposition follows directly from Proposition 1 in [ABP15], which established the claim when $\mathcal{P} = \sigma$.

2.1 Instance-Optimal Voting

An important observation is that under any input preferences $\mathcal{P}$ there exists a deterministic *instance-optimal* mechanism; this was noted by Goel et al. [GKM17] (see also [BCH+15]) when $\mathcal{P} = \sigma$, but their mechanism directly applies for our more general setting. We briefly present their idea, as we will also employ this mechanism for our experiments.

The first ingredient is an optimization problem that allows to compare a pair of distinct candidates, subject to the set of preferences given to the mechanism. Specifically, for $a, b \in C$, with $a \neq b$, consider the following *linear program* METRIC-LP(a, b):

$$
\begin{aligned}
\text{maximize} \quad & \sum_{i=1}^{n} x_{i,a} \\
\text{subject to} \quad & \sum_{i=1}^{n} x_{i,b} = 1; \\
& x_{i,p} \leq x_{i,q}, && \forall (p, q) \in \mathcal{P}_i, \forall i \in V; \\
& x_{i,i} = 0, && \forall i \in V \cup C; \\
& x_{i,j} = x_{j,i}, && \forall i, j \in V \cup C; \\
& x_{i,j} \leq x_{i,k} + x_{k,j}, \forall i, j, k \in V \cup C.
\end{aligned}
\tag{2}
$$

First of all, observe that the output of this linear program does not necessarily yield a metric space, but rather a *pseudo-metric* on $V \cup C$, given that it is possible that $x_{i,j} = 0$ for $i \neq j$. Nonetheless, this issue can be easily resolved by "merging" all the elements $i, j \in V \cup C$ such that $x_{i,j} = 0$; the consistency of this approach follows by the "triangle inequalities"—the final constraint of the linear program. It should be pointed out that some of the constraints in METRIC-LP are redundant, in the sense that they are implied by others, but we will not dwell on such optimizations here.

We will represent with $\mathfrak{D}(a|b)$ the value of the linear program METRIC-LP(a, b); if it is unbounded we let $\mathfrak{D}(a|b) = +\infty$. Note that the linear program is always feasible by virtue of Proposition 2.1. In this context, the mechanism of Goel et al. [GKM17] consists of the following steps:

- For any pair $a, b \in C$, with $a \neq b$, compute $\mathfrak{D}(a|b)$; also let $\mathfrak{D}(a|a) = 1$.

- Set $\mathfrak{D}(a) = \max_{b \in C} \mathfrak{D}(a|b)$.

- Return the candidate b with the minimum value $\mathfrak{D}(a)$ over all $a \in C$; ties are broken arbitrarily.

This mechanism will be referred to as MINIMAX-LP, to distinguish from the *minimax* voting scheme of Simpson and Kramer [LN95]. The MINIMAX-LP rule essentially performs brute-force search over all possible metrics in order to identify the candidate which minimizes the distortion; nonetheless, it can be solved in $\text{poly}(n, m)$ time given that the METRIC-LP admits a strongly polynomial time algorithm; this follows because the *bit complexity L*—the number of bits required to represent it [Kar84]—is small: $L = \mathcal{O}(\log(n + m))$. Moreover, it is easy to establish the following:

Theorem 2.2. *For any given preferences* $\mathcal{P}$, *the* MINIMAX-LP *rule is instance-optimal in terms of distortion.*

In particular, when $\mathcal{P} = \sigma$ note that MINIMAX-LP always yields distortion 3 by virtue of an upper-bound by Gkatzelis et al. [GHS20]. However, in this work we will be mostly interested in providing upper bounds on the distortion of MINIMAX-LP under incomplete rankings.

3 Sequence of Pairwise Comparisons

In this section we are considering voting rules which perform a sequence of pairwise comparisons between two candidates, with the result of each comparison being determined by the majority rule over the entire population of voters. To put it differently, consider the tournament graph $T = (C, E)$ where $(a, b) \in E$ if and only if candidate a pairwise-defeats candidate b; it will be tacitly assumed—without any loss of generality—that ties are broken arbitrarily so that T is indeed a tournament. We are studying mechanisms which elicit edges from T, and we are interested in establishing a trade-off between the number of elicited edges and the distortion of the mechanism. We commence with the following lower bound:

Proposition 3.1. *There are instances for which any deterministic mechanism which elicits (strictly) fewer than* $m - 1$ *edges from* T *has unbounded distortion.*

Sketch of Proof. Consider a family of tournaments $\mathcal{T}$ as illustrated in Figure 1, with the set C^* containing a single candidate. Then, there are metric spaces for which all the voters are arbitrarily close to the candidate in C^* and arbitrarily far from any other candidate. Thus, every mechanism with bounded distortion has to identify the candidate in C^*. However, it is easy to see that any pairwise comparison can eliminate at most one candidate from being in C^*. As a result, if $\widehat{T} = (C, \widehat{E})$ is the subgraph based on the elicited edges, there will be at least two distinct candidates which could lie in C^* for some tournament in $\mathcal{T}$ consistent with $\widehat{T}$, leading to the desired conclusion. □

In fact, the same limitation applies even if we allow randomization, either during the elicitation or the winner-determination phase. Importantly, we will show that $m - 1$ edges from T suffice to obtain near-optimal distortion. To this end, we will employ a powerful technical lemma by Kempe, proved via an LP-duality argument.
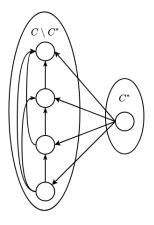
Figure 1: A hard instance when fewer than $m - 1$ pairwise comparisons are elicited.

Lemma 3.2 ([Kem20a])**.** *Let $a_1, a_2, \ldots a_\ell$ be a sequence of distinct candidates such that for every $i = 2, \ldots, \ell$ at least half of the agents prefer candidate a_{i-1} over candidate a_i. Then, $\mathrm{SC}(a_1) \leq (2\ell - 1)\,\mathrm{SC}(a_\ell)$.*

Armed with this important lemma we introduce the DOMINATIONROOT mechanism, which determines a winning candidate with access only to a pairwise comparison oracle; namely, $\mathfrak{D}$ takes as input two distinct candidates $a, b \in C$ and returns the *losing* candidate based on the voters' preferences (recall that in case of a tie the oracle returns an arbitrary candidate).

DOMINATIONROOT Mechanism
Input: Set of candidates C, Pairwise comparison oracle $\mathfrak{D}$
Output: Winner $w \in C$

  1. Initialize $S := C$

  2. Construct arbitrarily a set Π of $\lfloor S/2 \rfloor$ pairings from S

  3. For every $\{a, b\} \in \Pi$ remove $\mathfrak{D}(a, b)$ from S

  4. If $|S| = 1$ **return** $w \in S$; otherwise, continue from step 2

We refer to Figure 2 for an illustration of DOMINATIONROOT. The analysis of this mechanism boils down to the following simple claims:

Claim 3.3. DOMINATIONROOT *elicits exactly $m - 1$ edges from T.*

Proof. The claim follows given that for every elicited edge we remove a candidate for the rest of the mechanism, until only a single candidate survives. □

Claim 3.4. DOMINATIONROOT *returns a candidate w which can reach every other node in T in paths of length at most $\lceil \log m \rceil$.*

Proof. Consider the partition of candidates $C_1, \ldots, C_r$ such that C_i contains the candidates who were eliminated during the i-th round for $i \in \{1, 2, \ldots, r-1\}$, and $C_r = \{w\}$. Observe that every candidate $a \in C_i$ (with $i \in \{1, 2, \ldots, r-1\}$) was pairwise-defeated by some candidate in C_j for $j > i$; thus, the claim follows inductively since $r = \lceil \log m \rceil$. □

Theorem 3.5. DOMINATIONROOT *elicits only $m - 1$ edges from T and guarantees distortion at most $2\lceil \log m \rceil + 1$.*
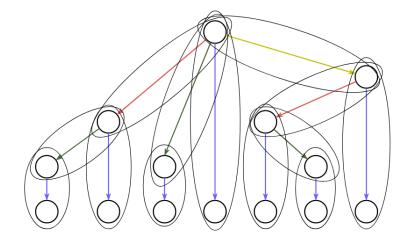
Figure 2: An implementation of DOMINATIONROOT for $m = 14$ candidates. We have high-lighted with different colors pairwise comparisons that correspond to different rounds of the mechanism; also notice that the "height" of every candidate indicates the order of elimination.
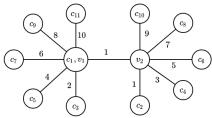
Proof. The theorem follows directly from Claim 3.3, Claim 3.4, and Lemma 3.2. □
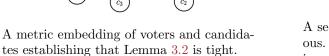
This theorem along with Proposition 3.1 imply a remarkable gap depending on whether the mechanism is able to elicit at least $m - 1$ pairwise comparisons. We also provide a matching lower bound for DOMINATIONROOT.
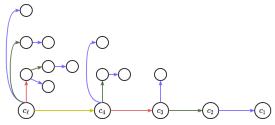
Proposition 3.6. *There exist instances for which* DOMINATIONROOT *yields distortion at least* $2 \log m + 1$.

Proof. We will first show that the bound established in Lemma 3.2 is tight. Indeed, consider a set of ℓ candidates $\{1, 2, \ldots, \ell\}$ and two voters (the instance directly extends to an arbitrary even number of voters) positioned according to the pattern of Figure 3a. Observe that—at least under some tie-breaking—candidate i pairwise-defeats candidate $i - 1$ for $i = 2, 3, \ldots, \ell$. Moreover, it follows that $\mathrm{SC}(i) = 2i - 1$, for all i, implying that $\mathrm{SC}(\ell) / \mathrm{SC}(1) = 2\ell - 1$, as desired.

Now consider m candidates such that m is a power of 2. We first consider $\ell = \log m + 1$ candidates positioned according to our previous argument (Figure 3a); the rest of the candidates are located arbitrarily far from the voters. It is easy to see that there exists a sequence of pairings (Figure 3b) such that c_ℓ will be declared victorious, leading to a distortion of $2 \log m + 1$ by virtue of our previous argument. □

A metric embedding of voters and candidates establishing that Lemma 3.2 is tight.

A sequence of pairings such that c_ℓ emerges victorious. We have highlighted with different colors pairings that correspond to different rounds.

4 Distortion of Deterministic Rules under Incomplete Orders

We commence this section with another useful lemma by Kempe [Kem20a].

Lemma 4.1 ([Kem20a])**.** *Consider three distinct candidates $w, y, x \in C$ such that $|wy| \geq \alpha n > 0$ and $|yx| \geq \alpha n > 0$. Then,*

$$\frac{\text{SC}(w)}{\text{SC}(x)} \leq \frac{2}{\alpha} + 1. \tag{3}$$

In particular, notice that if w is the winner in Copeland's rule it follows that for any candidate x there exists some other candidate y such that w pairwise-defeats y and y pairwise-defeats x [Mou86]; thus, applying Lemma 4.1 for $\alpha = 1/2$ implies that the winner in Copeland's rule has distortion upper-bounded by 5, as initially articulated in [ABP15].

As a warm-up, we first employ this lemma to characterize the distortion when for all pairs of candidates at least a small fraction of voters has provided their pairwise preferences. We stress that all of our upper bounds are attainable by the MINIMAX-LP rule, but nonetheless our proofs are constructive in the sense that we provide (efficient) mechanisms which obtain the desired bounds.

Proposition 4.2. *Consider an election $\mathcal{E} = (V, C, \mathcal{P})$ such that for every pair of distinct candidates $a, b \in C$, it holds that $\sum_{i=1}^{n} \mathbb{1}\{(a, b) \in \mathcal{P}_i \vee (b, a) \in \mathcal{P}_i\} \geq \alpha \cdot n$. Then, there exists a voting rule which obtains distortion at most $4/\alpha + 1$.*

Proof. Consider a complete, weighted and directed graph $G = (C, E, w)$ such that

$$w_{a,b} = \frac{\sum_{i=1}^{n} \mathbb{1}\{(a, b) \in \mathcal{P}_i\}}{n}. \tag{4}$$

In words, $w_{a,b}$ represents the fraction of voters who *certainly* prefer a to b; observe that if we had the complete rankings it would follow that $w_{a,b} + w_{b,a} = 1$, but here we can only say that $w_{a,b} + w_{b,a} \leq 1$. Moreover, by assumption we know that $w_{a,b} + w_{b,a} \geq \alpha$, implying that $w_{a,b} \geq \alpha/2$ or $w_{b,a} \geq \alpha/2$. With that in mind, we construct from G an unweighted and directed graph $\widehat{G} = (C, \widehat{E})$ according to the following threshold rule: $(a, b) \in \widehat{E} \iff w_{a,b} \geq \alpha/2$. We argued that our assumption implies that $(a, b) \in \widehat{E} \vee (b, a) \in \widehat{E}$. As a result, we can deduce that $\widehat{G}$ contains as a subgraph a tournament; thus, there exists a *king* vertex w so that every node $a \in C$ is reachable from w in at most 2 steps, and our claim follows from Lemma 4.1. □

We remark that this upper bound is tight up to constant factors, at least for certain instances. Indeed, if we only have an α fraction of the votes in the presence of 2 candidates it is easy to show an $\Omega(1/\alpha)$ lower bound for any mechanism, even if we allow randomization. Interestingly, Proposition 4.2 suggests one possible preference elicitation strategy: collect the information about the preferences in a "balanced" manner.

4.1 Missing Voters

Consider an election $\mathcal{E} = (V, C, \mathcal{P})$ and a mechanism which has access to the votes of only a subset $V \setminus Q$ of voters, where $Q \subset V$ is the set of *missing voters* such that $|Q| = \epsilon \cdot n$. If the mechanism performs well on $V \setminus Q$ can we characterize the distortion over the entire set of voters as ϵ increases? Observe that this setting is tantamount to $\mathcal{P}_i = \emptyset$ for all $i \in Q$. In the following theorem we provide a sharp bound:

Theorem 4.3. *Consider a mechanism with distortion at most ℓ w.r.t. an arbitrary subset with $(1 - \epsilon)$ fraction of all the voters, for some $\epsilon \in (0, 1)$. Then, the distortion of the mechanism w.r.t. the entire population is upper-bounded by*

$$\ell + \frac{\epsilon}{1 - \epsilon}(\ell + 1). \tag{5}$$

Proof. Consider a candidate $b \in C$ with distortion at most ℓ w.r.t. the agents in $V \setminus Q$. Moreover, consider some arbitrary candidate $a \in C$, and let $S_b = \sum_{i \in V \setminus Q} d(v_i, c_b)$, and $S_a = \sum_{i \in V \setminus Q} d(v_i, c_a)$; observe that (by assumption) $S_b / S_a \leq \ell$. Our analysis will distinguish between the following two cases:

Case I: $S_b \geq S_a > 0.$[2] Then, for all $i \in Q$ it follows that

$$S_b d(v_i, c_a) + S_a d(c_a, c_b) \geq S_a(d(v_i, c_a) + d(c_a, c_b)) \geq S_a d(v_i, c_b), \tag{6}$$

and hence,

$$S_a d(v_i, c_b) \leq S_a d(c_a, c_b) + S_b d(v_i, c_a) + d(c_a, c_b) d(v_i, c_a); \tag{7}$$

summing over all $i \in Q$ gives

$$S_a \sum_{i \in Q} d(v_i, c_b) \leq |Q| S_a d(c_a, c_b) + S_b \sum_{i \in Q} d(v_i, c_a) + d(c_a, c_b) \sum_{i \in Q} d(v_i, c_a)$$

$$\leq |Q| S_a d(c_a, c_b) + S_b \sum_{i \in Q} d(v_i, c_a) + |Q| d(c_a, c_b) \sum_{i \in Q} d(v_i, c_a). \tag{8}$$

Moreover, observe that

$$(8) \iff \frac{S_b + \sum_{i \in Q} d(v_i, c_b)}{S_a + \sum_{i \in Q} d(v_i, c_a)} \leq \frac{S_b + |Q| d(c_a, c_b)}{S_a}. \tag{9}$$

Next, we have that $d(c_a, c_b) \leq d(v_i, c_a) + d(v_i, c_b), \forall i$; summing over all $i \in V \setminus Q$ implies that $(n - |Q|) d(c_a, c_b) \leq S_a + S_b \leq (\ell + 1) S_a$. Therefore, along with (9) we obtain that

$$\frac{\mathrm{SC}(b)}{\mathrm{SC}(a)} \leq \ell + \frac{|Q|}{n - |Q|} (\ell + 1) = \ell + \frac{\epsilon}{1 - \epsilon} (\ell + 1). \tag{10}$$

Case II: $S_b < S_a$. In this case we can simply observe that

$$\frac{\mathrm{SC}(b)}{\mathrm{SC}(a)} \leq \frac{S_b + \sum_{i \in Q} d(v_i, c_a) + |Q| d(c_a, c_b)}{S_a + \sum_{i \in Q} d(v_i, c_a)} \leq 1 + |Q| \frac{d(c_a, c_b)}{S_a}. \tag{11}$$

Thus, the proof follows given that $(n - |Q|) d(c_a, c_b) \leq S_a + S_b < 2 S_a$. $\qquad\square$

A few remarks are in order. First of all, if all the voters in the set $V \setminus Q$ had provided their entire rankings we would derive a similar result via Proposition 4.2, but nonetheless Theorem 4.3 gives a more precise characterization. Indeed, we claim that the derived bound is tight. For example, consider an instance on the line with only two candidates a, b, so that every candidate receives half of the votes among the voters in $V \setminus Q$; assume without loss of generality that a is selected as the winning candidate, having distortion 3 w.r.t. the voters in $V \setminus Q$. However, we have to accept that $(1 - \epsilon)/2$ fraction of the voters could reside in the midpoint $(c_a + c_b)/2$, while the rest of the agents could all lie in c_b; thus, the distortion of candidate a is

$$\frac{\mathrm{SC}(a)}{\mathrm{SC}(b)} = \frac{\frac{1-\epsilon}{2} \frac{d(c_a, c_b)}{2} + \frac{1-\epsilon}{2} d(c_a, c_b) + \epsilon d(c_a, c_b)}{\frac{1-\epsilon}{2} \frac{d(c_a, c_b)}{2}} = 3 + 4 \frac{\epsilon}{1 - \epsilon}, \tag{12}$$

which matches our derived bound in Theorem 4.3.

---

[2]The case where $S_a = 0$ can be trivially handled. Indeed, it implies that $S_b \leq \ell \times S_a = 0$, which in turn yields that $d(c_a, c_b) = 0$; thus, $\mathrm{SC}(a) = \mathrm{SC}(b)$.

4.2 Top Preferences

In this subsection we investigate how the distortion increases when every voter provides only her k-top preferences, for some parameter $k \in [m]$. It should be noted that the two extreme cases are well understood. Specifically, when $k = m$ the mechanism has access to the entire rankings and we know that any deterministic mechanism has distortion at least 3, which is also the upper bound established in [GHS20]. On the other end of the spectrum, when $k = 1$ the plurality rule—which incidentally is the optimal deterministic mechanism when only the top preference is given—yields distortion at most $2m - 1$ [ABP15]. Consequently, the question is quantify the decay of distortion as we gradually increase k. We commence by reminding a lower bound given by Kempe [Kem20b]:

Proposition 4.4. *Any deterministic mechanism which elicits only the k-top preferences from every voter out of the m alternatives has distortion $\Omega(m/k)$.*

More precisely, the best lower bound is $2m/k$, ignoring some additive constant factors; for completeness we provide a proof in Appendix A. In the following theorem we come closer to matching this lower bound.

Theorem 4.5. *There exists a deterministic mechanism which elicits only the k-top preferences from every voter out of m candidates and has distortion at most $6m/k + 1$.*

Before we proceed with the proof it is important to point out that having only the k-top preferences is not subsumed by our previous consideration in Proposition 4.2; e.g., even if $k = m - 2$ there could be two candidates which lie on the last two positions of every voter's list, and hence, it is impossible to know which one is mostly preferred among the voters.

Proof of Theorem 4.5. Let $\mathcal{L}_i$ be the set with the k-top preferences of voter i. For a candidate $a \in C$ we let

$$\mathcal{V}_a = \frac{\sum_{i=1}^n \mathbb{1}\{a \in \mathcal{L}_i\}}{n}; \tag{13}$$

i.e. the fraction of voters for which a is among the k-top preferences. Notice that $\sum_{a \in C} \mathcal{V}_a = k$, and hence, by the pigeonhole principle there exists some candidate x such that $\mathcal{V}_x \geq k/m$. Similarly to Proposition 4.2 we consider the weighted, complete and directed graph $G = (C, E, w)$, so that

$$w_{a,b} = \frac{\sum_{i=1}^n \mathbb{1}\{(a, b) \in \mathcal{P}_i\}}{n}. \tag{14}$$

Moreover, based on G we construct the unweighted and directed graph $\widehat{G} = (C, \widehat{E})$, so that $(a, b) \in \widehat{E} \iff w_{a,b} \geq k/(3m)$; the constant $1/3$ in the threshold was selected as the largest number which makes the following argument work. In particular, we will show that $\widehat{G}$ has a king vertex, and then the theorem will follow by virtue of Lemma 4.1.

Let $C' = \{a \in C : \exists b \in C \setminus \{a\}.(a, b) \notin \widehat{E} \wedge (b, a) \notin \widehat{E}\}$ and $C^* = C \setminus C'$. Observe that the *induced* subgraph on C^* contains as a subgraph a tournament, and as such, it has a king vertex $w \in C^*$ (we will argue very shortly that indeed $C^* \neq \emptyset$). As a result, if $C' = \emptyset$ the theorem follows.

In the contrary case notice that C' contains at least two (distinct) nodes; let $a, b \in C'$ be such that $(a, b) \notin \widehat{E} \wedge (b, a) \notin \widehat{E}$. An important observation is that $\mathcal{V}_a \leq 2k/(3m)$ and $\mathcal{V}_b \leq 2k/(3m)$. Indeed, for the sake of contradiction let us assume that $\mathcal{V}_a > 2k/(3m)$. Given that $(a, b) \notin \widehat{E}$ we can infer that b is preferred over a in at least a $k/(3m)$ fraction of the voters; however, this would imply that $(b, a) \in \widehat{E}$, which is a contradiction. Similarly, we can show that $\mathcal{V}_b \leq 2k/(3m)$. Consequently, x cannot belong in the set C', where recall that x is a candidate for which $\mathcal{V}_x \geq k/m$, verifying that $C^* \neq \emptyset$.

Next, it is easy to see that for all $a \in C', (x, a) \in \widehat{E}$; this follows since $\mathcal{V}_a \leq 2k/(3m)$ for all $a \in C'$ while $\mathcal{V}_x \geq k/m$. As a result, if $x = w$ or if there exists the edge $(w, x) \in \widehat{E}$ then w can reach every node in at most 2 steps, and the theorem follows. Otherwise, it follows that there exists a path of length 2 from w to x since w is a king vertex in the induced subgraph on C^* and $x \in C^*$. We shall distinguish between two cases.

First, assume that for all $z \in C_1^*, (x, z) \in \widehat{E}$, where C_1^* is the subset of C^* which is reachable from w via a single edge. Then, given that we have assumed that $(w, x) \notin \widehat{E}$ and the induced graph on C^* is a tournament, it follows that $(x, w) \in \widehat{E}$ and subsequently x can reach every node in C in paths of length at most 2, as desired.

Finally, assume that there exists some $y \in C_1^*$ such that $(x, y) \notin \widehat{E}$. This implies that y is preferred over x in at least a $2k/(3m)$ fraction of the voters. If for every candidate $a \in C'$ it holds that $(w, a) \in \widehat{E}$ or $(z, a) \in \widehat{E}$ for some $z \in C_1^*$, we can conclude that w can reach every node in $\widehat{G}$ in at most 2 steps, again reaching the desired conclusion. On the other hand, assume that there exists $b \in C'$ such that $(w, b) \notin \widehat{E}$ and $(z, b) \notin \widehat{E}$ for all $z \in C_1^*$. By the definition of the set C^* we can infer that $(b, w) \in \widehat{E}$ and $(b, z) \in \widehat{E}$ for all $z \in C_1^*$. Moreover, we know that from all of the votes candidate y received candidate b was below in at most a $k/(3m)$ fraction (over all the voters); otherwise it would follow that $(y, b) \in \widehat{E}$. As a result, since y is preferred over x in at least a $2k/(3m)$ fraction of the voters we can conclude that $(b, x) \in \widehat{E}$, in turn implying that b can reach every node in $\widehat{G}$ in paths of length at most 2, concluding the proof. □
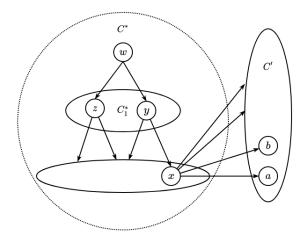

Figure 4: The anatomy of our proof for Theorem 4.5. The set of candidates is partitioned into a "good" set C^* and a "bad" set C'; C^* has a king vertex w, and we can essentially apply the reasoning of Proposition 4.2. A key observation is that C' is always dominated by some node in C^*, namely x.

Notably, we have shown that if $k = \gamma \cdot m$ for some $\gamma \in (0, 1)$, the distortion is at most $6/\gamma + 1$. Our analysis substantially improves over the previous best-known bound which was $12m/k$ [Kem20b, Kem20a], but nonetheless there is still a gap between the aforementioned lower bound. Before we conclude this section, we explain how one can further improve upon the bound obtained in Theorem 4.5.

Conjecture 4.6. *If we assume that every agent provides her k-top preferences for some $k \in [m]$, there is a candidate $a \in C$ and a subset $S \subseteq V$ such that*

- *There exists a perfect matching $M : S \mapsto S$ in the integral domination graph of a (see Definition 5.4 in the next section);*

- $|S| \geq n \times k/m$.

When $k = m$ this conjecture was shown to be true by Gkatzelis et al. [GHS20]. On the other end of the spectrum, when $k = 1$ it is easy to verify that the plurality winner establishes this conjecture.

Proposition 4.7. *If Conjecture 4.6 holds, then there exists a deterministic mechanism which elicits only the k-top preferences and yields distortion at most $4m/k - 1$.*

Proof. Let $a \in C$ be the candidate which satisfies Conjecture 4.6. Then, it follows that a yields distortion at most 3 w.r.t. the voters in the set S [GHS20]. As a result, Theorem 4.3 implies that the distortion of a is upper-bounded by

$$3 + 4\frac{n - |S|}{|S|} \leq \frac{4m}{k} - 1. \tag{15}$$

$\square$

5 Randomized Preference Elicitation & Sampling

Previously we characterized the distortion when only a deterministically (and potentially adversarially) selected subset of voters has provided information to the mechanism. This raises the question of bounding the distortion when the mechanism elicits information from only a small *random sample* of voters. Here a single sample corresponds to the *entire* ranking of a voter. We stress that randomization is only allowed during the preference elicitation process; for any given random sample as input the mechanism has to select a candidate *deterministically*. We commence this section with a simple lower bound, which essentially follows from a standard result by Canetti et al. [CEG95].

Proposition 5.1. *Any mechanism which yields distortion at most $3 + \epsilon$ with probability at least $1 - \delta$ requires $\Omega(\log(1/\delta)/\epsilon^2)$ samples, even for $m = 2$.*

Proof. Consider two candidates a, b, and assume that exactly $(1 - \epsilon)/2$ fraction of the voters prefer candidate a. It is easy to verify that a yields distortion strictly larger than $3 + \epsilon$; thus, any mechanism with distortion at most $3 + \epsilon$ has to return candidate b. However, we know from [CEG95] that the *winner determination* problem with *margin* ϵ requires $\Omega(\log(1/\delta)/\epsilon^2)$ samples, concluding the proof. $\square$

5.1 Approximating Copeland

Our main result in this subsection is the following:

Theorem 5.2. *For any $\epsilon \in (0, 4]$ and $\delta \in (0, 1)$ there exists a mechanism which takes a sample of size $c = \Theta(\log(m/\delta)/\epsilon^2)$ voters and yields at most $5 + \epsilon$ distortion with probability at least $1 - \delta$.*

In particular, the proof essentially analyzes a *sampling approximation* of Copeland's rule, which recall that yields at most 5 distortion when the entire input is available [ABP15]. As a result, it follows that $\widetilde{\Theta}(m/\epsilon^2)$ bits of information (in total) suffice to yield $5 + \epsilon$ distortion with high probability, where the notation $\widetilde{\Theta}(\cdot)$ suppresses poly-logarithmic factors. Before we proceed with the proof we state the following standard fact:

Lemma 5.3 (Chernoff-Hoeffding Bound). *Let $\{X_1, X_2, \ldots, X_c\}$ be a set of i.i.d. random variables with $X_i \sim Bern(p)$ and $X_\mu = (X_1 + X_2 + \cdots + X_c)/c$; then,*

$$\mathbb{P}(|X_\mu - p| \geq \epsilon) \leq 2e^{-2\epsilon^2 c}. \tag{16}$$

14

Proof of Theorem 5.2. Consider the complete, weighted and directed graph $G = (C, E, w)$ so that $w_{a,b} = |ab|/n$. We will show how to use the random sample in order to construct a graph $\widehat{G} = (C, E, \widehat{w})$ which approximately preserves the weights of G with high probability. In particular, consider some parameters $\epsilon \in (0, 1/2)$ and $\delta \in (0, 1)$, and take a sample S of size $|S| = c = \Theta(\log(m/\delta)/\epsilon^2)$ from the set of voters V; for simplicity we assume that the sampling occurs *with replacement* in order to guarantee independence, but the result holds even without replacement given that the dependencies are negligible (e.g., see [KR08]). Now we let $\widehat{w}_{a,b} = |\{i \in S : a \succ_i b\}|/c$. Lemma 5.3 implies that $|\widehat{w}_{a,b} - w_{a,b}| < \epsilon$ with probability at least $1 - \delta/m^2$. Thus, the union bound implies that for all distinct pairs a, b we have approximately preserved the weights: $|\widehat{w}_{a,b} - w_{a,b}| < \epsilon$ with probability at least $1 - \delta$.

From $\widehat{G}$ we construct the directed graph $T = (C, \widehat{E})$ so that $(a, b) \in \widehat{E} \iff \widehat{w}_{a,b} \geq 1/2$; if $\widehat{w}_{a,b} = \widehat{w}_{b,a}$ for some distinct candidates $a, b \in C$ we only retain one of the edges arbitrarily (this conundrum can be avoided by taking c to be odd). In this way T will be a tournament, and as such, there exists a candidate w which can reach every node in T in at most 2 steps. Thus, for any $a \in C$ there exists some intermediate candidate $b \in C$ so that $|wb| \geq 1/2 - \epsilon$ and $|ba| \geq 1/2 - \epsilon$. As a result, Lemma 4.1 implies that the distortion of w is upper-bounded by $4/(1 - 2\epsilon) + 1 \leq 5 + 16\epsilon$, for any $\epsilon \leq 1/4$. Finally, rescaling ϵ by a constant factor concludes the proof. $\qquad\square$

5.2 Approximating PluralityMatching

In light of Proposition 5.1 the main question that arises is whether we can asymptotically reach the optimal distortion bound of 3. To this end, we will analyze a sampling approximation of PluralityMatching, a deterministic mechanism introduced by Gkatzelis et al. [GHS20] which obtains the optimal distortion bound of 3. To keep the exposition reasonably self-contained we recall some basic facts about PluralityMatching.

Definition 5.4 ([GHS20], Definition 5). *For an election $\mathcal{E} = (V, C, \sigma)$ and a candidate $a \in C$, the integral domination graph of candidate a is the bipartite graph $G(a) = (V, V, E_a)$, where $(i, j) \in E_a$ if and only if $a \succeq_i \text{top}(j)$.*

Proposition 5.5 ([GHS20], Corollary 1). *There exists a candidate $a \in C$ whose integral domination graph $G(a)$ admits a perfect matching.*

Before we proceed let us first introduce some notation. For this subsection it will be convenient to use numerical values in the set $\{1, 2, \ldots, m\}$ to represent the candidates. We let $\Pi_j = \sum_{i \in V} \mathbb{1}\{\text{top}(i) = j\}$, i.e. the number of voters for which $j \in C$ is the top candidate. For candidate $j \in C$ we let $G(j)$ be the integral domination graph of j, and M_j be a maximum matching in $G(j)$. In the sequel, it will be useful to "decompose" M_j as follows. We consider the partition of V into $V_j^0, V_j^1, \ldots, V_j^m$ such that $V_j^k = \{i \in V : M_j(i) = k\}$ for all $k \in [m]$, while V_j^0 represents the subset of voters which remained unmatched under M_j; see Figure 5.

Moreover, consider a set $\mathcal{S} = S_j^0 \cup S_j^1 \cup \cdots \cup S_j^m$ such that $S_j^k \subseteq V_j^k$ for all k; we also let $c = |\mathcal{S}|$, and $\Pi_j' = c/n \times \Pi_j$. For now let us assume that $\Pi_j' \in \mathbb{N}$ for all j. We let $G^{\mathcal{S}}(j)$ represent the induced subgraph of $G(j)$ w.r.t. the subset $\mathcal{S} \subseteq V$ and the new plurality scores Π_j'. We start our analysis with the following observation:

Observation 5.6. *Assume that $\mathcal{S}$ is such that $|S_j^k|/c = |V_j^k|/n$ for all k. Then, if $M_j^{\mathcal{S}}$ represents the maximum matching in $G^{\mathcal{S}}(j)$, it follows that $|M_j^{\mathcal{S}}|/c = |M_j|/n$.*

Sketch of Proof. First, it is clear that $|M_j^{\mathcal{S}}| \geq \sum_{k=1}^m |S_j^k| = c/n \sum_{k=1}^m |V_j^k| = c/n \times |M_j|$. Thus, it remains to show that $|M_j^{\mathcal{S}}| \leq c/n \times |M_j|$. Indeed, if we assume otherwise we can infer via an exchange argument than M_j is not a maximum matching. $\qquad\square$
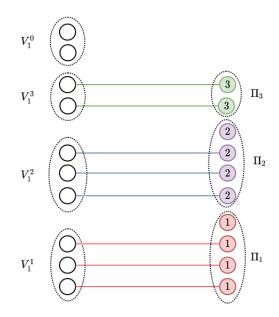
Figure 5: An example of a matching decomposition in the integral domination graph of candidate 1.

Let us denote with $\Phi_j = M_j/n$; roughly speaking, we know from [GHS20] that Φ_j is a good indicator of the "quality" of candidate j. Importantly, Observation 5.6 tells us that we can determine Φ_j in a much smaller graph, if only we had a sampling-decomposition that satisfied the "proportionality" condition of the claim. Of course, determining explicitly such a decomposition makes little sense given that we do not know the sets V_j^k, but the main observation is that we can approximately satisfy the condition of Observation 5.6 through sampling. It should be noted that we previously assumed that $\Pi_j' \in \mathbb{N}$, i.e. we ignored rounding errors. However, in the worst-case rounding errors can only induce an error of at most m/c in the value of Φ_j; thus, we remark that our subsequent selection of c will be such that this error will be innocuous, in the sense that it will be subsumed by the "sampling error" (see Lemma 5.8). Before we proceed, recall that for $\mathbf{p}, \widehat{\mathbf{p}} \in \Delta([k])$,

$$\mathrm{d}_{\mathrm{TV}}(\mathbf{p}, \widehat{\mathbf{p}}) \stackrel{\mathrm{def}}{=} \sup_{S \subseteq [k]} |\mathbf{p}(S) - \widehat{\mathbf{p}}(S)| = \frac{1}{2}||\mathbf{p} - \widehat{\mathbf{p}}||_1, \tag{17}$$

where $|| \cdot ||_1$ represents the ℓ_1 norm. In this context, we will use the following standard fact:

Lemma 5.7 ([Can20])**.** *Consider a discrete distribution $\mathbf{p} \in \Delta([k])$ and let $\widehat{\mathbf{p}}$ be the empirical distribution derived from N independent samples. For any $\epsilon > 0$ and $\delta \in (0,1)$, if $N = \Theta((k + \log(1/\delta))/\epsilon^2)$ it follows that $\mathrm{d}_{\mathrm{TV}}(\mathbf{p}, \widehat{\mathbf{p}}) \leq \epsilon$ with probability at least $1 - \delta$.*

As a result, if we draw a set $\mathcal{S}$ with $|\mathcal{S}| = c = \Theta((m + \log(1/\delta))/\epsilon^2)$ samples (*without* replacement[3]) we can guarantee that

$$\sum_{k=0}^{m} \left| \frac{|S_j^k|}{c} - \frac{|V_j^k|}{n} \right| \leq 2\epsilon; \tag{18}$$

$$\sum_{k=1}^{m} \left| \frac{\widehat{\Pi}_k}{c} - \frac{\Pi_k}{n} \right| \leq 2\epsilon, \tag{19}$$

---

[3]Although the samples are not independent since we are not replacing them, observe that the induced bias is negligible for n substantially larger than m.

where S_j^k represents the subset of $\mathcal{S}$ which intersects V_j^k, and $\widehat{\Pi}_k$ is the empirical plurality score of candidate k. Thus, the following lemma follows directly from Observation 5.6 and Lemma 5.7.

Lemma 5.8. *Let* $\widehat{\Phi}_j = |\widehat{M}_j|/c$, *where* $\widehat{M}_j$ *is the maximum matching in the graph* $G^{\mathcal{S}}(j)$. *Then, if* $|\mathcal{S}| = \Theta((m + \log(1/\delta))/\epsilon^2)$ *for some* $\epsilon, \delta \in (0, 1)$, *it follows that* $(1 - \epsilon)\Phi_j \leq \widehat{\Phi}_j \leq (1 + \epsilon)\Phi_j$ *with probability at least* $1 - \delta$.

Theorem 5.9. *For any* $\epsilon \in (0, 4]$ *and* $\delta \in (0, 1)$ *there exists a mechanism which takes a sample of size* $\Theta((m + \log(m/\delta))/\epsilon^2)$ *voters and yields distortion at most* $3 + \epsilon$ *with probability at least* $1 - \delta$.

Proof. Fix some $\epsilon \in (0, 1/4)$ and $\delta \in (0, 1)$. If we draw $\Theta((m + \log(m/\delta))/\epsilon^2)$ samples Lemma 5.8 along with the union bound imply that $(1-\epsilon)\Phi_j \leq \widehat{\Phi}_j \leq (1+\epsilon)\Phi_j$ for all $j \in [m]$, with probability at least $1 - \delta$, where $\widehat{\Phi}_j$ is defined as in Lemma 5.8. Conditioned on the success of this event, let $w = \arg\max_{j \in C} \widehat{\Phi}_j$. Proposition 5.5 implies that there exists some candidate x for which $\Phi_x = 1$; hence, we know that $\widehat{\Phi}_w \geq \widehat{\Phi}_x \geq 1 - \epsilon$, in turn implying that $\Phi_w \geq (1-\epsilon)/(1+\epsilon) \geq 1 - 2\epsilon$ (Lemma 5.8). As a result, it follows that there exists a subset of voters V' for which there exists a perfect matching in the integral domination graph $G(w)$, with $|V'| \geq n(1 - 4\epsilon)$. Thus, it follows that for the subset of voters in V' candidate w yields distortion at most 3 (see [GHS20]), and Theorem 4.3 leads to the desired conclusion. $\square$

6 Experiments

6.1 Synthetic Data

Here we present several synthetic experiments which illustrate the degradation of distortion under missing information. We will employ the MINIMAX-LP rule, which recall is instance-optimal under any given preferences $\mathcal{P}$ (Theorem 2.2). Also note that the MINIMAX-LP rule returns the exact distortion of *every* candidate w.r.t. the given preferences. The induced LPs in MINIMAX-LP rule are solved via the `Gurobi` software [GO21].

6.1.1 Top Preferences

First we assume that every voter provides only her k-top preferences to the mechanism, and we illustrate the decay of distortion while k gradually increases from 1 to m. Specifically, our experiments are conducted for $n = 50$ voters with preferences sampled from a uniform distribution over the space of permutations. In the parlor of social choice this probabilistic model is referred to as *impartial culture*,[4] and arguably it is unrealistic [TRG03]; nonetheless it will suffice for sketching the underlying qualitative behavior. Recall that according to Theorem 4.5 we expect the distortion to decay as $\mathcal{O}(m/k)$. The results are shown in Figure 6 for $m \in \{4, 6, 8, 10\}$; it should also be noted that for every case we presented 5 different random realizations in order to alleviate the "bias" in the input. Interestingly, the "bottom half" of the voters' preferences appears to offer no real improvement.

6.1.2 Missing Voters

Next we illustrate the decay of distortion as more voters provide information to the mechanism; here it will be assumed that (active) voters provide their total orders to the mechanism. As before, we consider $n = 50$ voters from an impartial culture probabilistic model with $m \in \{3, 4, 5, 6\}$; for every case we consider 5 different random realizations. The results are illustrated in Figure 7. Again, the observed curves closely match the theoretical predictions of Theorem 4.3.

---

[4]To put it differently, impartial culture corresponds to the *Mallows* model [MAL57] with unitary *spread parameter*.
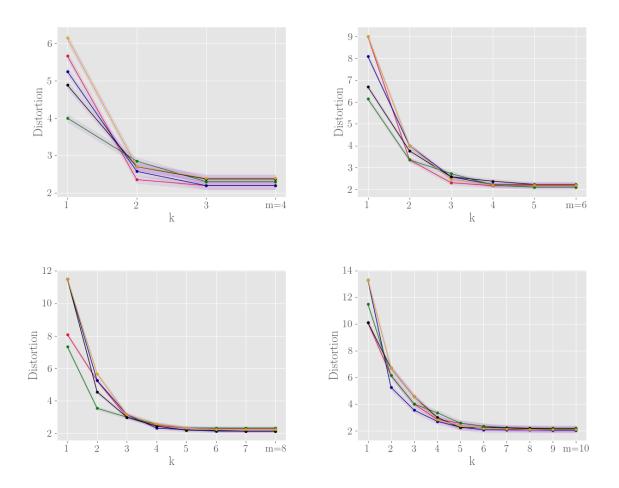
Figure 6: The decay of distortion as we gradually increase k from 1 to m.

6.2 Real Datasets

In this subsection we present some experiments conducted on real-life voting applications. In particular, we are interested in comparing the efficiency—measured in terms of distortion—of the voting system employed in practice with the instance-optimal mechanism, namely the Minimax-LP mechanism.

6.2.1 Eurovision Song Contest

Here we analyze the performance of the scoring system used in the Eurovision song contest, so let us first give a basic overview of the competition and the voting rule employed. Fist of all, we will only focus on the final stage of the competition, wherein a set of m countries compete amongst each other and a set of n countries—which is a strict superset of the contenders—provide their preferences over the finalists. Eurovision employs a specific *positional scoring* system which works as follows. Every country assigns 12 points to its highest preference, 10 points to its second-highest preference, and from $8-1$ points to each of its next 8 preferences; note that no country can vote for itself. This scoring system shall be referred to as the Scoring rule. It should be noted that the authors in [SE17] quantify the distortion for some specific scoring rules (e.g. the *harmonic rule*). We will make the working hypothesis that for every country the assigned scores correspond to its actual order of preferences. Nonetheless, we stress that the assigned scores of every country have been themselves obtained by *preference aggregation*[5], and

---

[5]For the years we are considering the scores were mainly determined by *televoting*, with some few exceptions.
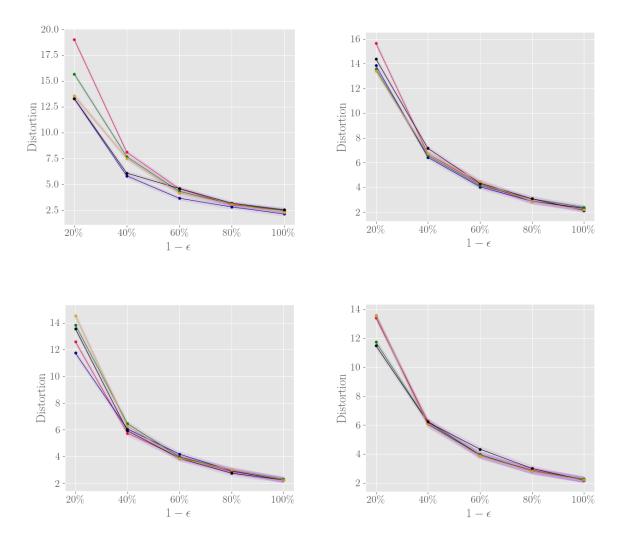
Figure 7: The decay of distortion as more voters (a $(1-\epsilon)$ fraction) provide their preferences to the mechanism, for $\epsilon \in \{0.8, 0.6, 0.4, 0.2, 0\}$. The illustrated curves correspond to $m \in \{3, 4, 5, 6\}$ from top-left, top-right, bottom-left, and bottom-right respectively.

as such they are themselves subject to *distortion*, but we will tacitly suppress this issue.[6]

We will focus on the competitions held between 2004 and 2008; during these years the number of finalists (or candidates) m was 24, with the exception of 2008 where 25 countries were represented in the final. We should note that for our experiments we used a dataset from `Kaggle`. Observe that every "voter" only provides its top $k = 10$ preferences, while the countries which are represented in the final are 0-*decisive* (see [AP16]). The main question that concerns us is whether the Scoring rule employed for the competition yields very different results from the optimal Minimax-LP mechanism. Our results are summarized in Table 1, while for more detailed findings we refer to Appendix C.

Perhaps surprisingly, on all occasions the winners in the two mechanisms coincide; on the other hand, there are generally substantial differences below the first position. It is also interesting to note that on all occasions the winner has a remarkably small distortion, at least compared to the theoretical bounds.

---

[6]We refer the interested reader to the work of Filos-Ratsikas and Voudouris [FV20].

Table 1: Summary of our findings for the Eurovision song contests held between 2004 and 2008. For every year we have indicated the top three countries according to the Minimax-LP rule and the Scoring system employed in the actual contest.

Year	Minimax-LP rule		Scoring rule		# of Countries
	Country	Distortion	Country	Score	
2004	Ukraine	**1.1786**	Ukraine	**280**	36
	Serbia & Montenegro	1.4444	Serbia & Montenegro	263	
	Turkey	1.4746	Greece	252	
2005	Greece	**1.4068**	Greece	**230**	39
	Switzerland	1.4127	Malta	192	
	Moldova	1.4194	Romania	158	
2006	Finland	**1.3000**	Finland	**292**	38
	Romania	1.4262	Russia	248	
	Russia	1.4407	Bosnia & Herzegovina	229	
2007	Serbia	**1.3235**	Serbia	**268**	42
	Ukraine	1.3667	Ukraine	235	
	Russia	1.5231	Russia	207	
2008	Russia	**1.3562**	Russia	**272**	43
	Greece	1.4507	Ukraine	230	
	Ukraine	1.4923	Greece	218	

6.2.2 Formula One

Moreover, we analyze the performance of the voting system employed in the Formula One (F1) world championship. In particular, we imagine that every competing driver constitutes a distinct candidate, while every race in the calendar corresponds to a "voter"; the "preferences" of every race are indicated by the order in which the drivers complete the race. We will assume that when two drivers fail to terminate they will not be comparable (in the spirit of partial orderings). The scoring rule employed in F1 assigns to the first 10 drivers the points $25, 18, 15, 12, 10, 8, 6, 4, 2, 1$ respectively, and the driver who manages to collect the most number of points throughout the championship is declared the winner; with a slight abuse of notation this rule will also be referred to as the Scoring rule. We will be analyzing the championships held between 2016 and 2020, using a dataset from `Kaggle`. A noteworthy detail is that for the last two years the scoring system assigned an additional point to the driver with the fastest lap, but for simplicity the Minimax-LP will not use any such information. Our results are summarized in Table 2. Again, the driver who won the championship is also the candidate who minimizes distortion, with the exception of 2016, where—if we are to accept the metric distortion framework—Lewis Hamilton should have won the championship.

---

[7]Tie with Lando Norris.
[8]Tie with Sebastian Vettel and Charles Leclerc.

Table 2: Summary of our findings for the F1 world championships held between 2016 and 2020. For every year we have indicated the top three drivers according to the MINIMAX-LP rule and the SCORING system employed.

Year	MINIMAX-LP rule		SCORING rule		# of Drivers	# of Races
	Driver	Distortion	Driver	Score		
2020	Lewis Hamilton	**1.6667**	Lewis Hamilton	**347**	23	17
	Valtteri Bottas	5	Valtteri Bottas	223		
	Max Verstappen[7]	5.6667	Max Verstappen	214		
2019	Lewis Hamilton	**1.7059**	Lewis Hamilton	**413**	20	21
	Valtteri Bottas	4	Valtteri Bottas	326		
	Max Verstappen[8]	4.4	Max Verstappen	278		
2018	Lewis Hamilton	**2**	Lewis Hamilton	**408**	20	21
	Sebastian Vettel	3.6	Sebastian Vettel	320		
	Kimi Räikkönen	4.4	Kimi Räikkönen	251		
2017	Lewis Hamilton	**2.2**	Lewis Hamilton	**363**	25	20
	Sebastian Vettel	3	Sebastian Vettel	317		
	Valtteri Bottas	3.1818	Valtteri Bottas	305		
2016	Lewis Hamilton	**2.8333**	Nico Rosberg	**385**	24	21
	Nico Rosberg	3	Lewis Hamilton	380		
	Daniel Ricciardo	3.9091	Daniel Ricciardo	256		

7 Open Problems

There are several compelling avenues for future research related to our work. First, it would be interesting to study the performance of the DOMINATIONROOT mechanism under randomized pairings; we suspect that this might lead to a substantial improvement since our lower bound (Proposition 3.6) is very brittle, but we did not pursue this direction. As we previously alluded to, in practice the pairings are typically constructed using some form of prior, so it might be interesting to formalize the guarantees of such techniques. It would also be meaningful to quantify the decay of distortion from $\mathcal{O}(\log m)$ to $\mathcal{O}(1)$ (which is the bound achievable when the mechanism has access to the entire tournament graph) if we gradually elicit more than $m - 1$ pairwise comparisons. With regards to the power of deterministic mechanisms which elicit only the k-top preferences, an obvious question is to settle Conjecture 4.6. As we showed in Proposition 4.7 this would immediately improve our upper bound, but it would still require some further work to close the gap for every value of $k \in [m]$. Finally, can we reduce the sample complexity established in Theorem 5.9 without sacrificing the efficiency? We argued that the dependence on ϵ and δ cannot be improved, but establishing the optimal dependence on the value of m requires future research.

References

[ABFV20]   Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. Peeking behind the ordinal curtain: Improving distortion via cardinal queries. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 1782–1789. AAAI Press, 2020.

[ABP15]   Elliot Anshelevich, Onkar Bhardwaj, and John Postl. Approximating optimal social choice under metric preferences. In Blai Bonet and Sven Koenig, editors, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, 2015*, pages 777–783. AAAI Press, 2015.

[AFRSV21]   Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A. Voudouris. Distortion in social choice problems: The first 15 years and beyond, 2021.

[AP16]   Elliot Anshelevich and John Postl. Randomized social choice functions under metric preferences. In Subbarao Kambhampati, editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016*, pages 46–59. IJCAI/AAAI Press, 2016.

[AS16a]   Elliot Anshelevich and Shreyas Sekar. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In Dale Schuurmans and Michael P. Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 390–396. AAAI Press, 2016.

[AS16b]   Elliot Anshelevich and Shreyas Sekar. Truthful mechanisms for matching and clustering in an ordinal world. In Yang Cai and Adrian Vetta, editors, *Web and Internet Economics - 12th International Conference, WINE 2016*, volume 10123 of *Lecture Notes in Computer Science*, pages 265–278. Springer, 2016.

[AZ18]   Elliot Anshelevich and Wennan Zhu. Ordinal approximation for social choice, matching, and facility location problems given candidate positions. In George Christodoulou and Tobias Harks, editors, *Web and Internet Economics - 14th International Conference, WINE 2018*, volume 11316 of *Lecture Notes in Computer Science*, pages 3–20. Springer, 2018.

[BCH⁺15]   Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. *Artif. Intell.*, 227:190–213, 2015.

[BCJ11]   Gleb Beliakov, Tomasa Calvo, and Simon James. *Aggregation of Preferences in Recommender Systems*, pages 705–734. Springer US, 2011.

[BDG18]   Umang Bhaskar, Varsha Dani, and Abheek Ghosh. Truthful and near-optimal mechanisms for welfare maximization in multi-winner elections. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, pages 925–932. AAAI Press, 2018.

[BLP04]   Salem Benferhat, Sylvain Lagrue, and Odile Papini. Reasoning with partially ordered information in a possibilistic logic framework. *Fuzzy Sets and Systems*, 144(1):25–41, 2004. Possibilistic Logic and Related Issues.

[BNPS17]   Gerdus Benade, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. Preference elicitation for participatory budgeting. In Satinder P. Singh and Shaul Markovitch,

editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 376–382. AAAI Press, 2017.

[BPQ19]   Gerdus Benadè, Ariel D. Procaccia, and Mingda Qiao. Low-distortion social welfare functions. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019*, pages 1788–1795. AAAI Press, 2019.

[Can20]   Clément L. Canonne. A short note on learning discrete distributions, 2020.

[CE03]   Sungdai Cho and James W. Endersby. Issues, the spatial theory of voting, and british general elections: A comparison of proximity and directional models. *Public Choice*, 114(3/4):275–293, 2003.

[CEG95]   Ran Canetti, Guy Even, and Oded Goldreich. Lower bounds for sampling algorithms for estimating the average. *Inf. Process. Lett.*, 53(1):17–25, 1995.

[CFF⁺16]   Ioannis Caragiannis, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Kristoffer Arnsfelt Hansen, and Zihan Tan. Truthful facility assignment with resource augmentation: An exact analysis of serial dictatorship. In Yang Cai and Adrian Vetta, editors, *Web and Internet Economics - 12th International Conference, WINE 2016*, volume 10123 of *Lecture Notes in Computer Science*, pages 236–250. Springer, 2016.

[CFNV18]   Ioannis Caragiannis, Aris Filos-Ratsikas, Swaprava Nath, and Alexandros A. Voudouris. Truthful mechanisms for ownership transfer with expert advice. *CoRR*, abs/1802.01308, 2018.

[CLWA13]   Shuwei Chen, Jun Liu, Hui Wang, and Juan Carlos Augusto. Ordering based decision making – a survey. *Information Fusion*, 14(4):521–531, 2013.

[CNPS16]   Ioannis Caragiannis, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. Subset selection via implicit utilitarian voting. In Subbarao Kambhampati, editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 151–157. IJCAI/AAAI Press, 2016.

[DB15]   Palash Dey and Arnab Bhattacharyya. Sample complexity for winner prediction in elections. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, page 1421–1430, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.

[FFG16]   Michal Feldman, Amos Fiat, and Iddan Golomb. On voting and facility location. In Vincent Conitzer, Dirk Bergemann, and Yiling Chen, editors, *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16*, pages 269–286. ACM, 2016.

[FFZ14]   Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. Social welfare in one-sided matchings: Random priority and beyond. In Ron Lavi, editor, *Algorithmic Game Theory - 7th International Symposium, SAGT 2014*, volume 8768 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2014.

[FGMP19]   Brandon Fain, Ashish Goel, Kamesh Munagala, and Nina Prabhu. Random dictators with a random referee: Constant sample complexity mechanisms for social choice. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019*, pages 1893–1900. AAAI Press, 2019.

[FGMS17]   Brandon Fain, Ashish Goel, Kamesh Munagala, and Sukolsak Sakshuwong. Sequential deliberation for social choice. In Nikhil R. Devanur and Pinyan Lu, editors, *Web and Internet Economics - 13th International Conference, WINE 2017*, volume 10660 of *Lecture Notes in Computer Science*, pages 177–190. Springer, 2017.

[FH10]   Johannes Frnkranz and Eyke Hllermeier. *Preference Learning.* Springer-Verlag, Berlin, Heidelberg, 1st edition, 2010.

[FKS21]   Dimitris Fotakis, Alkis Kalavasis, and Konstantinos Stavropoulos. Aggregating incomplete and noisy rankings. In Arindam Banerjee and Kenji Fukumizu, editors, *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, pages 2278–2286. PMLR, 2021.

[FV20]   Aris Filos-Ratsikas and Alexandros A. Voudouris. Approximate mechanism design for distributed facility location. *CoRR*, abs/2007.06304, 2020.

[GAX17]   Stephen Gross, Elliot Anshelevich, and Lirong Xia. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In Satinder P. Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, 2017*, pages 544–550. AAAI Press, 2017.

[GHS20]   Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. Resolving the optimal metric distortion conjecture. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020*, pages 1427–1438. IEEE, 2020.

[GKM17]   Ashish Goel, Anilesh K. Krishnaswamy, and Kamesh Munagala. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, page 287–304. Association for Computing Machinery, 2017.

[GO21]   LLC Gurobi Optimization. Gurobi optimizer reference manual, 2021.

[Kar84]   Narendra Karmarkar. A new polynomial-time algorithm for linear programming. *Comb.*, 4(4):373–396, 1984.

[Kem20a]   David Kempe. An analysis framework for metric voting based on LP duality. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 2079–2086. AAAI Press, 2020.

[Kem20b]   David Kempe. Communication, distortion, and randomness in metric voting. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 2087–2094. AAAI Press, 2020.

[KR08]   Leonid (Aryeh) Kontorovich and Kavita Ramanan. Concentration inequalities for dependent random variables via the martingale method. *The Annals of Probability*, 36(6):2126 – 2158, 2008.

[LN95]   Jonathan Levin and Barry Nalebuff. An introduction to vote-counting schemes. *Journal of Economic Perspectives*, 9(1):3–26, March 1995.

[LN05]   Gaëlle Legrand and Nicolas Nicoloyannis. Feature selection method using preferences aggregation. In Petra Perner and Atsushi Imiya, editors, *Machine Learning and Data Mining in Pattern Recognition*, pages 203–217, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

[LPR+07] Jérôme Lang, Maria Silvia Pini, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Winner determination in sequential majority voting. In Manuela M. Veloso, editor, *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 1372–1377, 2007.

[MAL57] C. L. MALLOWS. NON-NULL RANKING MODELS. I. *Biometrika*, 44(1-2):114–130, 06 1957.

[Mou86] H. Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.

[MPSW19] Debmalya Mandal, Ariel D. Procaccia, Nisarg Shah, and David P. Woodruff. Efficient and thrifty voting by any means necessary. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019*, pages 7178–7189, 2019.

[MSW20] Debmalya Mandal, Nisarg Shah, and David P. Woodruff. Optimal communication-distortion tradeoff in voting. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20, page 795–813. Association for Computing Machinery, 2020.

[MW19] Kamesh Munagala and Kangning Wang. Improved metric distortion for deterministic social choice rules. In Anna Karlin, Nicole Immorlica, and Ramesh Johari, editors, *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019*, pages 245–262. ACM, 2019.

[PR06] Ariel D. Procaccia and Jeffrey S. Rosenschein. The distortion of cardinal preferences in voting. In Matthias Klusch, Michael Rovatsos, and Terry R. Payne, editors, *Cooperative Information Agents X, 10th International Workshop, CIA 2006*, volume 4149 of *Lecture Notes in Computer Science*, pages 317–331. Springer, 2006.

[PZPR09] Ariel D. Procaccia, Aviv Zohar, Yoni Peleg, and Jeffrey S. Rosenschein. The learnability of voting rules. *Artif. Intell.*, 173(12-13):1133–1149, 2009.

[SE17] Piotr Krzysztof Skowron and Edith Elkind. Social choice under metric preferences: Scoring rules and STV. In Satinder P. Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 706–712. AAAI Press, 2017.

[SM96] Donald G. Saari and Vincent R. Merlin. The copeland method: I.: Relationships and the dictionary. *Economic Theory*, 8(1):51–76, 1996.

[TRG03] Ilia Tsetlin, Michel Regenwetter, and Bernard Grofman. The impartial culture maximizes the probability of majority cycles. *Social Choice and Welfare*, 21(3):387–398, 2003.

[VLZ12] Maksims Volkovs, Hugo Larochelle, and Richard S. Zemel. Learning to rank by aggregating expert preferences. In Xue-wen Chen, Guy Lebanon, Haixun Wang, and Mohammed J. Zaki, editors, *21st ACM International Conference on Information and Knowledge Management, CIKM'12, 2012*, pages 843–851. ACM, 2012.

[vNM44] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

[VZ14] Maksims Volkovs and Richard S. Zemel. New learning methods for supervised and unsupervised preference aggregation. *J. Mach. Learn. Res.*, 15(1):1135–1176, 2014.

A Lower Bound for k-top Preferences

In this section we present for completeness a lower bound for deterministic mechanisms which have access only to the k-top preferences of every voter.

Proposition A.1. *Any deterministic mechanism which elicits only the k-top preferences from every voter out of the m alternatives has distortion $\Omega(m/k)$.*

Proof. First of all, assume without any loss of generality that $k \mid (m-1)$[9], and let $n = (m-1)/k$ be the number of voters. For simplicity, let us enumerate the number of candidates as $C = \{1, 2, \ldots, n \times k\} \cup \{x\}$. Now consider some preference profile $\mathcal{P}$ in which the k-top preferences of voter i correspond to the of candidates $\{(i-1)k + 1, \ldots, (i-1)k + k\}$ according to some arbitrary order; observe that all of these sets are pairwise disjoint.

Based on these preferences the mechanism has to select some candidate. If x is selected the lower bound follows trivially since x could actually be the last choice for every voter. Therefore, let us assume that candidate 1 was selected by the mechanism; this hypothesis is without loss of generality due to the symmetry of the input $\mathcal{P}$. However, the agents and the candidates could be located on the metric space of Figure 8; indeed, it is easy to check that the induced metric space is consistent with the given preferences. As a result, it follows that

$$\frac{\mathrm{SC}(1)}{\mathrm{SC}(x)} = \frac{D + (n-1) \times (\delta + 2D)}{D + (n-1) \times \delta} = \frac{1 + (n-1) \times (\delta/D + 2)}{1 + (n-1) \times \delta/D}. \tag{20}$$

Thus, for $\delta/D \downarrow 0$ we obtain that $\mathrm{SC}(1)/\mathrm{SC}(x) \to 2n - 1 = \Omega(m/k)$. $\square$
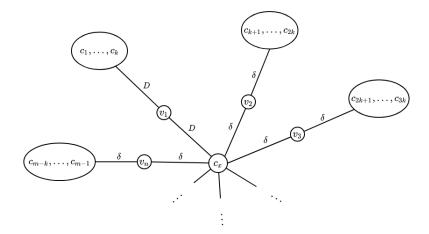

Figure 8: The metric space considered for the proof of Proposition A.1, where $\delta/D \downarrow 0$ for some positive numbers δ and D. Naturally, the distance between two points is simply the shortest path in the graph.

We should note that although in our worst-case example the number of voters n is smaller than the number of candidates m, which is not the canonical case, our argument directly extends whenever n is a multiple of $(m-1)/k$, allowing n to be arbitrarily large. Moreover, a similar construction shows an $\Omega(m/k)$ lower bound for α-*decisive* metrics [AP16], for any $\alpha \in [0, 1]$; indeed, it suffices to place the voters within the "cluster" of their k-most preferred candidates.

---

[9]If it is not the case that $k \mid (m-1)$ take k' to be the smallest number larger than k such that $k' \mid (m-1)$, and apply our argument for k'; given that $k' < 2k$ we will establish again a lower bound of $\Omega(m/k)$ even though the mechanism had more information than the k-top preferences.

B PluralityMatching vs Minimax-LP

In this section we compare the PluralityMatching mechanism of Gkatzelis et al. [GHS20] with the instance-optimal mechanism, namely Minimax-LP; in this section we tacitly posit that $\mathcal{P} = \sigma$, i.e. all the agents provide their entire rankings to the mechanism.

B.1 Instance Optimality

The first question that arises is how far could the distortion of PluralityMatching be with respect to the instance-optimal candidate; to this end, we commence with the following proposition:

Lemma B.1 (Lemma 6, [GHS20]). *For any election $\mathcal{E} = (V, C, \sigma)$, a candidate $a \in C$ can be selected by* PluralityMatching *only if* $\mathrm{plu}(a) \geq \mathrm{veto}(a)$.

With this lemma in mind, we consider an instance with a set of m voters $V = \{1, 2, \ldots, m\}$, and a set of m candidates $C = \{a, \ldots\}$. We assume that for every voter $i \in [n-1], \mathrm{sec}(i) = a$, while the (single) top-preferences of all the voters $i \in [n-1]$ are assumed to be pairwise-distinct. Finally, the last voter places candidate a in her last place, while her preferences are otherwise arbitrary. An example with four candidates $\{a, b, e, f\}$ corresponds to the following input:

- $b \succ_1 a \succ_1 e \succ_1 f$;

- $f \succ_2 a \succ_2 b \succ_2 e$;

- $e \succ_3 a \succ_3 f \succ_3 b$;

- $b \succ_4 f \succ_4 e \succ_4 a$.

In general, observe that for any candidate $b \in C \setminus \{a\}$ it follows that $|ab| = (m-2)/m$. Moreover, we will use the following standard lemma:

Lemma B.2. *Consider two (distinct) candidates $a, b \in C$ such that $|ab| \geq \alpha n > 0$. Then,*

$$\frac{\mathrm{SC}(a)}{\mathrm{SC}(b)} \leq \frac{2}{\alpha} - 1. \tag{21}$$

This implies that the distortion of candidate a is $1 + \mathcal{O}(1/m)$. However, given that $\mathrm{plu}(a) = 0 < 1 = \mathrm{veto}(a)$, we know from Lemma B.1 that a cannot be selected by PluralityMatching. We will show that every other candidate yields distortion close to 3. In particular, consider the metric space illustrated in Figure 9. It is easy to verify that the induced metric space is consistent with the given preferences. But, it follows that $\mathrm{SC}(a) = m$, while $\mathrm{SC}(b) = 2 + 3(m-2) = 3m - 4$ for any $b \neq a$, implying that $\mathrm{SC}(b)/\mathrm{SC}(a) = 3 - \mathcal{O}(1/m)$. As a result, we have arrived at the following conclusion:

Proposition B.3. *For any sufficiently small $\epsilon > 0$ and $m = \mathcal{O}(1/\epsilon)$ there exists a preferences profile σ such that* Minimax-LP *yields distortion $1 + \epsilon$, while* PluralityMatching *yields distortion at least $3 - \epsilon$.*

B.2 Decisive Metrics

Moreover, it is natural to compare these mechanisms in more refined metrics. Specifically, we espouse the α-decisiveness assumption of Anshelevich and Postl [AP16], according to which $d(v_i, c_p) \leq \alpha \cdot d(v_i, c_q)$, where $p = \mathrm{top}(i)$ and $q = \mathrm{sec}(i)$, and $\alpha \in [0, 1]$ some parameter; notice that the general case corresponds to $\alpha = 1$, while for $\alpha = 0$ every voter also serves as a
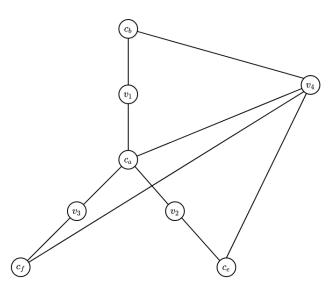
Figure 9: A metric space embedded on an *unweighted* and *undirected* graph; this example corresponds to $m = n = 4$, but the pattern should be already clear.

candidate. The first observation is that this particular refinement can be addresses by simply incorporating some additional constraints in the linear program. More precisely, for a pair of distinct candidates a, b this leads to the following linear program $\text{METRIC}^\alpha\text{-LP}(a, b)$:

$$
\begin{aligned}
\text{maximize} \quad & \sum_{i=1}^n x_{i,a} \\
\text{subject to} \quad & \sum_{i=1}^n x_{i,b} = 1; \\
& x_{i,\text{top}(i)} \leq \alpha \cdot x_{i,\text{sec}(i)}, \forall i \in V; \\
& x_{i,p} \leq x_{i,q}, \qquad \forall (p, q) \in \mathcal{P}_i, \forall i \in V; \\
& x_{i,i} = 0, \qquad \forall i \in V \cup C; \\
& x_{i,j} = x_{j,i}, \qquad \forall i, j \in V \cup C; \\
& x_{i,j} \leq x_{i,k} + x_{k,j}, \quad \forall i, j, k \in V \cup C.
\end{aligned}
\tag{22}
$$

Here we have assumed that every agent i provides her most preferred candidate $\text{top}(i)$, as well as her second most preferred candidate $\text{sec}(i)$. Having solved the $\text{METRIC}^\alpha\text{-LP}(a, b)$ for every distinct pair of candidates a, b, we simply select the candidate who minimizes the maximum cost obtained over all other candidates; this mechanism shall be referred to as the $\text{MINIMAX}^\alpha\text{-LP}$. Similarly to Theorem 2.2 we can establish the following:

Proposition B.4. *For any given preferences $\mathcal{P}$ and any $\alpha \in [0, 1]$ the $\text{MINIMAX}^\alpha\text{-LP}$ rule is instance-optimal in terms of distortion under α-decisive metrics.*

We should point out that for α-decisive metrics PLURALITYMATCHING always yields a candidate with distortion $2 + \alpha$. Moreover, Gkatzelis et al. [GHS20] showed a lower bound of $2 + \alpha - 2(1-\alpha)/m'$ for deterministic mechanisms, where $m' = 2\lfloor m/2 \rfloor$; thus, they showed that their mechanism obtains the optimal distortion only when $m \to \infty$ or when $\alpha = 1$, leaving a substantial gap.

We will show that $\text{MINIMAX}^\alpha\text{-LP}$ can substantially outperform PLURALITYMATCHING even for α-decisive metrics with α close to 0. Specifically, consider an election with 3 candidates and

2 voters[10] with the following preferences: $\sigma_1 = a \succ b \succ e$, and $\sigma_2 = e \succ b \succ a$. For this election, b could be returned by PLURALITYMATCHING (see [GHS20]). However, we claim that b yields distortion $2 + \alpha$, while a and e have distortion $1 + 2\alpha$. Indeed, we will show that candidate a has always distortion upper-bounded by $1 + 2\alpha$ (by symmetry, the same holds for e), while for candidate b there exists a metric space for which b yields distortion $2 + \alpha$. Specifically, we have that $d(c_a, c_b) \leq d(v_1, c_a) + d(v_1, c_b) \leq (1 + \alpha)d(v_1, c_b)$; thus we obtain that

$$d(v_1, c_a) \leq \alpha d(v_1, c_b), \tag{23}$$

$$d(v_2, c_a) \leq d(v_2, c_b) + d(c_a, c_b) \leq (1 + \alpha)d(v_1, c_b) + d(v_2, c_b). \tag{24}$$

Summing these inequalities yields that $\text{SC}(a) \leq (1 + 2\alpha)d(v_1, c_b) + d(v_2, c_b) \leq (1 + 2\alpha)\,\text{SC}(b)$. Similarly, we can prove that $\text{SC}(a) \leq (1 + 2\alpha)\,\text{SC}(e)$. On the other hand, for candidate b Gkatzelis et al. [GHS20] considered the following metric space:
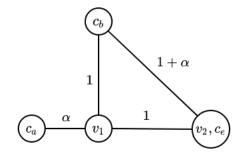

Figure 10: A metric space embedded on a graph.

Naturally, the distance between a pair of nodes is the corresponding shortest path in the graph. Thus, for this instance it follows that $\text{SC}(e) = 1$, while $\text{SC}(b) = 2 + \alpha$, implying that the distortion of b is $2 + \alpha$. Thus, for $\alpha \to 0$ PLURALITYMATCHING loses a factor of 2 with respect to the optimal candidate, which would be identified by the MINIMAX$^\alpha$-LP rule by virtue of Proposition B.4.

Proposition B.5. *There exists a preference profile σ such that* MINIMAX$^\alpha$-LP *yields distortion* $1 + 2\alpha$, *while* PLURALITYMATCHING *yields distortion at least* $2 + \alpha$ *under α-decisive metrics.*

Nonetheless we should point out that PLURALITYMATCHING does *not* require knowing the value of parameter α, unlike the instance-optimal mechanism.

C   Additional Experiments

In this section we provide additional details about our empirical findings. Specifically, in Figure 11 we juxtapose the scores of the finalists in the Eurovision song contest (based on the SCORING rule) with their distortion as determined by MINIMAX-LP. We note that we have removed the following outliers:

- For the year 2006 the countries Malta and Spain which incurred a distortion of 24.3333 and 18.0000 respectively.

- For the year 2007 the countries Ireland and United Kingdom which incurred a distortion of 27.0000 and 13.6667 respectively.
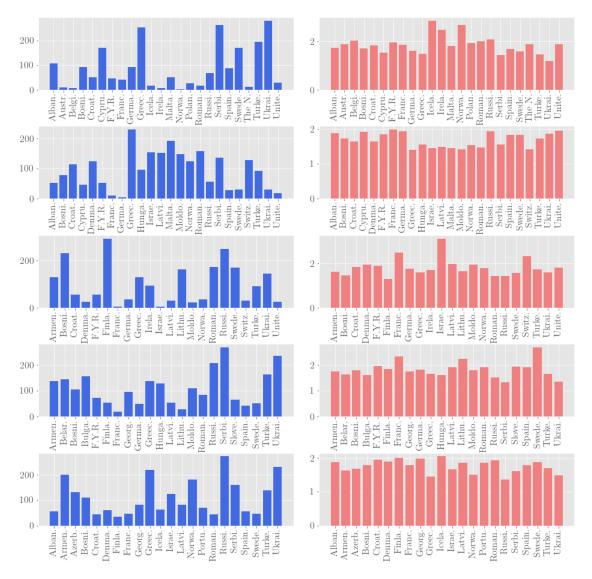
---

[10]This example is taken from [GHS20].

Figure 11: Scoring rule vs Minimax-LP for the Eurovision song contests during the years between 2004 and 2008. For every country we have indicated only the first 5 letters according to the entry in the dataset.

- For the year 2008 the countries United Kingdom, Germany and Poland which incurred a distortion of 20.5000, 20.5000 and 14 respectively.

Similarly, in Figure 12 we juxtapose the scores accumulated by the drives based on the Scoring rule with their distortion as determined by Minimax-LP. These results follow after removing the following outliers:

- For the year 2016 the driver Stoffel Vandoorne who incurred a distortion of 41.

- For the year 2017 the drivers Paul di Resta, Jenson Button, Brendon Hartley, and Antonio Giovinazzi who incurred a distortion of ∞, ∞, 19, and 39 respectively.

- For the year 2020 the drivers Pietro Fittipaldi, Jack Aitken, and Nico Hulkenberg who incurred a distortion of 16, 33, and 16 respectively.
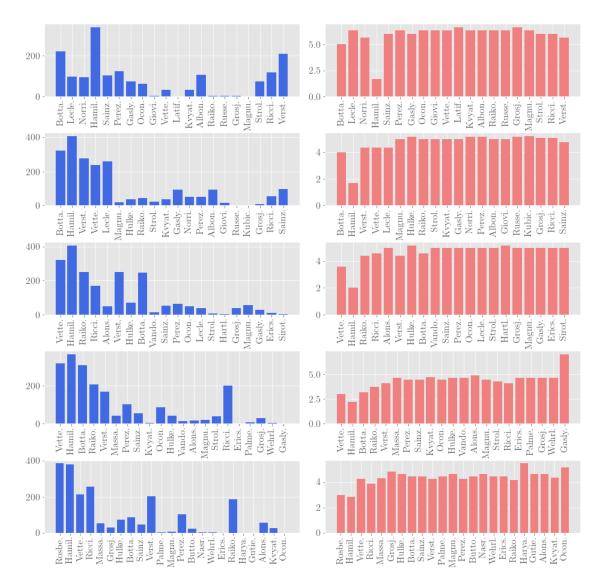
Figure 12: SCORING rule vs MINIMAX-LP for the F1 world championships during the years between 2020 and 2016. For every driver we have indicated only the first 5 letters of his/her (last) name according to the entry in the dataset.