

Peeking Behind the Ordinal Curtain: Improving Distortion via Cardinal Queries*

Georgios Amanatidis¹, Georgios Birmpas², Aris Filos-Ratsikas³ and Alexandros A. Voudouris²

¹Centrum Wiskunde & Informatica

²University of Oxford

³École Polytechnique Fédérale de Lausanne

Abstract

Aggregating the preferences of individuals into a collective decision is the core subject of study of social choice theory. In 2006, Procaccia and Rosenschein considered a utilitarian social choice setting, where the agents have explicit numerical values for the alternatives, yet they only report their linear orderings over them. To compare different aggregation mechanisms, Procaccia and Rosenschein introduced the notion of *distortion*, which quantifies the inefficiency of using only ordinal information when trying to maximize the social welfare, i.e., the sum of the underlying values of the agents for the chosen outcome. Since then, this research area has flourished and bounds on the distortion have been obtained for a wide variety of fundamental scenarios. However, the vast majority of the existing literature is focused on the case where nothing is known beyond the ordinal preferences of the agents over the alternatives. In this paper, we take a more expressive approach, and consider mechanisms that are allowed to further ask a *few cardinal queries* in order to gain partial access to the underlying values that the agents have for the alternatives. With this extra power, we design new *deterministic* mechanisms that achieve significantly improved distortion bounds and, in many cases, outperform the best-known randomized ordinal mechanisms. We paint an almost complete picture of the number of queries required to achieve specific distortion bounds.

1 Introduction

Social choice theory [Brandt *et al.*, 2016] is concerned with aggregating the preferences of individuals into a joint decision. In an election, for instance, the winner should represent well (in some precise sense) the viewpoints of the voters. Similarly, the expenditure of public funds is typically geared towards projects that increase the well-being of society. Most traditional models assume that the preferences of individuals are expressed through *ordinal preference rankings*, where each agent sorts all alternatives from the most to the least favorable according to her. Underlying these ordinal preferences, it is often assumed that there

*Georgios Amanatidis is supported by the NWO Gravitation Project NETWORKS, Grant Number 024.002.003, the ERC Advanced Grant 788893 AMDROMA "Algorithmic and Mechanism Design Research in Online Markets" and MIUR PRIN project ALGADIMAR "Algorithms, Games, and Digital Markets". Georgios Birmpas is supported by the European Research Council (ERC) under the advanced grant number 321171 (ALGAME). Aris Filos-Ratsikas is supported by the Swiss National Science Foundation under contract number 200021_165522. Alexandros Voudouris is supported by the European Research Council (ERC) under grant number 639945 (ACCORD).

exists a *cardinal* utility structure, which further specifies the intensity of the preferences [Von Neumann and Morgenstern, 1947; Bogomolnaia and Moulin, 2001; Barbera *et al.*, 1998]. That is, there exist numerical values that indicate how much an agent prefers an outcome to another. Given this cardinal utility structure, usually expressed via *valuation functions*, one can define meaningful quantitative objectives, with the most prominent one being the maximization of the *utilitarian* (or *social*) *welfare*, i.e., the sum of the values of the agents for the chosen outcome.

The main rationale justifying the dominance of ordinal preferences in the classical economics literature is that the task of asking individuals to express their preferences in terms of numerical values is arguably quite demanding for them. In contrast, performing simple comparisons between the different options is certainly more easily conceivable. To quantify how much the lack of cardinal information affects the maximization of quantitative objectives like the social welfare, Procaccia and Rosenschein [2006] defined the notion of *distortion* for mechanisms as the worst-case ratio between the optimal social welfare (which would be achievable using cardinal information) and the social welfare of the outcome selected by the mechanism, which has access only to the preference rankings of the agents. Following their agenda, a plethora of subsequent works studied the distortion of mechanisms in several different settings, such as normalized valuation functions [Boutilier *et al.*, 2015], metric preferences [Anshelevich *et al.*, 2018; Anshelevich and Postl, 2017], committee elections [Caragiannis *et al.*, 2017], and participatory budgeting [Benade *et al.*, 2017].

Somewhat surprisingly, the different variants of the distortion framework studied in this rich line of work differentiate between two extremes: we either have complete cardinal information or only ordinal information. Driven by the original motivation for using ordinal preferences, it seems quite meaningful to ask whether improved distortion guarantees can be obtained if one has access to *limited* cardinal information, especially in settings for which the worst-case distortion bounds are already quite discouraging [Boutilier *et al.*, 2015]. We formulate this idea via the use of *cardinal queries*, which elicit cardinal information from the agents. These queries can be as simple as asking the value of an agent for a possible outcome, or even asking an agent *whether an outcome is at least x times better than some other outcome*, according to her underlying valuation function. Note that questions of the latter form are much less demanding than eliciting a complete cardinal utility structure, and thus are much more realistic as an elicitation device (see also the discussion below).

In this paper, we enhance the original distortion setting of Procaccia and Rosenschein [2006] and Boutilier *et al.* [2015] on single winner elections, by allowing the use of cardinal queries. In their setting, there are n agents that have cardinal values over m alternatives, and the goal is to elect a single alternative that (approximately) maximizes the social welfare, while having access *only* to ordinal information. Procaccia and Rosenschein [2006] proved that no deterministic mechanism can achieve a distortion better than $\Omega(m)$ when agents have *unit-sum* normalized valuation functions (i.e., the sum of the values of each agent for all possible alternatives is 1), which was later on improved to $\Omega(m^2)$ by Caragiannis *et al.* [2017]. Under the same assumption, Boutilier *et al.* [2015] proved that the distortion of any (possibly randomized) mechanism is between $\Omega(\sqrt{m})$ and $O(\sqrt{m} \cdot \log^* m)$. Here we show how – with only a limited number of cardinal queries – deterministic mechanisms can significantly outperform any mechanism that has access only to ordinal information, even randomized ones.

1.1 Our Contributions

We initiate the study of trade-offs between the number of cardinal queries per agent that a mechanism uses and the distortion that it can achieve. In particular, we show results of the following type:

The distortion $\mathcal{D}(\mathcal{M})$ of a mechanism \mathcal{M} that makes at most λ queries per agent is $O(g(m, \lambda))$.

What our results suggest is that we can drastically reduce the distortion by exploiting only a small amount of cardinal information.

Query Model

We consider two different types of cardinal queries, namely *value queries* and *comparison queries*.

- A *value query* takes as input an agent i and an alternative j , and returns the agent’s value for that alternative.
- A *comparison query* takes as input an agent i , two alternatives j, ℓ and a real number d and returns “yes” if the value of agent i for alternative j is at least d times her value for alternative ℓ , and “no” otherwise.

Note that value queries are in general stronger than comparison queries, as they reveal much more detailed information. On the other hand, comparison queries are quite attractive as an elicitation device, since the cognitive complexity of the question that they pose is not much higher than that of forming a preference ranking. Additionally, comparison queries can also be interpreted under the original utility framework defined by [Von Neumann and Morgenstern \[1947\]](#). The idea there is that a cardinal scale for utility is possible because agents are capable of not only performing comparisons between alternatives, but also between lotteries over alternatives. For example, an agent i should be able to tell whether she prefers alternative a with certainty, or alternative b with probability $1/2$. Assuming *risk-neutrality*, this is equivalent to asking the comparison query with parameters $(i, a, b, 1/2)$.

It should be clear that upper bounds (distortion guarantees) for comparison queries are stronger than those for value queries, whereas lower bounds (inapproximability bounds) are stronger when proven for value queries. All of our lower bounds are for value queries, while our main upper bounds extend to comparison queries as well.

Results and Techniques

We warm-up in [Section 3](#) by using λ simple prefix value queries per agent (i.e., ask her at the first λ positions of her preference ranking). By selecting the alternative with the highest social welfare restricted to the query answers (the *revealed welfare*), we obtain a linear improvement in the distortion, specifically $1 + (m - 1)/\lambda$. We show that this result is asymptotically optimal, among all mechanisms that use λ prefix queries per agent.

In [Section 4](#), we devise a class of more sophisticated mechanisms that achieve much improved trade-offs between the distortion and the number of queries. In particular, our class contains

- a mechanism that achieves *constant* distortion using at most $O(\log^2 m)$ queries per agent, and
- a mechanism that achieves a distortion of $O(\sqrt{m})$ using $O(\log m)$ queries, matching the performance of the best possible randomized mechanism in the setting of [\[Boutilier et al., 2015\]](#), and outperforming all known randomized mechanisms for that setting.

Our mechanisms are based on a binary search procedure, which for every agent finds the last alternative a in the agent’s preference ranking such that the agent’s value for a is at least $1/k$ times the value for her most-preferred alternative α^* , for some chosen parameter k . Then, the mechanism *simulates* the value of the agent for all alternatives that the agent ranks between α^* and a by her value for a , and outputs the alternative that maximizes the *simulated welfare*. By repeatedly applying this idea for appropriately chosen values of

k , we explore the trade-offs between the distortion and the number of queries, when the latter range from $\log m$ to $\log^2 m$ per agent.

In [Section 5](#), we significantly improve on our $O(\sqrt{m})$ result (second bullet above) for the fundamental case $n = \Theta(m)$. We present a mechanism that achieves a distortion of $O(\sqrt{m})$, using only 2 queries per agent. The core idea behind this mechanism is choosing an appropriate threshold τ and then carefully querying the agents based on this threshold. First, the mechanism queries every agent at the first position, and the remaining queries (one per agent) are made in successive steps. During the ℓ -th step, the mechanism queries about alternatives that are ranked at the first ℓ positions by at least τ agents, but *only* if such a query is meaningful and possible; we never repeat a query and we never ask an agent more than twice. The query process terminates in at most m steps and the mechanism returns an alternative with maximum revealed welfare. This result demonstrates that with the clever use of very limited cardinal information, one can outperform all possible mechanisms (even randomized) in the purely ordinal setting.

In [Section 6](#) we extend the ideas of [Section 4](#) to show that the mechanism which achieves a constant distortion using $O(\log^2 m)$ value queries, can actually be transformed into a mechanism which uses the same number of *comparison queries*. In particular, we show how to approximate an agent’s value for her most-preferred alternative using only $O(\log^2 m)$ comparison queries.

In [Section 7](#) we present several lower bounds on the possible achievable trade-offs between the number of queries and distortion. These bounds follow by explicit instances where we carefully define a single ordinal preference profile as well as the cardinal information that may be revealed by the value queries of any mechanism. This information is defined in such a way so that, no matter how the mechanism makes its selection, it is always possible to create a superconstant gap between the optimal social welfare and the social welfare of the winning alternative.

We conclude the paper in [Section 8](#) with several interesting open problems, and a particular set of very challenging conjectures about the tight trade-offs between the number of queries and distortion.

An overview of our main results can be found in [Table 1](#). An alternative representation of our results is given in [Figure 1](#), which depicts the trade-offs between the number of queries and the distortion.

Remark 1 (Normalization assumptions). We remark here that all of our upper bounds for value queries hold *without any normalization assumption on the cardinal values*, contrary to the results of [[Procaccia and Rosenschein, 2006](#)] and almost all subsequent works in the related literature, which typically assume that values are normalized according to the unit-sum normalization. We do use the unit-sum normalization in [Section 6](#), where we use comparison queries.¹ For the lower bounds, we prove bounds both for normalized and unrestricted values.

Remark 2 (Noisy queries). Throughout this work we implicitly assume that agents can accurately answer all value or comparison queries. In fact, this is not necessary for any of our positive results! That is, we may assume that the answers to the queries are *noisy*, e.g., because it requires extra effort for the agents to precisely determine these answers. As long as each inaccurate answer is at most a (multiplicative) constant factor away from the truth, all our upper bound proofs go through, at the expense of worse constants. Note that lower bounds are stronger when proven for exact queries, as is the case here.

¹Actually, our results hold even if one uses other reasonable normalizations. In particular, for the other common normalization assumption in the literature [[Caragiannis et al., 2018](#); [Feige and Tennenholtz, 2010](#); [Filos-Ratsikas and Miltersen, 2014](#)], the unit-range normalization, where the value of an agent for her most-preferred alternative is 1 and all other values are in the interval $[0, 1]$, the results of [Section 4](#) obviously extend verbatim to the case of comparison queries.

Number of queries	Upper Bounds	Lower Bounds
0 (ordinal, deterministic)	$O(m^2)$ [Caragiannis and Procaccia, 2011]	$\Omega(m^2)$ [Caragiannis <i>et al.</i> , 2017]
0 (ordinal, randomized)	$O(\sqrt{m} \log^* m)$ [Boutilier <i>et al.</i> , 2015]	$\Omega(\sqrt{m})$ [Boutilier <i>et al.</i> , 2015]
1 (value)	$O(m)$ [1-PRV, Theorem 1]	$\Omega(m)$ [Theorem 9] $\Omega(\sqrt{m})$ [Theorem 11]
$\lambda \geq 2$, constant (value)	$O(\sqrt{m})$ [\sqrt{m} -TRV, Theorem 6]*	$\Omega(m^{1/2(\lambda+1)})$ [Corollary 5]
$O\left(\frac{\log m}{\log \log m}\right)$ (value)	$O(\sqrt{m})$ [\sqrt{m} -TRV, Theorem 6]*	$\Omega(\log \log m)$ [Corollary 5]
$O(\log m)$ (value)	$O(\sqrt{m})$ [$O(1)$ -ARV, Corollary 2]	$\Omega(1)$
$O(\log^2 m)$ (value)	$O(1)$ [$O(\log m)$ -ARV, Corollary 2]	$\Omega(1)$
$O(\log^2 m)$ (comparison)	$O(1)$ [$O(\log m)$ -ARV, Corollary 4]	$\Omega(1)$

Table 1: A table showing the most important results in the paper. All our results are for *deterministic* mechanisms. Results marked by \star hold for $n = \Theta(m)$. Results for unit-sum valuation functions are highlighted; everything else is for unrestricted valuation functions.




Figure 1: A graphical representation of the trade-offs between the number of queries and the distortion. Points and lines in red represent upper bounds, while points and lines in blue represent lower bounds. The black points correspond to tight bounds, and the red square point corresponds to the distortion of the best-known randomized ordinal mechanism.

1.2 Related Work

The distortion framework was introduced by Procaccia and Rosenschein [2006], and has been studied subsequently in a series of papers, most prominently by Boutilier *et al.* [2015], who consider a general social choice setting, under the unit-sum normalization; this general model was also previously studied by Caragiannis and Procaccia [2011] who considered different methods to translate the values of the agents for the

alternatives into rankings (embeddings), and more recently by [Filos-Ratsikas et al. \[2019\]](#) who bounded the distortion of deterministic mechanisms in district-based elections. A related model is that of distortion of social choice functions in a metric space, which was initiated by [Anshelevich et al. \[2018\]](#), and has since then been studied extensively [[Anshelevich and Postl, 2017](#); [Goel et al., 2017](#); [Fain et al., 2019](#); [Goel et al., 2018](#); [Anshelevich and Zhu, 2018](#); [Pierczynski and Skowron, 2019](#); [Gross et al., 2017](#); [Cheng et al., 2017, 2018](#); [Feldman et al., 2016](#); [Ghods et al., 2019](#); [Borodin et al., 2019](#); [Munagala and Wang, 2019](#)]. In this setting, there is no normalization of values (or costs), but the valuation (or cost) functions are assumed to satisfy the triangle inequality. Similar distortion frameworks, in a metric space or under normalizations, have also been studied for other related problems, such as matching and clustering [[Anshelevich and Sekar, 2016](#); [Abramowitz and Anshelevich, 2018](#); [Anshelevich and Zhu, 2017, 2018](#); [Filos-Ratsikas et al., 2014](#)].

Two related variants of the problem are k -winner elections, where k alternatives are to be elected instead of one [[Caragiannis et al., 2017](#); [Benade et al., 2019](#)], and participatory budgeting, where every alternative is associated with a cost, and one or more alternatives have to be elected in a manner that ensures that the total cost does not exceed a pre-specified budget constraint [[Lu and Boutilier, 2011](#)]. [Benade et al. \[2017\]](#) studied the k -winner participatory budgeting problem, but interestingly, they considered a more expressive model for the preferences of the agents, compared to simple preference rankings. In particular, they considered the knapsack votes model of [[Goel et al., 2016](#)], rankings by value, rankings by value-for-money and threshold votes. While the first three are not very relevant for our purposes, the latter one can be thought of as a different type of (more expressive) query, in which a numerical value is specified, and every agent is asked to return the set of alternatives for which her value is above this threshold. [Bhaskar et al. \[2018\]](#) used a different model with thresholds drawn from $\mathcal{U}[0, 1]$ to construct a randomized social choice function that approaches a distortion of 1 as the number of agents approaches infinity.

Very recently, [Mandal et al. \[2019\]](#) study a related model to ours, in which agents are asked to provide cardinal information, but there is a restriction on the number of bits to be communicated to the mechanism. Hence, they study trade-offs between the number of transmitted bits and distortion. This is markedly quite different from what we do here, as a query in their setting has access to the (approximate) values of an agent for many alternatives simultaneously, and is therefore much too expressive when translated in our setting. On the other hand, the setting of [Mandal et al. \[2019\]](#) does not assume “free” access to the ordinal preferences, which are also considered as part of the elicitation process. We consider our work complementary to theirs, as they are mostly motivated by the computational limitations of elicitation (corresponding to a communication complexity approach), whereas we are motivated by the cognitive limitations of eliciting cardinal values, as often highlighted in the classical literature of social choice (corresponding to a query complexity approach).

Finally, at the same time and independently of our work, [Abramowitz et al. \[2019\]](#) also introduce a setting in which the mechanism designer has access to some cardinal information on top of the ordinal preferences. This enables the design of improved mechanisms in terms of distortion. While the motivation of their paper is the same as ours, the approaches are inherently different. Besides the fact that [Abramowitz et al. \[2019\]](#) study a metric distortion setting, whereas we study a general setting with valuation functions that which are either unrestricted or normalized according to unit-sum, there is another fundamental distinction. The access to the cardinal information in [[Abramowitz et al., 2019](#)] is not via queries. Instead, it is given explicitly as part of the input in terms of a threshold τ , which allows the designer to know the number of agents for which the distance to an alternative a is at most $1/\tau$ times their distance to another alternative b .

2 The model

We consider a standard *social choice setting*, in which there is a set A of m alternatives and a set N of n agents. Our goal is to elect a *single* alternative based on the preferences of the agents, which are expressed through *valuation functions* $v_i : A \rightarrow \mathbb{R}_{\geq 0}$ that map alternatives to non-negative real numbers. For notational convenience, we use v_{ij} instead of $v_i(j)$ to denote the *cardinal value* of agent i for alternative j , and refer to the matrix $\mathbf{v} = (v_{ij})_{i \in N, j \in A}$ as a *valuation profile*. By \mathbf{V} we denote the set of all possible valuation profiles. Clearly, the valuation function v_i also defines a *preference ranking* for agent i , i.e., a linear ordering $>_i$ of A such that $j >_i j'$ if $v_{ij} \geq v_{ij'}$; we assume that ties are broken according to a deterministic tie-breaking rule, e.g., according to a fixed global ordering of the alternatives.² We refer to $\succ_{\mathbf{v}} = (>_1, \dots, >_n)$ as an (ordinal) *preference profile*.

In this work, we consider the following two families of valuation functions:

- *Unrestricted valuation functions*, which may take any non-negative real values.
- *Unit-sum valuation functions*, which are such that $\sum_{j \in A} v_{ij} = 1$ for every agent $i \in N$.

The *social welfare* of alternative $j \in A$ with respect to \mathbf{v} is the total value of the agents for j : $\text{SW}(j|\mathbf{v}) = \sum_{i \in N} v_{ij}$. Our goal is to output one of the alternatives who maximize the social welfare, i.e., an alternative in $\arg \max_{j \in A} \text{SW}(j|\mathbf{v})$. This is clearly a trivial task if one has full access to the valuation profile. However, we assume *limited access* to these cardinal values. In particular, we assume that we only have access to the preference profile $\succ_{\mathbf{v}}$ and can also learn cardinal information by asking queries. We consider two types of queries: *value queries* that reveal the value of an agent for a given alternative, and *comparison queries* that reveal whether the value of an agent for an alternative is a multiplicative factor larger than her value for some other alternative.

Definition 1. Given a preference profile, a query about the underlying cardinal values is called

- A *value query*, if it takes as input an agent i and an alternative j and returns the agent's value v_{ij} for that alternative. This is implemented via the function $\mathcal{V} : N \times A \rightarrow \mathbb{R}_{\geq 0}$. We say that *agent i is queried at position k* , if alternative j is ranked k -th in $>_i$ and we make the query $\mathcal{V}(i, j)$.
- A *comparison query*, if it takes as input an agent i , two alternatives j, ℓ and a real number d , and returns YES if $v_{ij} \geq d \cdot v_{i\ell}$, and NO otherwise. This is implemented via the function $\mathcal{C} : N \times A \times A \times \mathbb{R}_{\geq 0} \rightarrow \{\text{YES}, \text{NO}\}$.

Clearly, value queries reveal more information than comparison queries. Note that the information obtained by a comparison query can be obtained by at most two value queries. On the other hand, however, without any cardinal information or any normalization assumption, it is impossible to even approximate the information obtained by a value query using only comparison queries. In this sense, value queries are considerably stronger than comparison queries.

Definition 2. A *mechanism* $\mathcal{M} = (Q, f)$ with access to a (value or comparison) oracle takes as input a preference profile $\succ_{\mathbf{v}}$ and returns an alternative. In particular, it consists of the following two parts:

- An algorithm Q that takes as input the preference profile $\succ_{\mathbf{v}}$, adaptively makes queries to the oracle, and returns the set of answers to these queries.

²It would be equivalent to allow ties at this point, get pre-linear orderings instead, and leave the tie-breaking to the mechanisms when necessary.

- A mapping f that takes as input the preference profile $\succ_{\mathbf{v}}$ and the set $Q(\succ_{\mathbf{v}})$ of answers to the queries above, and outputs a single alternative $j \in A$. Such a mapping is called a *social choice function*.

By the description of Q above, it is clear that the mechanism is free to choose the positions at which each agent will be queried, and those can depend not only on $\succ_{\mathbf{v}}$, but on the answers to the queries already asked as well. The performance of a mechanism is measured by its *distortion*.

Definition 3. The *distortion* of a mechanism \mathcal{M} is

$$\mathcal{D}(\mathcal{M}) = \sup_{\mathbf{v} \in \mathbf{V}} \frac{\max_{j \in A} \text{SW}(j|\mathbf{v})}{\text{SW}(\mathcal{M}(\succ_{\mathbf{v}})|\mathbf{v})},$$

where $\text{SW}(j|\mathbf{v})$ is the social welfare of alternative j given a particular valuation profile, and $\mathcal{M}(\succ_{\mathbf{v}})$ is the output of the mechanism on input $\succ_{\mathbf{v}}$.

Throughout our proofs, it will be useful to partition the quantity $\text{SW}(j|\mathbf{v})$, into two separate quantities depending on the cardinal information we obtain from the queries. This is particularly relevant when we deal with value queries, but even for comparison queries we use a similar decomposition in Section 6.

Definition 4. The *revealed welfare* $\text{SW}_r(j|\mathbf{v})$ of j is the contribution to $\text{SW}(j|\mathbf{v})$ of agents that have been queried for alternative j via value queries, i.e., $\text{SW}_r(j|\mathbf{v}) = \sum_{i \in N: \mathcal{V}(i,j) \in Q(\succ_{\mathbf{v}})} v_{ij}$. The remaining quantity $\text{SW}(j|\mathbf{v}) - \text{SW}_r(j|\mathbf{v})$ is called the *concealed welfare* $\text{SW}_c(j|\mathbf{v})$ of j .

3 Warm-Up: Mechanisms Using Fixed-Position Value Queries

We start the presentation of our technical results with the class of mechanisms that query every agent at the first $\lambda \geq 1$ positions. A particular member of this class is the mechanism that uses the *Range Voting* (RV) social choice function to decide the outcome. Formally, RV takes as input the whole valuation profile \mathbf{v} and elects an alternative x with maximum social welfare: $x \in \arg \max_{j \in A} \text{SW}(j|\mathbf{v})$. In our case, since \mathbf{v} is not fully known, we deploy RV only on the revealed valuation profile, where any unknown value is assumed to be zero.

To be more specific, let $T_k(j)$ be the set of agents that rank alternative $j \in A$ at position $k \in [m]$. Our mechanism first queries every agent at each of the first λ positions of her preference ranking. Then, it elects the alternative y that maximizes the revealed welfare: $y \in \arg \max_{j \in A} \text{SW}_r(j|\mathbf{v})$. We refer to this mechanism as *λ -Prefix Range Voting* (λ -PRV).

Mechanism λ -PRV($\succ_{\mathbf{v}}$)

```

1 for  $j \in A$  do
2    $\text{SW}_r(j|\mathbf{v}) = 0$ 
3 for  $i \in N$  do
4   for  $k \in [\lambda]$  do
5     Make a query to learn  $\mathcal{V}(i, j_i(k))$ , where  $j_i(k)$  is the  $k$ -th favorite alternative of agent  $i$ 
6      $\text{SW}_r(j_i(k)|\mathbf{v}) = \text{SW}_r(j_i(k)|\mathbf{v}) + \mathcal{V}(i, j_i(k))$ 
7 Let  $y \in \arg \max_{j \in A} \text{SW}_r(j|\mathbf{v})$  be an alternative achieving the best revealed welfare
8 return  $y$ 

```

Theorem 1. *The distortion of λ -PRV is $\mathcal{D}(\lambda\text{-PRV}) \leq 1 + \frac{m-1}{\lambda}$, even for unrestricted valuation functions.*

Proof. Consider some instance with valuation profile \mathbf{v} . Let x be an alternative that maximizes the social welfare according to \mathbf{v} , and let y be the alternative that is elected by λ -PRV. Recall that here the revealed welfare of any alternative $j \in A$ is $\text{SW}_r(j|\mathbf{v}) = \sum_{k=1}^{\lambda} \sum_{i \in T_k(j)} v_{ij}$. Since $\text{SW}(y|\mathbf{v}) \geq \text{SW}_r(y|\mathbf{v})$, it suffices to show that $\text{SW}(x|\mathbf{v}) \leq (1 + \frac{m-1}{\lambda}) \text{SW}_r(y|\mathbf{v})$. To this end, we will bound the revealed and the concealed welfare of x separately.

Since y is an alternative that maximizes the revealed welfare, we have that $\text{SW}_r(y|\mathbf{v}) \geq \text{SW}_r(j|\mathbf{v})$ for every $j \in A$, and therefore

$$\text{SW}_r(x|\mathbf{v}) \leq \text{SW}_r(y|\mathbf{v}). \quad (1)$$

Now, consider the agents in $\bigcup_{k=\lambda+1}^m T_k(x)$. They are not queried about their value for x , and therefore contribute to the concealed welfare of x . For every such agent i there exist λ different alternatives $j_i(1), \dots, j_i(\lambda)$ that i ranks above x , and for whom she has value $v_{i,j_i(1)}, \dots, v_{i,j_i(\lambda)} \geq v_{ix}$.³ Consequently, we have that

$$\begin{aligned} \text{SW}_c(x|\mathbf{v}) &= \sum_{k=\lambda+1}^m \sum_{i \in T_k(x)} v_{ix} \leq \sum_{k=\lambda+1}^m \sum_{i \in T_k(x)} \frac{v_{i,j_i(1)} + \dots + v_{i,j_i(\lambda)}}{\lambda} \\ &= \frac{1}{\lambda} \sum_{j \in A \setminus \{x\}} \text{SW}_r(j|\mathbf{v}) \leq \frac{1}{\lambda} \sum_{j \in A \setminus \{x\}} \text{SW}_r(y|\mathbf{v}) = \frac{m-1}{\lambda} \text{SW}_r(y|\mathbf{v}). \end{aligned} \quad (2)$$

The statement now follows by (1) and (2). \square

Clearly, the distortion guarantee of λ -PRV improves linearly in the number of queries λ . Nevertheless, it is interesting to see for which values of λ the mechanism achieves distortion $O(\sqrt{m})$ and $O(1)$. These are given by the following statement.

Corollary 1. *The distortion of λ -PRV is*

$$\mathcal{D}(\lambda\text{-PRV}) = \begin{cases} O(\sqrt{m}), & \text{for } \lambda = \Theta(\sqrt{m}) \\ O(1), & \text{for } \lambda = \Theta(m) \end{cases}$$

Next, we show that, in terms of distortion, λ -PRV is the best possible mechanism among those that make at most λ prefix value queries.

Theorem 2. *Any mechanism that makes λ prefix value queries per agent has distortion $\Omega(m/\lambda)$, even for unit-sum valuation functions.*

Proof. Consider an instance with n agents and $m = n$ alternatives a_1, \dots, a_m . Let $\lambda \leq m/2$. We define the following ordinal profile:

- The λ favorite alternatives of agent i are $a_i, a_{i+1}, \dots, a_{i+\lambda-1}$ in decreasing order, where all the indices are considered modulo m . Hence, all alternatives appear exactly once at each of the first λ positions.
- Alternatives $x = a_1$ and $y = a_{\lambda+1}$ appear $m/2$ times each at position $(\lambda + 1)$, in the $m - \lambda \geq m/2$ agent rankings in which they do not appear at the first λ positions. Observe that, by definition, x and y do not appear together at the first λ positions in any preference ranking, and there are multiple ways to decide in which rankings each of them appears at position $(\lambda + 1)$; any such construction works for our purposes.

³When the subscripts have subscripts themselves, we follow the common practice of separating them with commas.

- For every agent, the remaining alternatives are arbitrarily ordered at positions $(\lambda + 2)$ up to m .

See Table 2 for a specific example of the ordinal profile.

agent	ranking					
1	a_1	a_2	a_3	a_4	a_5	a_6
2	a_2	a_3	a_1	a_4	a_5	a_6
3	a_3	a_4	a_1	a_2	a_5	a_6
4	a_4	a_5	a_3	a_1	a_2	a_6
5	a_5	a_6	a_1	a_3	a_2	a_4
6	a_6	a_1	a_3	a_2	a_4	a_5

Table 2: An example of the ordinal profile used in the proof of Theorem 2 with $m = n = 6$ and $\lambda = 2$, where $x = a_1$, $y = a_3$, $T_{\lambda+1}(x) = \{2, 3, 5\}$ and $T_{\lambda+1}(y) = \{1, 4, 6\}$.

The valuation profile \mathbf{v} is such that each agent has value $\frac{1}{\lambda+1}$ for her first λ favorite alternatives. It is without loss of generality to assume that any mechanism that knows the ordinal information of this instance and also makes λ prefix value queries, must elect either x or y . To see this, first notice that given the revealed cardinal information, the revealed welfare of all alternatives is the same. Further, given the particular preference profile, it is easy to always complete the valuation profile \mathbf{v} in a way that guarantees that no alternative has more concealed welfare than x and y ; indeed, the two possible valuation profiles we consider have this property.

So, assume that the mechanism selects alternative y (the case of x being completely symmetric). Now the remaining values of the agents are such that the $m/2$ agents in $T_{\lambda+1}(x)$ have value $\frac{1}{\lambda+1}$ for x and 0 for the remaining alternatives, while the $m/2$ agents in $T_{\lambda+1}(y)$ have value $\frac{1}{(m-\lambda)(\lambda+1)}$ for all alternatives at positions $\lambda + 1$ up to m .

Given this valuation profile \mathbf{v} , the social welfare of the winner y is

$$\text{SW}(y|\mathbf{v}) = \frac{\lambda}{\lambda + 1} + \frac{\frac{m}{2}}{(m - \lambda)(\lambda + 1)} \leq 1.$$

In contrast, the social welfare of the optimal alternative x is

$$\text{SW}(x|\mathbf{v}) = \frac{\lambda + \frac{m}{2}}{\lambda + 1} + \frac{\frac{m}{2}}{(m - \lambda)(\lambda + 1)} \geq \frac{m}{2(\lambda + 1)}.$$

Therefore, the distortion of any mechanism is at least $\frac{m}{2(\lambda+1)}$. \square

We now turn our attention to a slightly more general class of mechanisms which query all agents at the same fixed positions, and show that λ -PRV remains best possible among the mechanisms of this class for unrestricted valuation functions. In Section 7 we further show that 1-PRV is best possible among all mechanisms that make one query per agent for unrestricted valuation functions.

Theorem 3. *For unrestricted valuation functions, any mechanism that makes λ fixed-position value queries per agent has distortion $\Omega(m/\lambda)$.*

Proof. Let $\lambda \leq m/2$. Consider any mechanism of this class, and let ℓ be the first position at which it does not query the agents. Observe that if $\ell > \lambda$, then the mechanism only makes prefix value queries. In this case, the bound follows by Theorem 2, which holds for unit-sum valuation functions, and thus for unrestricted ones as well. So, we may assume that $\ell \in [\lambda]$.

Now, we consider an instance with $n = m$ that is very similar to the one presented in the proof of Theorem 2. Essentially, we substitute $(\lambda + 1)$ with ℓ , and we have that all alternatives appear exactly once at each of the first $\ell - 1$ positions, while two alternatives x and y appear $m/2$ times each at position ℓ . The remaining alternatives for every agent are arbitrarily ordered at position $\ell + 1$ up to m .

The valuation profile \mathbf{v} is such that each agent has value 1 for her first $\ell - 1$ favorite alternatives, and value 0 for the alternatives at positions $\ell + 1$ up to m . Observe that the revealed welfare of all alternatives is exactly equal to $\ell - 1$. Given the revealed cardinal information and the particular ordinal profile, we can argue exactly like we did in the proof of Theorem 2 about fact that it is without loss of generality to assume that the mechanism elects either x or y . So, assume that the mechanism selects alternative y ; the case of x is symmetric. The remaining values of the agents are such that the $m/2$ agents in $T_\ell(x)$ have value 1 for x , while the $m/2$ agents in $T_\ell(y)$ have value 0 for y .

Given this valuation profile \mathbf{v} , the social welfare of the winner y is $\text{SW}(y|\mathbf{v}) = \ell - 1 \leq \lambda - 1$, while the social welfare of the optimal alternative x is $\text{SW}(x|\mathbf{v}) = \ell - 1 + \frac{m}{2} \geq \frac{m}{2}$. Therefore, the distortion of the mechanism is $\Omega(m/\lambda)$. \square

4 Achieving Constant Distortion

Our goal in this section is to further explore the additional power that cardinal queries provide, and focus on the design of mechanisms with improved distortion guarantees. Mechanism λ -PRV is a good first step in this direction, but it needs to make a quite large number of queries per agent in order to do so; in particular, by Corollary 1, it achieves distortion $O(\sqrt{m})$ for $\lambda = \Theta(\sqrt{m})$ and constant distortion for $\lambda = \Theta(m)$. Therefore, it is natural to ask whether it is possible to design mechanisms that achieve similar distortion bounds, but require much less queries per agent. We manage to answer this question positively.

For any $k \in [m]$, we define a mechanism which we call *k-Acceptable Range Voting (k-ARV)*. Let $\lambda_1, \dots, \lambda_k$ be k thresholds such that $\lambda_\ell = m^{\frac{\ell}{k+1}}$ for $\ell \in [k]$. For every agent $i \in N$, we first query her value v_i^* for her favorite alternative $j_i(1)$ and let $S_{i,0} = \{j_i(1)\}$. Then, using binary search we compute the maximal λ_ℓ -acceptable set $S_{i,\ell} = \{j \in A : v_{ij} \geq v_i^*/\lambda_\ell\}$ for every $\ell \in [k]$. We continue by constructing a new approximate valuation profile $\tilde{\mathbf{v}}$, where the values of every agent i are

- $\tilde{v}_i^* = v_i^*$;
- $\tilde{v}_{ij} = v_i^*/\lambda_\ell$ for every $j \in S_{i,\ell} \setminus S_{i,\ell-1}$ with $\ell \in [k]$;
- $\tilde{v}_{ij} = 0$ for every $j \in A \setminus S_{i,k}$.

We finally elect the alternative $z \in A$ that maximizes the social welfare according to the approximate valuation profile: $z \in \arg \max_{j \in A} \sum_{i \in N} \tilde{v}_{ij}$.

Now, we proceed by proving an upper bound on the distortion achieved by k -ARV as a function of k .

Theorem 4. *The mechanism k -ARV makes $O(k \log m)$ direct value queries per agent, and has distortion $\mathcal{D}(k\text{-ARV}) = O(k^{\frac{1}{k+1}}\sqrt{m})$.*

Proof. Consider any instance with valuation profile \mathbf{v} . Since mechanism k -ARV executes a binary search in order to compute the λ_ℓ -acceptable sets for each $\ell \in [k]$, it requires a total of $O(k \log m)$ value queries per

Mechanism k -ARV(\succ_v)

```
1 for  $i \in N$  do
2   Let  $v_i^* = \mathcal{V}(i, j_i(1))$ , where  $j_i(p)$  is the alternative that agent  $i$  ranks at position  $p$ 
3   Let  $\tilde{v}_{i, j_i(1)} = v_i^*$ 
4   Let  $S_{i,0} = \{j_i(1)\}$ 
5   for  $\ell \in \{1, 2, \dots, k\}$  do
6     Let  $\lambda_\ell = m^{\frac{\ell}{k+1}}$ 
7     Set  $p^* = \text{BSearch}(1, m, \lambda_\ell, v_i^*)$ 
8     Let  $S_{i,\ell} = \{j \in A : j \succ_i j_i(p^*)\}$  /* define the  $\lambda_\ell$ -acceptable set of agent  $i$  */
9     for  $j \in S_{i,\ell} \setminus S_{i,\ell-1}$  do
10      Let  $\tilde{v}_{ij} = v_i^* / \lambda_\ell$  /* define the approximate valuation profile */
11  for  $j \in A \setminus S_{i,k}$  do
12    Let  $\tilde{v}_{ij} = 0$ 
13 for  $j \in A$  do
14   Let  $\text{SW}_s(j|\tilde{\mathbf{v}}) = 0$  /* compute the simulated welfare of alternative  $j$  */
15   for  $j \in A$  do
16     Let  $\text{SW}_s(j|\tilde{\mathbf{v}}) = \text{SW}_s(j|\tilde{\mathbf{v}}) + \tilde{v}_{ij}$ 
17 Let  $z \in \arg \max_{j \in A} \text{SW}_s(j|\tilde{\mathbf{v}})$  be an alternative achieving the best simulated welfare
18 return  $z$ 

19 Procedure  $\text{BSearch}(\alpha, \beta, \lambda, v)$ 
20   if  $\alpha = \beta$  then
21     return  $\alpha$ 
22   Let  $u = \mathcal{V}(i, j_i(\frac{\alpha+\beta}{2}))$ 
23   if  $u \geq v/\lambda$  then
24      $\text{BSearch}(\frac{\alpha+\beta}{2}, \beta, \lambda, v)$ 
25   else
26      $\text{BSearch}(\alpha, \frac{\alpha+\beta}{2}, \lambda, v)$ 
```

agent. The rest of the proof is dedicated in bounding the distortion of k -ARV. First, we define some useful notation:

- z is the alternative elected by k -ARV;
- y is a welfare-maximizing alternative for the valuation profile $\hat{\mathbf{v}}$, which is such that the value of agent $i \in N$ for alternative $j \in A$ is

$$\hat{v}_{ij} = \begin{cases} 0, & \text{if } j \in A \setminus S_{i,k} \\ v_{ij}, & \text{otherwise.} \end{cases}$$

That is, $y \in \arg \max_{j \in A} \sum_{i \in N} \hat{v}_{ij}$.

- x is the welfare-maximizing alternative for the true valuation profile \mathbf{v} . That is, $x \in \arg \max_{j \in A} \sum_{i \in N} v_{ij}$.

Also, let $N_j(\mathbf{v}) = \{i \in N : v_{ij} > 0\}$ be the set of agents with strictly positive value for alternative $j \in A$. We use the following easy fact about welfare-maximizing alternatives.

Lemma 1. *If $j \in \arg \max_{j \in A} \sum_{i \in N} v_{ij}$, then $j \in \arg \max_{j \in A} \sum_{i \in N_j(\mathbf{v})} v_{ij}$.*

To prove the statement, we will bound the social welfare of x in terms of the social welfare of z for the true valuation profile \mathbf{v} . In particular, we will show that

$$\text{SW}(x|\mathbf{v}) \leq \left(\lambda_1 + \frac{m}{\lambda_k} \right) \text{SW}(z|\mathbf{v}). \quad (3)$$

Then, the approximation ratio of k -ARV will be

$$\frac{\text{SW}(x|\mathbf{v})}{\text{SW}(z|\mathbf{v})} \leq \lambda_1 + \frac{m}{\lambda_k} = 2 \cdot m^{\frac{1}{k+1}} = O(k^{\frac{1}{k+1}} \sqrt{m}).$$

We partition the social welfare of x into the following two quantities: the contribution of the agents i that place x in the λ_k -acceptable set $S_{i,k}$, and the contribution of the remaining agents that have small value for x . By definition, we have that $i \in N_x(\hat{\mathbf{v}})$ for any agent i such that $x \in S_{i,k}$, and therefore

$$\text{SW}(x|\mathbf{v}) = \sum_{i \in N_x(\hat{\mathbf{v}})} v_{ix} + \sum_{i \notin N_x(\hat{\mathbf{v}})} v_{ix}$$

We first consider the term $\sum_{i \in N_x(\hat{\mathbf{v}})} v_{ix}$, and have that

$$\sum_{i \in N_x(\hat{\mathbf{v}})} v_{ix} \leq \sum_{i \in N_y(\hat{\mathbf{v}})} v_{iy} \leq \lambda_1 \cdot \sum_{i \in N_y(\hat{\mathbf{v}})} \tilde{v}_{iy} \leq \lambda_1 \cdot \sum_{i \in N_z(\tilde{\mathbf{v}})} \tilde{v}_{iz} \leq \lambda_1 \cdot \sum_{i \in N_z(\tilde{\mathbf{v}})} v_{iz} \leq \lambda_1 \cdot \text{SW}(z|\mathbf{v}), \quad (4)$$

where

- the first inequality follows by the definition of y , the simple fact that $N_x(\hat{\mathbf{v}}) = N_y(\hat{\mathbf{v}})$, and Lemma 1;
- for the second inequality it suffices to notice that for any $i \in N_y(\hat{\mathbf{v}})$ there exists an $\ell \in [k]$ such that $y \in S_{i,\ell} \setminus S_{i,\ell-1}$, and thus $v_{ij} \leq \frac{v_i^*}{\lambda_{\ell-1}} = \lambda_1 \cdot \frac{v_i^*}{\lambda_\ell} = \lambda_1 \cdot \tilde{v}_{ij}$;
- the third inequality follows by the definition of z , the simple fact that $N_y(\hat{\mathbf{v}}) = N_z(\tilde{\mathbf{v}})$, and Lemma 1;

- the fourth inequality follows by the fact that $v_{ij} \geq \tilde{v}_{ij}$, for every $i \in N$ and $j \in A$;
- the last inequality is obvious.

Next, we consider the term $\sum_{i \notin N_x(\hat{\mathbf{v}})} v_{ix}$. By the definition of $N_x(\hat{\mathbf{v}})$, for every $i \notin N_x(\hat{\mathbf{v}})$ it holds that $x \notin S_{i,k}$, and hence $v_{ix} < v_i^*/\lambda_k$. Using this, we obtain

$$\sum_{i \notin N_x(\hat{\mathbf{v}})} v_{ix} < \sum_{i \notin N_x(\hat{\mathbf{v}})} \frac{v_i^*}{\lambda_k} = \frac{1}{\lambda_k} \sum_{i \notin N_x(\hat{\mathbf{v}})} v_i^* \leq \frac{1}{\lambda_k} \sum_{j \in A \setminus \{x\}} \sum_{i \in T_1(j)} v_{ij}, \quad (5)$$

where $T_1(j)$ is the set of agents whose favorite alternative is j , and for whom $v_i^* = \tilde{v}_i^* = v_{ij} = \tilde{v}_{ij}$. Since z is the alternative that maximizes the quantity $\sum_{i \in N} \tilde{v}_{ij}$, for every $j \neq z$ we have that

$$\sum_{i \in N} \tilde{v}_{iz} \geq \sum_{i \in N} \tilde{v}_{ij} = \sum_{i \in T_1(j)} v_{ij} + \sum_{i \in N \setminus T_1(j)} \tilde{v}_{ij} \geq \sum_{i \in T_1(j)} v_{ij}.$$

Combining the above inequality together with the fact that $v_{iz} \geq \tilde{v}_{iz}$ for every agent $i \in N$, we have that

$$\sum_{i \in N} v_{iz} \geq \sum_{i \in T_1(j)} v_{ij}.$$

Using this last inequality, (5) becomes

$$\sum_{i \notin N_x(\hat{\mathbf{v}})} v_{ix} \leq \frac{1}{\lambda_k} \sum_{j \in A \setminus \{x\}} \sum_{i \in T_1(j)} v_{ij} \leq \frac{1}{\lambda_k} \sum_{j \in A \setminus \{x\}} \sum_{i \in N} v_{iz} = \frac{m-1}{\lambda_k} \text{SW}(z|\mathbf{v}). \quad (6)$$

Finally, the desired inequality (3) follows by combining inequalities (4) and (6). \square

The next statement follows by appropriately setting the value of the parameter k in Theorem 4, and shows how mechanism k -ARV improves upon the distortion guarantees of λ -PRV using way less value queries per agent.

Corollary 2. *We have that*

- 1-ARV achieves distortion $O(\sqrt{m})$ using $O(\log m)$ values queries per agent;
- $\log m$ -ARV achieves distortion $O(1)$ using $O(\log^2 m)$ value queries.

We conclude this section by showing that the analysis of k -ARV is tight.

Theorem 5. *The distortion of k -ARV is $\Omega(\sqrt[k+1]{m})$.*

Proof. Recall that $\lambda_1 = m^{\frac{1}{k+1}} = \lambda$ and consider the following instance with m alternatives $A = \{a_1, \dots, a_m\}$ and $n = m - 2$ agents. To simplify our discussion, let $z = a_{m-1}$ and $x = a_m$. The valuation profile \mathbf{v} is such that the values of agent i are

- $v_{i,a_i} = v_{ix} = \frac{\lambda}{2\lambda+1}$,
- $v_{iz} = \frac{1}{2\lambda+1}$, and
- $v_{i,a_j} = 0$ for $j \in [m] \setminus \{i, m-1, m\}$.

In the ordinal profile \succ_v which is given as input to the mechanism, we assume without loss of generality that agent i ranks alternative a_i ahead of x .

Since $\frac{1}{2\lambda+1} = \frac{1}{\lambda} \cdot \frac{\lambda}{2\lambda+1}$, k -ARV defines only one acceptable set per agent using λ . In particular, the algorithm sets $S_{i,1} = \{x, z\}$ for every agent $i \in [m-2]$. Then, the approximate valuation profile \tilde{v} is such that the values of agent i are

- $\tilde{v}_{i,a_i} = \frac{\lambda}{2\lambda+1}$,
- $\tilde{v}_{ix} = \tilde{v}_{iz} = \frac{1}{2\lambda+1}$, and
- $\tilde{v}_{i,a_j} = 0$ for $j \in [m] \setminus \{i, m-1, m\}$

For the approximate valuation profile \tilde{v} , the social welfare of both alternatives x and z is

$$\text{SW}(x|\tilde{v}) = \text{SW}(z|\tilde{v}) = \frac{m-2}{2\lambda+1},$$

while any other alternative $j \in A \setminus \{x, z\}$ has social welfare

$$\text{SW}(j|\tilde{v}) = \frac{\lambda}{2\lambda+1}.$$

Hence, k -ARV might select alternative z as the winner instead of x , and the distortion is then

$$\frac{(m-2)\frac{\lambda}{2\lambda+1}}{(m-2)\frac{1}{2\lambda+1}} = \lambda = \sqrt[k+1]{m},$$

as desired. □

5 Achieving \sqrt{m} Distortion with Two Value Queries

Here we present a more sophisticated mechanism, which makes two value queries per agent, and refines the main ideas of the previous two sections. Like before, the first query is used to learn the value of each agent for her favorite alternative. However, we would like to avoid making a naive second query as we do with 2-PRV. Ideally, we would like to ask each agent about an alternative that is qualitatively similar to the one identified by 1-ARV; in other words, we would like to reveal for each agent the position where her value is roughly $1/\sqrt{m}$ of that for her favorite alternative. Although maintaining the same guarantee as 1-ARV, while substituting each binary search with a single query seems far-fetched, we do come very close. By utilizing the available ordinal information globally rather than per agent, our mechanism achieves distortion $O(\sqrt{m})$ with just two value queries, under the assumption that $n = O(m)$. The crucial idea is that the second query for each agent depends on the number of appearances of the alternatives in the whole ordinal preference profile.

For any *threshold* $\tau \in \mathbb{N}$ we define a mechanism called τ -Threshold Range Voting (τ -TRV). Before stating the mechanism formally, we give a short high-level description. As noted above, the first query for each agent is used to ask about her favorite alternative. The remaining queries are made in successive steps. During the ℓ -th step we make queries about alternatives that are ranked within the first ℓ positions by at least τ agents. These queries are made *only* if they are meaningful and possible: we never repeat a query and we never ask an agent more than twice. After at most m steps, τ -TRV returns an alternative with maximum revealed welfare.

We state the mechanism τ -TRV, as well as the main result of this section, Theorem 6, with respect to a general threshold τ . Depending on the ranges of n and m we appropriately set τ to get Corollary 3. Recall that $T_k(j)$ is the set of agents that rank alternative j at position k . Thus, the set \mathcal{E}_ℓ of *eligible alternatives at ℓ -th step* contains all alternatives j for which $|\bigcup_{k=1}^\ell T_k(j)| \geq \tau$.

τ -TRV($\succ_{\mathbf{v}}$)

```

1 for  $j \in A$  do
2    $\lfloor$   $\text{SW}_r(j|\mathbf{v}) = 0$ 
3 for  $i \in N$  do
4    $\lfloor$  Make a query to learn  $\mathcal{V}(i, j_i)$ , where  $j_i$  is the top alternative of agent  $i$ 
5    $\lfloor$   $\text{SW}_r(j_i|\mathbf{v}) = \text{SW}_r(j_i|\mathbf{v}) + \mathcal{V}(i, j_i)$ 
6  $Q_1 = N$ 
7 for  $\ell \in \{2, \dots, m\}$  do
8    $\mathcal{E}_\ell = \{j \in A : |\bigcup_{k=1}^\ell T_k(j)| \geq \tau\}$  /* eligible alternatives of  $\ell$ -th step */
9    $Q_\ell = Q_{\ell-1}$  /* maintain the set of active agents who have been queried once */
10  for  $j \in \mathcal{E}_\ell$  do
11    for  $i \in Q_\ell$  do
12      if  $i \in \bigcup_{k=2}^\ell T_k(j)$  then
13         $\lfloor$  Make a query to learn  $\mathcal{V}(i, j)$ 
14         $\lfloor$   $\text{SW}_r(j|\mathbf{v}) = \text{SW}_r(j|\mathbf{v}) + \mathcal{V}(i, j)$ 
15         $\lfloor$   $Q_\ell = Q_\ell \setminus \{i\}$  /* agent  $i$  becomes permanently inactive */
16 Let  $y \in \arg \max_{j \in M} \text{SW}_r(j|\mathbf{v})$  be an alternative achieving the best revealed welfare
17 return  $y$ 

```

Theorem 6. *The mechanism τ -TRV has distortion $\mathcal{D}(\tau\text{-TRV}) \leq \tau + \min\{m-1, n/\tau\}$, even for unrestricted valuation functions.*

Proof. Consider any instance with valuation profile \mathbf{v} . If $\tau > n$, then no alternative will ever be eligible, and hence τ -TRV coincides with mechanism 1-PRV from Section 3; by Theorem 1, the distortion is at most m . Consequently, in the rest of the proof we focus on the case $\tau \leq n$. Let y be the alternative elected by τ -TRV, and x be an alternative that maximizes the social welfare. By the definition of the revealed welfare, we have $\text{SW}(y|\mathbf{v}) \geq \text{SW}_r(y|\mathbf{v})$. We are going to show that $\text{SW}(x|\mathbf{v}) \leq \text{SW}_r(y|\mathbf{v}) \cdot (\tau + \min\{m-1, n/\tau\})$; then the statement follows.

By partitioning the optimal welfare into the revealed and the concealed welfare of x , we have

$$\begin{aligned} \text{SW}(x|\mathbf{v}) &= \text{SW}_r(x|\mathbf{v}) + \text{SW}_c(x|\mathbf{v}) \\ &\leq \text{SW}_r(y|\mathbf{v}) + \text{SW}_c(x|\mathbf{v}), \end{aligned} \tag{7}$$

where the inequality follows by the optimality of y with respect to the revealed welfare.

Next, we focus on bounding the quantity $\text{SW}_c(x|\mathbf{v})$. Let $\ell^* \in \{2, \dots, m\}$ be such that $x \in \mathcal{E}_{\ell^*}$ but $x \notin \mathcal{E}_{\ell^*-1}$. By definition, for any ℓ we have that $x \in \mathcal{E}_\ell \Rightarrow x \in \mathcal{E}_{\ell+1}$, and thus ℓ^* is unique. In fact, ℓ^* is the smallest step in which alternative x becomes eligible. We can further partition the concealed welfare of x as

$$\text{SW}_c(x|\mathbf{v}) = \text{SW}_c^{<\ell^*}(x|\mathbf{v}) + \text{SW}_c^{\geq\ell^*}(x|\mathbf{v}), \tag{8}$$

where

- $\text{SW}_c^{<\ell^*}(x|\mathbf{v})$ is the contribution to $\text{SW}_c(x|\mathbf{v})$ of agents who rank x at some position $k < \ell^*$;
- $\text{SW}_c^{\geq\ell^*}(x|\mathbf{v})$ is the contribution to $\text{SW}_c(x|\mathbf{v})$ of agents who rank x at some position $k \geq \ell^*$.

If $\ell^* = 2$, then it is straightforward that $\text{SW}_c^{<\ell^*}(x|\mathbf{v}) = 0$. So, assume that $\ell^* \geq 3$. Observe that since $x \notin \mathcal{E}_k$ for any $k \in \{2, \dots, \ell^* - 1\}$, there can be at most $\tau - 1$ agents who rank x before position ℓ^* . If $\hat{i} \in N$ is such an agent, an obvious upper bound for the corresponding concealed value $v_{i,x}$ is $\max_{i \in N, j \in A} v_{ij}$, and by the queries in lines 3–5 of mechanism τ -TRV, we further know that $\max_{i \in N, j \in A} v_{ij} \leq \text{SW}_r(y|\mathbf{v})$. Therefore, we obtain the following upper bound

$$\text{SW}_c^{<\ell^*}(x|\mathbf{v}) \leq (\tau - 1) \cdot \text{SW}_r(y|\mathbf{v}). \quad (9)$$

Finally, we consider the quantity $\text{SW}_c^{\geq\ell^*}(x|\mathbf{v})$. Let S be the set of agents who rank x at position ℓ^* or later, but were never queried their value for x . That is, S contains all agents that contribute to $\text{SW}_c^{\geq\ell^*}(x|\mathbf{v})$. We claim that any agent in S must be queried for an alternative that she prefers more to x .

Claim 1. *If $i \in S$, then during some step, mechanism τ -TRV asks agent i in line 13 for some alternative j such that $j \succ_i x$.*

Proof of claim. Let $i \in S$. By the definition of S , we know that $i \in T_k(x)$ for some $k \in \{\ell^*, \dots, m\}$. If agent i has not been queried twice before step k , then $i \in Q_k$ and therefore it is guaranteed that τ -TRV queries agent i in line 13 during step k . Note that this query is not necessarily about alternative x . During step k it is possible that several other alternatives become eligible as well. If one of them appears at some of the first k positions in the ranking of agent i and is examined before x in line 10, then i will be queried about this alternative instead.

Now suppose that this second query is $\mathcal{V}(i, j)$ and is made during step $s \leq k$. The condition in line 12 guarantees that i ranks j within the first s positions. Towards a contradiction, suppose that $j \not\succeq_i x$. We know that $j \neq x$ by the definition of S . So, it must be $x \succ_i j$. Since i ranks x at position k , this implies that i ranks j at position $k + 1 > s$ or later, leading to a contradiction. \square

From the above claim, we obtain that the contribution of each agent $i \in S$ to the *revealed* welfare of the alternative j that i was queried for must be at least as high as her contribution to the *concealed* welfare of x . Let $S_j \subseteq S$ be the set of agents of S who were queried for alternative j instead of x . The total value of the agents in S_j for j is only a part of the revealed welfare of j . So, we have

$$\begin{aligned} \text{SW}_c^{\geq\ell^*}(x|\mathbf{v}) &\leq \sum_{j \in A \setminus \{x\}} \sum_{i \in S_j} v_{ij} \leq \sum_{j \in A \setminus \{x\}} \text{SW}_r(j|\mathbf{v}) \\ &\leq |\{j \in A \setminus \{x\}\}| \cdot \text{SW}_r(y|\mathbf{v}). \end{aligned}$$

What remains to be bounded is $|\{j \in A \setminus \{x\}\}|$. Clearly the possible alternatives are bounded by $m - 1$. Furthermore, since there are n agents, there can be at most n/τ such alternatives, as an alternative has to be eligible when queried for, and each agent can only be queried for a single alternative at positions $\{2, \dots, m\}$. From this, it follows that

$$\text{SW}_c^{\geq\ell^*}(x|\mathbf{v}) \leq \min\{m - 1, n/\tau\} \cdot \text{SW}_r(y|\mathbf{v}). \quad (10)$$

Combining (7), (8), (9), and (10), we have that

$$\begin{aligned} \text{SW}(x|\mathbf{v}) &\leq \text{SW}_r(y|\mathbf{v}) + (\tau - 1) \cdot \text{SW}_r(y|\mathbf{v}) + \min\{m, n/\tau\} \cdot \text{SW}_r(y|\mathbf{v}) \\ &= \text{SW}_r(y|\mathbf{v}) \cdot (\tau + \min\{m - 1, n/\tau\}) \\ &\leq \text{SW}(y|\mathbf{v}) \cdot (\tau + \min\{m - 1, n/\tau\}), \end{aligned}$$

and the statement follows. \square

It is not hard to see that Theorem 6 essentially translates to distortion $O(\min\{m, \sqrt{n}\})$ with just two queries, by setting $\tau = \min\{m, \sqrt{n}\}$. While this result is not always an improvement over the performance of 1-PRV, it is interesting to interpret it for different ranges of n and m .

Corollary 3. *The mechanism τ -TRV has distortion*

$$\mathcal{D}(\tau\text{-TRV}) \leq \begin{cases} m, & \text{when } n = \omega(m^2), \text{ for } \tau = 1 \\ \sqrt{n}, & \text{when } n = O(m^2), \text{ for } \tau = \sqrt{n} \\ \sqrt{m}, & \text{when } n = \Theta(m), \text{ for } \tau = \sqrt{m} \end{cases}$$

Of course, the cases of Corollary 3 are neither exhaustive nor disjoint, yet they all deserve some attention. The highlight of this section, is the fact that with only two value queries per agent the distortion is $o(m)$, as long as $n = o(m^2)$. In particular, for the special case where $n = \Theta(m)$ we deterministically match the lower bound of $\Omega(\sqrt{m})$ for any randomized ordinal mechanism, and beat the best known randomized ordinal mechanism of [Boutilier et al. \[2015\]](#) which achieves a distortion of $O(\sqrt{m} \cdot \log^* m)$.

However, even for the worst case of the corollary, it is interesting to note how exactly we recover the guarantee of 1-PRV when $n \gg m^2$. Does this mean that the threshold approach we take here essentially wastes the second query per agent as long as n is sufficiently large? Our analysis seems to suggest so as well. In fact, for each range mentioned in the corollary it is not hard to construct instances showing that the analysis of the mechanism is asymptotically tight. Below we provide such a construction for the case where $n \leq m^2$, which further illustrates that the suggested threshold of Corollary 3 is asymptotically optimal, even when one assumes normalized valuation functions.

Theorem 7. *For $n = o(m^2)$, the best threshold for τ -TRV is $\tau = \Theta(\sqrt{n})$, even for unit-sum valuation functions, leading to $\Theta(\sqrt{n})$ distortion.*

Proof. We assume that $\tau \leq n$, since otherwise the distortion of τ -TRV is m , while the distortion of \sqrt{n} -TRV is $\sqrt{n} \leq m$. We first consider the case where $m \leq n$, and define the following two (families of) instances.

First instance:

- Besides alternative x , every other alternative appears *first* in roughly $\frac{n}{m-1}$ rankings (either in $\lfloor \frac{n}{m-1} \rfloor$ or $\lceil \frac{n}{m-1} \rceil$ rankings each);
- There are $r = \min\{m - 1, n/\tau\}$ alternatives $\{y_1, \dots, y_r\}$ who appear *second* in roughly n/r rankings each;
- Alternative x appears third in every ranking;
- The remaining ordinal profile is filled arbitrarily.

The valuation profile is such that the alternatives in the first three positions get a value of $1/3$ for each of their appearances in those positions. Hence, each alternative y_j for $j \in [r]$ has social welfare $\frac{1}{3}(n/m + n/r)$, while alternative x has social welfare $n/3$. Notice that $n/r \geq \tau$; indeed, when $r = n/\tau$ this is trivial, and when $r = m-1$ we have $m-1 < n/\tau \Rightarrow \tau < n/(m-1) = n/r$. So, by definition, τ -TRV will query all agents at the second position and will therefore select some alternative y_j , yielding distortion $\mathcal{D} = \frac{mr}{m+r}$. Unless $\tau = \omega(n/m)$, there is an infinite subsequence of pairs of values for n and m for which $\frac{mr}{m+r} = \Theta(m) = \omega(\sqrt{n})$. This would imply a worst case distortion $\mathcal{D} = \omega(\sqrt{n})$ in general. Since \sqrt{n} -TRV already has a strictly better distortion guarantee than that, and we are looking for an optimal threshold, we may assume that $\tau = \omega(n/m)$. Therefore, subject to the possible optimal choices of τ , this first instance forces τ -TRV to have distortion $\mathcal{D} \geq \frac{m\frac{n}{\tau}}{m+\frac{n}{\tau}} = \Theta(n/\tau)$.

Second instance:

- Besides alternatives x , y and z , every other alternative appears *first* in roughly $\frac{n}{m-3}$ rankings;
- Alternative x appears *second* in $\tau - 1$ rankings;
- Alternative y appears *second* in $n - \tau + 1$ rankings;
- Alternative z appears *third* in every ranking;
- The remaining ordinal profile is filled arbitrarily.

Let $\mu = \max\{n, 10m\}$. The alternatives at the first position get a value of $\frac{1}{2}$ in the rankings where x is second, and $1 - \frac{m-1}{\mu}$ otherwise. Alternative x gets a value of $\frac{1}{2} - \frac{m-2}{\mu}$ for every appearances at a second position. The alternatives in the remaining positions equally share the remaining value (out of total value 1 per agent). That is, they all get a value of $\frac{1}{\mu}$. By definition, τ -TRV will query for y at the second position (if $n - \tau + 1 \geq \tau$) or for z at the third position (if $n - \tau + 1 < \tau$). At this point, the revealed welfare of x is 0, while the revealed welfare of any other alternative is positive, and so the mechanism will elect some alternative $w \neq x$. Observe that the social welfare of w is no more than $1 + \frac{n-1}{\mu} \leq 2$ in any case, while the social welfare of x is at least $(\frac{1}{2} - \frac{m-2}{\mu})(\tau - 1) \geq \frac{2}{5}(\tau - 1)$. Therefore, this second instance forces τ -TRV to have distortion $\mathcal{D} \geq \frac{1}{5}(\tau - 1) = \Theta(\tau)$.

To balance the distortion between these two instances, we must set $\tau \approx \sqrt{5n} = \Theta(\sqrt{n})$.

For the case $m > n$, only n alternatives are relevant. Therefore, we can construct the above instances for $m' = n$ and then add $m - n$ dummy alternatives at the end of each ranking such that all agents have zero value for them. Again, we get that the best threshold for τ -TRV is $\tau = \Theta(\sqrt{n})$. \square

6 Constant Distortion Using Only Comparison Queries

A straightforward, yet crucial, observation is that mechanism k -ARV can actually be implemented using just *one* value query. We can ask the value of each agent for her favorite alternative, and then ask $O(k \log m)$ comparison queries that guide the binary search in computing the maximal acceptable sets. Hence, $\log m$ -ARV achieves constant distortion using only one value query and $O(\log^2 m)$ comparison queries.

Therefore, it is natural to ask whether we can avoid this single value query entirely, and rely only on comparison queries instead. Surprisingly, for unit-sum valuation functions, we show that this is indeed possible at no extra cost! More precisely, we show that we can approximate the value that an agent has for her favorite alternative within a factor of $1 \pm \varepsilon$, using $O(\log^2 m)$ comparison queries. Note that this is the first time that we assume the unit-sum normalization.

For the sake of readability, we focus on a single agent and write u_j for her value for the alternative that she ranks at position $j \in [m]$. We take the same approach as in the proof of Theorem 4 in order to build an approximate valuation profile. Since everything in this profile is expressed in terms of the largest value u_1 , we utilize the unit-sum assumption to approximately solve for u_1 .

Theorem 8. *For any constant $\varepsilon \geq 1/m$, it is possible to compute some u^* such that $(1 - \varepsilon)u^* \leq u_1 \leq (1 + \varepsilon)u^*$, using $O(\log^2 m)$ comparison queries.*

Proof. Let $\kappa = \lceil \log_{1+\varepsilon} m^2 \rceil = \Theta(\log m)$. We define κ thresholds $\lambda_\ell = (1 + \varepsilon)^\ell$ for $\ell \in [\kappa]$; observe that $\lambda_i = \lambda_{i-1} \cdot \lambda_1$. For each $\ell \in [\kappa]$, we perform a binary search using $\Theta(\log m)$ comparison queries to find the maximum integer ξ_ℓ such that $u_{\xi_\ell} \geq \frac{u_1}{\lambda_\ell}$; we also set $\xi_0 = 1$ and $\xi_{\kappa+1} = m$. Hence, we have that

$$u_j \in \begin{cases} \left[\frac{u_1}{\lambda_\ell}, \frac{u_1}{\lambda_{\ell-1}} \right), & \text{for each } j \in (\xi_{\ell-1}, \xi_\ell], \ell \in [\kappa] \\ \left[0, \frac{u_1}{\lambda_\kappa} \right), & \text{for each } j \in (\xi_\kappa, \xi_{\kappa+1}]. \end{cases} \quad (11)$$

To simplify the notation, let $g_i = \xi_i - \xi_{i-1}$ for $i \in [\kappa + 1]$. Note that $g_i \leq m$.

By the unit-sum normalization we have

$$1 = \sum_{j=1}^m u_j = \sum_{\ell=1}^{\kappa+1} \sum_{j \in (\xi_{\ell-1}, \xi_\ell]} u_j.$$

Using (11), we can now upper- and lower-bound the above expression. We start with the upper bound:

$$1 = \sum_{\ell=1}^{\kappa+1} \sum_{j \in (\xi_{\ell-1}, \xi_\ell]} u_j \leq \sum_{\ell=1}^{\kappa+1} \sum_{j \in (\xi_{\ell-1}, \xi_\ell]} \frac{u_1}{\lambda_{\ell-1}} = \sum_{\ell=1}^{\kappa+1} u_1 \frac{g_\ell}{\lambda_{\ell-1}} = u_1 \sum_{\ell=1}^{\kappa} \frac{g_\ell}{\lambda_{\ell-1}} + u_1 \frac{g_{\kappa+1}}{\lambda_\kappa}.$$

By the definition of λ_κ we have that $\frac{g_{\kappa+1}}{\lambda_\kappa} \leq \frac{1}{m} \leq \varepsilon$ and hence,

$$u_1 \geq (1 - \varepsilon) \cdot \left(\sum_{\ell=1}^{\kappa} \frac{g_\ell}{\lambda_{\ell-1}} \right)^{-1}.$$

Similarly, by using the lower bounds in (11), we have that

$$1 = \sum_{\ell=1}^{\kappa+1} \sum_{j \in (\xi_{\ell-1}, \xi_\ell]} u_j \geq \sum_{\ell=1}^{\kappa} \sum_{j \in (\xi_{\ell-1}, \xi_\ell]} \frac{u_1}{\lambda_\ell} = \sum_{\ell=1}^{\kappa} u_1 \frac{g_\ell}{\lambda_\ell} = u_1 \sum_{\ell=1}^{\kappa} \frac{g_\ell}{\lambda_{\ell-1}} \cdot \frac{1}{\lambda_1} = \frac{u_1}{1 + \varepsilon} \sum_{\ell=1}^{\kappa} \frac{g_\ell}{\lambda_{\ell-1}}.$$

or, equivalently,

$$u_1 \leq (1 + \varepsilon) \cdot \left(\sum_{\ell=1}^{\kappa} \frac{g_\ell}{\lambda_{\ell-1}} \right)^{-1}.$$

Hence, the theorem follows by setting $u^* = \left(\sum_{\ell=1}^{\kappa} \frac{g_\ell}{\lambda_{\ell-1}} \right)^{-1}$. Indeed, to compute u^* only the integers ξ_ℓ , $\ell \in [\kappa]$ are used. Those, in turn, are computed via $\Theta(\log m)$ binary searches. Hence, we only need $O(\log^2 m)$ comparison queries. \square

By inspecting the proof of Theorem 4, it is easy to see that knowing the approximate valuation profile \tilde{v} exactly or perturbed within a multiplicative constant factor, makes no difference asymptotically. Therefore, we augment k -ARV with a pre-processing step where each maximum value v_i^* is approximated according to Theorem 8 above, and these approximations are used in line 2 of the mechanism. For $k = \log m$, this new mechanism, which we call *modified* $(\log m)$ -ARV, achieves the same distortion guarantee and asks the same number of queries (asymptotically) with $(\log m)$ -ARV.

Corollary 4. *Modified* $(\log m)$ -ARV achieves distortion $O(1)$ using $O(\log^2 m)$ comparison queries per agent.

7 Lower Bounds

We now present general lower bounds on the distortion which depend on the number of value queries the mechanisms are allowed to ask per agent, but are unconditional on how and where they decide to ask these queries. In particular, we show that the distortion of any mechanism that makes one value query per agent is $\Omega(m)$ when the agents have unrestricted valuation functions, and $\Omega(\sqrt{m})$ when the agents have unit-sum valuation functions. Moreover, for mechanisms that are allowed to make λ queries per agents, we show a weaker lower bound of $\Omega\left(\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}\right)$ for unrestricted valuation functions. This shows that in order to achieve constant distortion, we need to necessarily make more than $O\left(\frac{\log m}{\log \log m}\right)$ queries per agent. Closing the gap between this lower bound and the upper bound of $O(\log^2 m)$ queries that the mechanism $O(\log m)$ -ARV from Section 4 requires in order to achieve constant distortion is one of the most interesting open problems that our work leaves open; see the discussion in Section 8.

Before we proceed with the presentation of the results of this section, let us give a very brief roadmap of the proofs. The high-level idea is similar to the that used in the proofs of lower bounds presented in previous sections (like Theorems 2 and 3 in Section 3), but the particular constructions and arguments exploited in the proofs below are more delicate; this is an aftermath of the fact that we aim to lower-bound the distortion of any mechanism. To this end, assuming an arbitrary mechanism (that is allowed to make a specific number of queries per agent), we first define a single ordinal preference profile which is given as input to the mechanism, and also carefully define the cardinal information that is revealed from all possible queries of the mechanism. This cardinal information is such that it is always possible to complete the valuation profile in a way that allows us to define the social welfare of the optimal alternative to be much higher than that of the alternative selected by the mechanism. Since we do not know how the mechanism makes its selection, we need to take into account every possible scenario, and therefore define many different valuation profiles that can be used in different cases.

To simplify our discussion when we deal with unrestricted valuation functions in this section, we assume that the values are normalized and lie in the interval $[0, 1]$. This is without loss of generality since we make no other assumptions on the way the mechanism behave, other than that they are allowed to ask a particular number of queries.

7.1 One-Query Mechanisms with Unrestricted Valuations

We start by showing that, for unrestricted valuations, any mechanism that makes one value query per agent has linear distortion. This also shows that the mechanism 1-PRV from Section 3 is the best possible mechanism among such mechanisms.

Theorem 9. *For unrestricted valuation functions, the distortion of any mechanism that uses one value query per agent is $\Omega(m)$.*

Proof. Let \mathcal{M} be an arbitrary mechanism that makes one value query per agent, and consider an instance with $m \geq 4$ alternatives and $n = m - 2$ agents, where m is an even number. We denote the set of alternatives as $A = \{a_1, \dots, a_{m-2}, x, y\}$. Using the notation $[z, w]$ to denote the fact that alternatives z and w are ordered arbitrarily in the ranking of an agent, we define the ordinal profile as follows:

- The ranking of agent $i \leq \frac{n}{2}$ is $a_i \succ_i x \succ_i y \succ_i [a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_{m-2}]$;
- The ranking of agent $i > \frac{n}{2}$ is $a_i \succ_i y \succ_i x \succ_i [a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_{m-2}]$.

Depending on the positions at which \mathcal{M} queries, we reveal the following cardinal information:

- For every query at a first position we reveal a value of m^{-1} ;
- For every query at a second or third position we reveal a value of m^{-2} ;
- For any other position we reveal a value of 0.

We claim that \mathcal{M} must query all agents at the first position, as otherwise its distortion is $\Omega(m)$. Assume otherwise that \mathcal{M} does not query agent 1 her value for alternative a_1 ; this is without loss of generality due to symmetry. We now define two valuation profiles \mathbf{v}_1 and \mathbf{v}_2 , which are both consistent to the ordinal profile and the revealed information, but differ on the value that agent 1 has for alternative a_1 . In particular:

- In both \mathbf{v}_1 and \mathbf{v}_2 , every agent $i \geq 2$ has value m^{-1} for alternative a_i , m^{-2} for alternatives x and y , and 0 for everyone else;
- In both \mathbf{v}_1 and \mathbf{v}_2 , agent 1 has value m^{-2} for alternatives x and y , and 0 for every alternative a_i for $i \geq 2$. The value of agent 1 for alternative a_1 is m^{-2} in \mathbf{v}_1 , and 1 in \mathbf{v}_2 .

These two profiles are utilized in the following way: If \mathcal{M} selects a_1 , then the valuation profile is set to be \mathbf{v}_1 , while if \mathcal{M} selects some other alternative, then the valuation profile is set to be \mathbf{v}_2 . Now, observe that

$$\text{SW}(a_i|\mathbf{v}_1) = \text{SW}(a_i|\mathbf{v}_2) = m^{-1} \quad \text{for every } i \geq 2,$$

and

$$\text{SW}(x|\mathbf{v}_1) = \text{SW}(x|\mathbf{v}_2) = \text{SW}(y|\mathbf{v}_1) = \text{SW}(y|\mathbf{v}_2) = (m - 2) \cdot m^{-2} \leq m^{-1}.$$

If \mathcal{M} selects a_1 , the social welfare of a_1 is $\text{SW}(a_1|\mathbf{v}_1) = m^{-2}$ and therefore any alternative a_i for $i \geq 2$ is optimal, yielding distortion equal to m . Similarly, when \mathcal{M} selects some alternative different than a_1 , then a_1 is optimal with social welfare $\text{SW}(a_1|\mathbf{v}_2) = 1$, yielding distortion at least m .

Hence, \mathcal{M} must query all agents at the first position in order to learn a value of m^{-1} for every alternative a_i , $i \in [n]$. We now define three valuation profiles \mathbf{v}_3 , \mathbf{v}_4 and \mathbf{v}_5 , which are consistent to the ordinal profile and this revealed information, but differ on the values that the agents have for alternatives x and y ; in particular, \mathbf{v}_4 and \mathbf{v}_5 are symmetric.

- In all three profiles, every agent $i \in [n]$ has value m^{-1} for alternative a_i , and 0 for any alternative a_j such that $j \neq i$;
- In \mathbf{v}_3 , all agents have value m^{-1} for alternatives x and y ;
- In \mathbf{v}_4 , all agents have value m^{-2} for alternative y , every agent $i > n/2$ (who ranks x after y) has value m^{-2} for x , and every agent $i \leq n/2$ (who ranks x before y) has value m^{-1} for x .

- In \mathbf{v}_5 , all agents have value m^{-2} for alternative x , every agent $i \leq n/2$ (who ranks y after x) has value m^{-2} for y , and every agent $i > n/2$ (who ranks y before x) has value m^{-1} for y .

If \mathcal{M} selects some alternative a_i for $i \in [n]$, then the valuation profile is set to be \mathbf{v}_3 , while if \mathcal{M} selects alternative y or x , then the valuation profile is set to be \mathbf{v}_4 or \mathbf{v}_5 , respectively. Given this, observe that if \mathcal{M} decides to select alternative a_i for some $i \in [n]$, then since

$$\text{SW}(a_i|\mathbf{v}_3) = \text{SW}(a_i|\mathbf{v}_4) = \text{SW}(a_i|\mathbf{v}_5) = m^{-1} \quad \text{for every } i \in [n],$$

and

$$\text{SW}(x|\mathbf{v}_3) = \text{SW}(y|\mathbf{v}_3) = (m-2) \cdot m^{-1} = 1 - 2m^{-1},$$

the distortion is at least $m-2$. Similarly, if \mathcal{M} decides to select alternative y , then since

$$\text{SW}(y|\mathbf{v}_4) = (m-2)m^{-2} \leq m^{-1}$$

and

$$\text{SW}(x|\mathbf{v}_4) = \left(\frac{m}{2} - 1\right) m^{-1} + \left(\frac{m}{2} - 1\right) m^{-2} = \frac{1}{2} \left(1 - m^{-1} - 2m^{-2}\right),$$

the distortion is at least $\frac{1}{2}(m-1-2m^{-1}) \geq \frac{m}{4}$ for any $m \geq 4$; the case where \mathcal{M} selects x is symmetric and follows by \mathbf{v}_5 . In any case, \mathcal{M} has distortion $\Omega(m)$ and the theorem follows. \square

7.2 General Mechanisms with Unrestricted Valuations

We will now focus on mechanisms that make a number $\lambda \geq 1$ of queries per agent, and will show a weaker lower bound on the distortion which depends on λ .

Theorem 10. *For unrestricted valuation functions, the distortion of any mechanism that uses $\lambda \geq 1$ value queries per agent is $\Omega\left(\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}\right)$.*

Proof. Our instance consists of $m \geq \lambda$ alternatives and $n = m$ agents. We partition the set A of alternatives into the following $(\lambda+2)$ sets:

- A_j with $|A_j| = m^{1-\frac{j}{\lambda+1}}$, for every $j \in [\lambda]$;
- X with $|X| = 2$;
- Y with $|Y| = m - 2 - \sum_{j=1}^{\lambda} m^{1-\frac{j}{\lambda+1}}$.

We will now define the ordinal profile. Let $X = \{x_1, x_2\}$ and denote by $[z, w]$ the fact that alternatives z and w are ordered arbitrarily in the ranking of an agent. For every agent i there exists an alternative $a_{ij} \in A_j$ for each $j \in [\lambda]$ such that:

- The ranking of agent $i \leq \frac{m}{2}$ is $a_{i1} \succ_i \dots \succ_i a_{i\lambda} \succ_i x_1 \succ_i x_2 \succ_i [Y] \succ_i [\cup_{j \in [\lambda]} A_j \setminus \{a_{ij}\}]$;
- The ranking of agent $i > \frac{m}{2}$ is $a_{i1} \succ_i \dots \succ_i a_{i\lambda} \succ_i x_2 \succ_i x_1 \succ_i [Y] \succ_i [\cup_{j \in [\lambda]} A_j \setminus \{a_{ij}\}]$;

In words, every agent i ranks some alternative $a_{ij} \in A_j$ at position $j \in [\lambda]$, followed by the two alternatives of $X = \{x_1, x_2\}$ at positions $(\lambda + 1)$ and $(\lambda + 2)$, followed by all alternatives of Y (in an arbitrary order), followed by the alternatives of $\cup_{j \in [\lambda]} A_j \setminus \{a_{ij}\}$ (in an arbitrary order). Observe that the alternatives of Y are all dominated by the alternatives of X in the sense that both x_1 and x_2 have at least as much social welfare as any alternative of Y . The choices as to how the alternatives of $\cup_{j \in [\lambda]} A_j$ are distributed in the rankings of the agents are such that:

- Each alternative of A_j appears $m^{\frac{j}{\lambda+1}}$ times at position $j \in [\lambda]$;
- For any $j \in [\lambda - 1]$ and pair of agents i, i' such that $a_{ij} = a_{i'j}$, it holds that $a_{i,j+1} = a_{i',j+1}$.

Hence, the agents with the same favorite alternative have exactly the same ranking. To simplify our discussion in what follows, we refer to the alternatives in Y and $\cup_{j \in [\lambda]} A_j \setminus \{a_{ij}\}$ as the *tail* alternatives of agent i . Let $T_j(z)$ be the set of the $m^{\frac{j}{\lambda+1}}$ agents that rank alternative $z \in A_j$ at position $j \in [\lambda]$. Figure 2 depicts the ordinal profile of our instance for $\lambda = 2$.

Let \mathcal{M} be an arbitrary mechanism that makes λ value queries per agent. Naturally, we assume that \mathcal{M} does not elect any dominated alternative from Y . The cardinal information that is revealed due to the queries of \mathcal{M} is as follows:

- Each of the first $\frac{\lambda}{\lambda+1} \cdot m^{\frac{j}{\lambda+1}}$ queries of \mathcal{M} for any alternative $z \in A_j$, $j \in [\lambda]$ reveals a value of $m^{-\frac{j+1/2}{\lambda+1}}$, while each of the remaining $\frac{1}{\lambda+1} \cdot m^{\frac{j}{\lambda+1}}$ queries of \mathcal{M} for z reveals a value of $m^{-\frac{j}{\lambda+1}}$.
- Every query of \mathcal{M} for an alternative of X reveals a value of m^{-1} .
- Every query of \mathcal{M} for a tail alternative reveals zero value.

To simplify our notation in the rest of the proof, let δ_{ij} be the indicator variable:

$$\delta_{ij} = \begin{cases} 1, & \text{if } \mathcal{M} \text{ asks agent } i \text{ for } a_{ij} \text{ and has previously asked} \\ & \text{strictly less than } \frac{\lambda}{\lambda+1} \cdot m^{\frac{j}{\lambda+1}} \text{ other agents of } T_j(a_{ij}) \text{ for } a_{ij} \\ 0, & \text{otherwise.} \end{cases}$$

Now, assume towards a contradiction that \mathcal{M} has distortion $\mathcal{D}(\mathcal{M}) \notin \Omega\left(\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}\right)$. Using the next two claims, we will show by induction that \mathcal{M} must query a large proportion of the agents at the first λ positions, since otherwise the distortion of \mathcal{M} would be $\Omega\left(\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}\right)$.

Claim 2. *The mechanism \mathcal{M} must ask at the first position strictly more than $\frac{\lambda}{\lambda+1} \cdot m^{\frac{1}{\lambda+1}}$ of the agents in $T_1(z)$ for every alternative $z \in A_1$.*

Claim 3. *Given that for every alternative $z \in A_j$, $j \in [\lambda - 1]$ the mechanism \mathcal{M} asks at the first j positions strictly more than $\left(1 - \frac{j}{\lambda+1}\right) \cdot m^{\frac{j}{\lambda+1}}$ of the agents in $T_j(z)$, \mathcal{M} must ask at the first $j + 1$ positions strictly more than $\left(1 - \frac{j+1}{\lambda+1}\right) \cdot m^{\frac{j+1}{\lambda+1}}$ of the agents in $T_{j+1}(w)$ for every alternative $w \in A_{j+1}$.*

By Claims 2 and 3, \mathcal{M} must ask at the first λ positions strictly more than $\left(1 - \frac{\lambda}{\lambda+1}\right) \cdot m^{\frac{\lambda}{\lambda+1}} = \frac{1}{\lambda+1} \cdot m^{\frac{\lambda}{\lambda+1}}$ of the agents in $T_\lambda(z)$ for every alternative $z \in A_\lambda$. Consequently, since A_λ consists of exactly $m^{\frac{1}{\lambda+1}}$ alternatives, there are at least $m^{\frac{1}{\lambda+1}} \cdot \frac{1}{\lambda+1} \cdot m^{\frac{\lambda}{\lambda+1}} = \frac{1}{\lambda+1} \cdot m$ agents that are *not* queried at positions $(\lambda + 1)$ and $(\lambda + 2)$ for

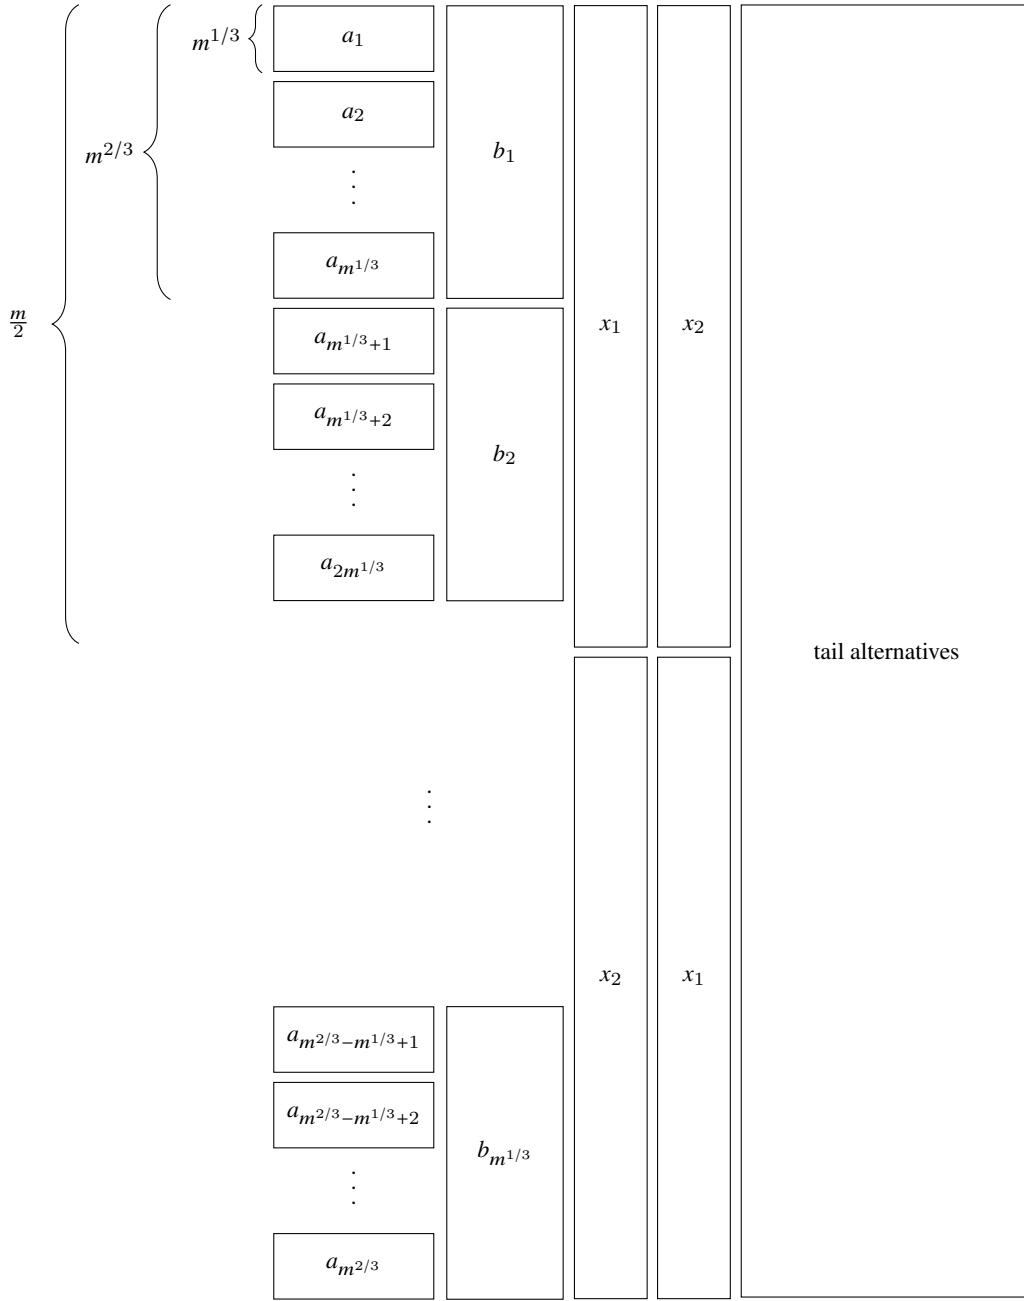


Figure 2: An example of the instance used in the proof of Theorem 10 for $\lambda = 2$; for convenience, we denote here the alternatives of A_1 and A_2 as $A_1 = \{a_1, \dots, a_{m^{2/3}}\}$ and $A_2 = \{b_1, \dots, b_{m^{1/3}}\}$.

the alternatives of $X = \{x_1, x_2\}$. Let S be the set of these $\frac{1}{\lambda+1} \cdot m$ agents; observe that half of them rank x_1 ahead of x_2 and half of them rank x_1 below x_2 , which follows by the fact that S includes the same number of agents per alternative of A_λ and the definition of the ordinal profile. Further, we define two more sets of agents: $S_{12} = \{i \in S : i \leq \frac{m}{2}\}$ and $S_{21} = S \setminus S_{12}$. Observe that all agents of S_{12} rank alternative x_1 ahead of x_2 , and all agents of S_{21} rank x_2 ahead of x_1 .

Now, we define three valuation profiles \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 , which are consistent to the ordinal profile and the cardinal information revealed by the queries of \mathcal{M} , but differ on the values that the agents in S have for the alternatives in X ; in particular, \mathbf{v}_2 and \mathbf{v}_3 are symmetric.

- In all three profiles, every agent $i \in [m]$ has value $\delta_{ij} \cdot m^{-\frac{j+1/2}{\lambda+1}} + (1 - \delta_{ij}) \cdot m^{-\frac{j}{\lambda+1}}$ for the alternative $a_{ij} \in A_j$ that she ranks at position $j \in [\lambda]$, and zero value for her tail alternatives;
- In all three profiles, every agent $i \notin S$ has value m^{-1} for both x_1 and x_2 ;
- In \mathbf{v}_1 , every agent $i \in S$ has value $m^{-\frac{\lambda+1/2}{\lambda+1}}$ for both x_1 and x_2 ;
- In \mathbf{v}_2 , every agent $i \in S_{12}$ has value m^{-1} for both x_1 and x_2 , while every agent $i \in S_{21}$ has value m^{-1} for x_1 and value $m^{-\frac{\lambda+1/2}{\lambda+1}}$ for x_2 .
- In \mathbf{v}_3 , every agent $i \in S_{21}$ has value m^{-1} for both x_1 and x_2 , while every agent $i \in S_{12}$ has value m^{-1} for x_2 and value $m^{-\frac{\lambda+1/2}{\lambda+1}}$ for x_1 .

Next, we compute the social welfare of each alternative for the different valuation profiles:

- The social welfare of every alternative $z \in Y$ is

$$\text{SW}(z|\mathbf{v}_1) = \text{SW}(z|\mathbf{v}_2) = \text{SW}(z|\mathbf{v}_3) = 0.$$

- The social welfare of every alternative $z \in \cup_{j \in [\lambda]} A_j$ is

$$\begin{aligned} \text{SW}(z|\mathbf{v}_1) = \text{SW}(z|\mathbf{v}_2) = \text{SW}(z|\mathbf{v}_3) &= \frac{\lambda}{\lambda+1} \cdot m^{\frac{j}{\lambda+1}} \cdot m^{-\frac{j+1/2}{\lambda+1}} + \frac{1}{\lambda+1} \cdot m^{\frac{j}{\lambda+1}} \cdot m^{-\frac{j}{\lambda+1}} \\ &= \frac{\lambda}{\lambda+1} \cdot m^{-\frac{1}{2(\lambda+1)}} + \frac{1}{\lambda+1} \leq 1. \end{aligned}$$

- The social welfare of x_1 is

$$\begin{aligned} \text{SW}(x_1|\mathbf{v}_1) &\geq \frac{1}{\lambda+1} \cdot m \cdot m^{-\frac{\lambda+1/2}{\lambda+1}} = \frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}, \\ \text{SW}(x_1|\mathbf{v}_2) &= m \cdot m^{-1} = 1, \\ \text{SW}(x_1|\mathbf{v}_3) &\geq \frac{1}{2} \cdot \frac{1}{\lambda+1} \cdot m \cdot m^{-\frac{\lambda+1/2}{\lambda+1}} = \frac{1}{2(\lambda+1)} \cdot m^{\frac{1}{2(\lambda+1)}}. \end{aligned}$$

- The social welfare of x_2 is

$$\begin{aligned} \text{SW}(x_2|\mathbf{v}_1) &\geq \frac{1}{\lambda+1} \cdot m \cdot m^{-\frac{\lambda+1/2}{\lambda+1}} = \frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}, \\ \text{SW}(x_2|\mathbf{v}_2) &\geq \frac{1}{2} \cdot \frac{1}{\lambda+1} \cdot m \cdot m^{-\frac{\lambda+1/2}{\lambda+1}} = \frac{1}{2(\lambda+1)} \cdot m^{\frac{1}{2(\lambda+1)}}, \\ \text{SW}(x_2|\mathbf{v}_3) &= m \cdot m^{-1} = 1. \end{aligned}$$

Depending on the choices of the mechanism \mathcal{M} , we set the valuation profile to be one of \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 so that the distortion is as high as possible. In particular, we have:

- If \mathcal{M} selects any alternative $z \in \cup_{j \in [\lambda]} A_j$, we set the valuation profile to be \mathbf{v}_1 . Hence, the social welfare of the winner z is at most 1, while any alternative of X is optimal with social welfare at least $\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}$, yielding distortion at least $\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}$.
- If \mathcal{M} selects alternative x_1 , we set the valuation profile to be \mathbf{v}_2 . Hence, the social welfare of the winner x_1 is exactly 1, while x_2 is the optimal alternative with social welfare at least $\frac{1}{2(\lambda+1)} \cdot m^{\frac{1}{2(\lambda+1)}}$, yielding distortion at least $\frac{1}{2(\lambda+1)} \cdot m^{\frac{1}{2(\lambda+1)}}$.
- If \mathcal{M} selects alternative x_2 , we set the valuation profile to be \mathbf{v}_3 , which is symmetric to the previous case and again yields distortion at least $\frac{1}{2(\lambda+1)} \cdot m^{\frac{1}{2(\lambda+1)}}$.

Therefore, the distortion of \mathcal{M} is $\Omega\left(\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}\right)$ and the proof of the theorem is now complete; the proofs of Claims 2 and 3 can be found in the appendix. \square

Using Theorem 10, we can show several lower bounds on the distortion of any mechanism, depending on the number of queries that it makes per agent. In particular, we have the following statement.

Corollary 5. *For unrestricted valuation functions, the distortion of any mechanism \mathcal{M} that uses λ queries per agent is*

$$\mathcal{D}(\mathcal{M}) = \begin{cases} \Omega\left(m^{\frac{1}{2(\lambda+1)}}\right), & \text{for any constant } \lambda \geq 1 \\ \Omega(\log \log m), & \text{for } \lambda = O\left(\frac{\log m}{\log \log m}\right). \end{cases}$$

7.3 One-Query Mechanisms with Unit-Sum Valuations

Next, we turn our attention to unit-sum valuation functions and mechanisms that are allowed to make only one value query per agent. For this case, we are able to show a weaker lower bound of $\Omega(\sqrt{m})$, which indicates (but does not prove) some separation between unrestricted and unit-sum valuation functions.

Theorem 11. *For unit-sum valuation functions, the distortion of any mechanism that uses only one direct value query per agent is $\Omega(\sqrt{m})$.*

Proof. Consider an instance with $m \geq 4$ alternatives and $n = \sqrt{m}$ agents. We partition the set A of alternatives into the following four sets:

- $B = \{b_1, \dots, b_{\sqrt{m}}\}$ with $|B| = \sqrt{m}$;
- $C = \{c_1, c_2\}$ with $|C| = 2$;
- $D = \{d_1, \dots, d_{\sqrt{m}-3}\}$ with $|D| = \sqrt{m} - 3$;
- $E = \{e_1, \dots, e_{m-2\sqrt{m}+1}\}$ with $|E| = m - 2\sqrt{m} + 1$.

Using the notation $[z, w]$ to denote the fact that alternatives z and w are ordered arbitrarily in the ranking of an agent, we define the following ordinal profile:

- The ranking of agent $i \leq \frac{n}{2}$ is $b_i \succ_i c_1 \succ_i c_2 \succ_i [D] \succ_i [E] \succ_i [B \setminus \{b_i\}]$
- The ranking of agent $i > \frac{n}{2}$ is $b_i \succ_i c_2 \succ_i c_1 \succ_i [D] \succ_i [E] \succ_i [B \setminus \{b_i\}]$

Observe that the alternatives of $D \cup E$ are all dominated by alternatives c_1 and c_2 in the sense that electing c_1 or c_2 always yields social welfare that is at least as much as the social welfare of any alternative in $D \cup E$. We refer to the alternatives of $E \cup B \setminus \{b_i\}$ as the tail alternatives of agent i .

Let \mathcal{M} be any mechanism that makes one query per agent. Naturally, we assume that \mathcal{M} does not elect any dominated alternative from $D \cup E$. The possible queries of \mathcal{M} reveal the following cardinal information:

- A query for the favorite alternative of an agent (at the first position) reveals a value of $\frac{1}{\sqrt{m}}$;
- A query for an alternative in $C \cup D$ reveals a value of $\frac{1}{m}$;
- Any other query (for a tail alternative) reveals zero value.

We define the following sets of agents, depending on the function of \mathcal{M} :

- S_1 is the set of agents queried at the first position (for their favorite alternative);
- S_C is the set of agents queried for some alternative in C ;
- S_D is the set of agents queried for some alternative in D ;
- $S_{>}$ is the set of agents queried for some tail alternative.

Next, we distinguish between three cases, depending on the alternative that \mathcal{M} elects.

Case I: \mathcal{M} selects alternative c_1 (the case of c_2 is symmetric)

If $|S_1| < n$, we define the following valuation profile \mathbf{v} :

- For every agent $i \in S_C$, we set $v_{i,b_i} = 1 - \frac{2}{m}$ and $v_{i,c_1} = v_{i,c_2} = \frac{1}{m}$; the value for all other alternatives is zero.
- For every agent $i \in S_D$, we set $v_{i,b_i} = 1 - \frac{\sqrt{m}-1}{m}$ and $v_{i,c_1} = v_{i,c_2} = v_{i,d_j} = \frac{1}{m}$ for $j \in [\sqrt{m} - 3]$; the value for all other alternatives is zero.
- For every agent $i \in S_{>}$, we set $v_{i,b_i} = 1$; the value for all other alternatives is zero.
- For every agent $i \in S_1$, we set $v_{i,b_i} = \frac{1}{\sqrt{m}}$, and split the remaining value of $1 - \frac{1}{\sqrt{m}}$ equally among all other $m - 1$ alternatives so that for each of them the value of agent i is $\frac{\sqrt{m}-1}{\sqrt{m}(m-1)}$.

Hence, alternative c_1 has social welfare

$$\begin{aligned}
\text{SW}(c_1|\mathbf{v}) &= \left(|S_C| + |S_D| \right) \cdot \frac{1}{m} + |S_{>}| \cdot 0 + |S_1| \cdot \frac{\sqrt{m}-1}{\sqrt{m}(m-1)} \\
&\leq \left(|S_1| + |S_C| + |S_D| \right) \cdot \frac{1}{m} \\
&\leq \frac{1}{\sqrt{m}},
\end{aligned}$$

where the first inequality follows since $\frac{\sqrt{m}-1}{\sqrt{m(m-1)}} \leq \frac{1}{m} \Leftrightarrow \sqrt{m} \leq m$, and the second follows by the fact that $|S_1| + |S_C| + |S_D| \leq n = \sqrt{m}$. Since $|S_1| < n$, there exists an agent $i^* \in S_C \cup S_D \cup S_>$ such that her favorite alternative b_{i^*} has social welfare

$$\text{SW}(b_{i^*}|\mathbf{v}) \geq 1 - \frac{\sqrt{m}-1}{m} \geq 1 - \frac{1}{\sqrt{m}}.$$

As a result, the distortion is at least $\sqrt{m} - 1$.

If $|S_1| = n$, we define the following valuation profile \mathbf{v} :

- For every agent $i \leq \frac{n}{2}$, we set $v_{i,b_i} = \frac{1}{\sqrt{m}}$, and split the remaining value of $1 - \frac{1}{\sqrt{m}}$ equally among all other $m - 1$ alternatives so that for each of them the value of agent i is $\frac{\sqrt{m}-1}{\sqrt{m(m-1)}}$.
- For every agent $i > \frac{n}{2}$, we set $v_{i,b_i} = v_{i,c_2} = \frac{1}{\sqrt{m}}$, and split the remaining value of $1 - \frac{2}{\sqrt{m}}$ equally among all other $m - 2$ alternatives so that for each of them the value of agent i is $\frac{\sqrt{m}-2}{\sqrt{m(m-2)}}$.

Hence, alternative c_1 has social welfare

$$\begin{aligned} \text{SW}(c_1|\mathbf{v}) &= \frac{n}{2} \cdot \left(\frac{\sqrt{m}-1}{\sqrt{m(m-1)}} + \frac{\sqrt{m}-2}{\sqrt{m(m-2)}} \right) \\ &\leq \frac{n}{2} \cdot \frac{2\sqrt{m}-3}{\sqrt{m(m-2)}}. \end{aligned}$$

On the other hand, alternative c_2 has social welfare

$$\begin{aligned} \text{SW}(c_2|\mathbf{v}) &= \frac{n}{2} \cdot \left(\frac{\sqrt{m}-1}{\sqrt{m(m-1)}} + \frac{1}{\sqrt{m}} \right) \\ &= \frac{n}{2} \cdot \frac{m + \sqrt{m} - 2}{\sqrt{m(m-1)}}. \end{aligned}$$

Consequently, the distortion is at least

$$\frac{\text{SW}(c_2|\mathbf{v})}{\text{SW}(c_1|\mathbf{v})} \geq \frac{m-2}{m-1} \cdot \frac{m + \sqrt{m} - 2}{2\sqrt{m}-3} \geq \frac{1}{2}\sqrt{m},$$

where the last inequality holds for any $m \geq 3$.

Case II: $|S_1| \geq 1$ and \mathcal{M} selects some alternative b_{i^*} for $i^* \in S_1$.

If $|S_1| < n$, we define the following valuation profile \mathbf{v} :

- For every agent $i \in S_C$, we set $v_{i,b_i} = 1 - \frac{2}{m}$ and $v_{i,c_1} = v_{i,c_2} = \frac{1}{m}$; the value for all other alternatives is zero.
- For every agent $i \in S_D$, we set $v_{i,b_i} = 1 - \frac{1}{\sqrt{m}} + \frac{1}{m}$ and $v_{i,c_1} = v_{i,c_2} = v_{i,d_j} = \frac{1}{m}$ for $j \in [\sqrt{m} - 3]$; the value for all other alternatives is zero.
- For every $i \in S_>$, we set $v_{i,b_i} = 1$; the value for all other alternatives is zero.

- For every $i \in S_1$, we set $v_{i,b_i} = \frac{1}{\sqrt{m}}$, and split the remaining value of $1 - \frac{1}{\sqrt{m}}$ equally among the $m - \sqrt{m}$ alternatives of $C \cup D \cup E$, while the value for the alternatives of $B \setminus \{b_i\}$ is zero.

Hence, the social welfare of alternative b_{i^*} is

$$\text{SW}(b_{i^*}|\mathbf{v}) = \frac{1}{\sqrt{m}}.$$

Since $|S_1| < n$, there exists an agent $i \in S_C \cup S_D \cup S_>$ such that alternative b_i has social welfare

$$\text{SW}(b_i|\mathbf{v}) \geq 1 - \frac{1}{\sqrt{m}} + \frac{1}{m} \geq 1 - \frac{1}{\sqrt{m}},$$

and therefore the distortion is at least $\sqrt{m} - 1$.

If $|S_1| = n$, we define the following valuation profile \mathbf{v} :

- For every agent $i \in [n]$, we set $v_{i,b_i} = v_{i,c_1} = v_{i,c_2} = \frac{1}{\sqrt{m}}$, and split the remaining value of $1 - \frac{3}{\sqrt{m}}$ equally among the $m - \sqrt{m} - 2$ alternatives of $D \cup E$; the value for the alternatives of $B \setminus \{b_i\}$ is zero. This is a valid valuation definition since the value for each alternative in $D \cup E$ is $\frac{\sqrt{m}-3}{\sqrt{m}(m-\sqrt{m}-2)} \leq \frac{1}{\sqrt{m}} \Leftrightarrow (\sqrt{m}-1)^2 \geq 0$.

Hence, alternative b_{i^*} has social welfare

$$\text{SW}(b_{i^*}|\mathbf{v}) = \frac{1}{\sqrt{m}}.$$

But now, the social welfare of c_1 and c_2 is equal to

$$\text{SW}(c_1|\mathbf{v}) = \text{SW}(c_2|\mathbf{v}) = n \frac{1}{\sqrt{m}} = 1,$$

yielding distortion that is at least \sqrt{m} .

Case III: $|S_C| + |S_D| + |S_>| \geq 1$ and \mathcal{M} selects some alternative b_{i^*} for $i^* \in S_C \cup S_D \cup S_>$

We define the following valuation profile \mathbf{v} :

- If $i^* \in S_C \cup S_D$, we set $v_{i^*,j} = \frac{1}{m}$ for every alternative $j \in A$. If $i^* \in S_>$, we split the total value of 1 equally among the $m - \sqrt{m} + 1$ alternatives in $\{b_{i^*}\} \cup C \cup D \cup E$ so that the value of agent i for each such alternative is $\frac{1}{m - \sqrt{m} + 1} \leq \frac{1}{\sqrt{m}}$.
- For every agent $i \in S_C \setminus \{i^*\}$, we set $v_{i,b_i} = 1 - \frac{2}{m}$ and $v_{i,c_1} = v_{i,c_2} = \frac{1}{m}$; the value for all other alternatives is zero.
- For every agent $i \in S_D \setminus \{i^*\}$, we set $v_{i,b_i} = 1 - \frac{1}{\sqrt{m}} + \frac{1}{m}$ and $v_{i,c_1} = v_{i,c_2} = v_{i,d_j} = \frac{1}{m}$ for $j \in [\sqrt{m} - 3]$; the value for all other alternatives is zero.
- For every agent $i \in S_> \setminus \{i^*\}$, we set $v_{i,b_i} = 1$; the value for all other alternatives is zero.

- For every agent $i \in S_1$, we set $v_{i,b_i} = v_{i,c_1} = v_{i,c_2} = \frac{1}{\sqrt{m}}$, and split the remaining value of $1 - \frac{3}{\sqrt{m}}$ equally among the $m - \sqrt{m} - 2$ alternatives in $D \cup E$; the value of agent i for the alternatives of $B \setminus \{b_i\}$ is zero. This is a valid valuation definition since the value for each alternative in $D \cup E$ is $\frac{1 - \frac{3}{\sqrt{m}}}{m - 2\sqrt{m} + 1} \leq \frac{1}{\sqrt{m}} \Leftrightarrow (\sqrt{m} - 1)^2 \geq 0$.

Hence, in any case, the social welfare of alternative b_{i^*} is

$$\text{SW}(b_{i^*} | \mathbf{v}) \leq \frac{1}{\sqrt{m}}.$$

We now distinguish between a couple more cases:

- If $|S_C| + |S_D| + |S_{>}| \geq 2$, then there exists an agent $i \in S_C \cup S_D \cup S_{>} \setminus \{i^*\}$ such that the social welfare of alternative b_i is

$$\text{SW}(b_i | \mathbf{v}) \geq 1 - \frac{1}{\sqrt{m}} + \frac{1}{m} \geq 1 - \frac{1}{\sqrt{m}},$$

yielding distortion at least $\sqrt{m} - 1$.

- If $|S_C| + |S_D| + |S_{>}| = 1$, then since $|S_1| = n - 1 = \sqrt{m} - 1$, alternatives c_1 and c_2 both have social welfare

$$\text{SW}(c_1 | \mathbf{v}) = \text{SW}(c_2 | \mathbf{v}) \geq (\sqrt{m} - 1) \frac{1}{\sqrt{m}} = 1 - \frac{1}{\sqrt{m}}$$

and the distortion is at least $\sqrt{m} - 1$.

The proof is now complete. □

8 Conclusions and Directions for Future Research

In this paper, we studied mechanisms for general single winner elections. In particular, we explored the potential of improving the distortion of deterministic ordinal mechanisms by making a *limited* number of cardinal queries per agent. On this front, we obtained a definitive positive answer. Among other positive results, we showed that it is possible to achieve *constant* distortion by making $O(\log^2 m)$ value or comparison queries per agent, while only *two* value queries are enough to guarantee distortion $O(\sqrt{m})$ in the fundamental case of $n = \Theta(m)$, thus outperforming the best known randomized ordinal mechanism of [Boutilier et al. \[2015\]](#). Quite interestingly, our positive results for value queries hold without any normalization assumptions, which makes them even stronger.

We complemented these results by showing (nearly) tight lower bounds for many interesting cases. For one-query mechanisms we showed a linear lower bound for unrestricted valuation functions and a lower bound of $\Omega(\sqrt{m})$ for unit-sum valuations. Further, for mechanisms that make $O\left(\frac{\log m}{\log \log m}\right)$ queries we showed a superconstant lower bound for unrestricted valuation functions.

Possibly the most obvious open problem is to fill in the gaps between our upper and lower bounds. To this end, we make the following two conjectures.

1-Query Conjecture. *There exists a mechanism that achieves a distortion of $O(\sqrt{m})$ using 1 value query per agent, for unit-sum valuation functions.*

log m -Queries Conjecture. *There exists a mechanism that achieves a constant distortion, using $O(\log m)$ value queries per agent, for unit-sum valuation functions.*

We consider settling these two conjectures the most interesting problems left open in our work. Since our upper bounds for value queries do not make use of the unit-sum normalization, it is conceivable that some clever use of that extra information could possibly lead to better trade-offs.

A natural direction for future work is to consider randomization. Intriguingly, one could consider two different levels of randomization. The first level consists of mechanisms that decide randomly what queries to make to the agents, yet the winning alternative is chosen deterministically. The second level consists of mechanisms that use randomization for both querying and making the final decision. Both of these two classes of randomized mechanisms are very natural and may lead to similar distortion bounds but potentially using fewer queries.

Our work takes a first step towards exploring how powerful ordinal mechanisms with limited access to cardinal information can actually be. Of course, the same idea can be applied to many different contexts, such as participatory budgeting, multi-winner elections, or the metric distortion setting, which has been extensively studied over the past years. As we mentioned in the introduction, [Abramowitz *et al.* \[2019\]](#) already take a step in this direction in the metric setting.

References

- Ben Abramowitz and Elliot Anshelevich. Utilitarians without utilities: Maximizing social welfare for graph problems using only ordinal preferences. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 894–901, 2018.
- Ben Abramowitz, Elliot Anshelevich, and Wennan Zhu. Awareness of voter passion greatly improves the distortion of metric social choice. *CoRR*, abs/1906.10562, 2019.
- Elliot Anshelevich and John Postl. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research*, 58:797–827, 2017.
- Elliot Anshelevich and Shreyas Sekar. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 390–396, 2016.
- Elliot Anshelevich and Wennan Zhu. Tradeoffs between information and ordinal approximation for bipartite matching. In *Proceedings of the 10th International Symposium on Algorithmic Game Theory (SAGT)*, pages 267–279, 2017.
- Elliot Anshelevich and Wennan Zhu. Ordinal approximation for social choice, matching, and facility location problems given candidate positions. In *Proceedings of the 14th International Conference on Web and Internet Economics (WINE)*, pages 3–20, 2018.
- Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.
- Salvador Barbera, Anna Bogomolnaia, and Hans van der Stel. Strategy-proof probabilistic rules for expected utility maximizers. *Mathematical Social Sciences*, 35(2):89–103, 1998.

- Gerardus Benade, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 376–382, 2017.
- Gerardus Benade, Ariel D. Procaccia, and Mingda Qiao. Low-distortion social welfare functions. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- Umang Bhaskar, Varsha Dani, and Abheek Ghosh. Truthful and near-optimal mechanisms for welfare maximization in multi-winner elections. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 925–932, 2018.
- Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100:295–328, 2001.
- Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. Primarily about primaries. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- Ioannis Caragiannis and Ariel D. Procaccia. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9-10):1655–1671, 2011.
- Ioannis Caragiannis, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research*, 58:123–152, 2017.
- Ioannis Caragiannis, Aris Filos-Ratsikas, Swaprava Nath, and Alexandros A. Voudouris. Truthful mechanisms for ownership transfer with expert advice. *CoRR*, abs/1802.01308, 2018.
- Yu Cheng, Shaddin Dughmi, and David Kempe. Of the people: Voting is more effective with representative candidates. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, pages 305–322, 2017.
- Yu Cheng, Shaddin Dughmi, and David Kempe. On the distortion of voting with multiple representative candidates. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 973–980, 2018.
- Brandon Fain, Ashish Goel, Kamesh Munagala, and Nina Prabhu. Random dictators with a random referee: Constant sample complexity mechanisms for social choice. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- Uriel Feige and Moshe Tennenholtz. Responsive lotteries. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*, pages 150–161, 2010.
- Michal Feldman, Amos Fiat, and Iddan Golomb. On voting and facility location. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*, pages 269–286, 2016.

- Aris Filos-Ratsikas and Peter Bro Miltersen. Truthful approximations to range voting. In *Proceedings of the 10th International Conference on Web and Internet Economics (WINE)*, pages 175–188, 2014.
- Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *Proceedings of the 7th Symposium of Algorithmic Game Theory (SAGT)*, pages 1–12, 2014.
- Aris Filos-Ratsikas, Evi Micha, and Alexandros A. Voudouris. The distortion of distributed voting. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*, 2019.
- Mohammad Ghodsi, Mohamad Latifian, and Masoud Seddighin. On the distortion value of the elections with abstention. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- Ashish Goel, Anilesh K Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. Knapsack voting for participatory budgeting, 2016.
- 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)*, pages 287–304, 2017.
- Ashish Goel, Reyna Hulett, and Anilesh K. Krishnaswamy. Relating metric distortion and fairness of social choice rules. In *Proceedings of the 13th Workshop on Economics of Networks (NetEcon)*, page 4:1, 2018.
- Stephen Gross, Elliot Anshelevich, and Lirong Xia. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 544–550, 2017.
- Tyler Lu and Craig Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 280–286, 2011.
- Debmalya Mandal, Ariel D. Procaccia, Nisarg Shah, and David P. Woodruff. Efficient and thrifty voting by any means necessary. 2019.
- Kamesh Munagala and Kangning Wang. Improved metric distortion for deterministic social choice rules. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, pages 245–262, 2019.
- Grzegorz Pierczynski and Piotr Skowron. Approval-based elections and distortion of voting rules. *CoRR*, abs/1901.06709, 2019.
- Ariel D. Procaccia and Jeffrey S. Rosenschein. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents (CIA)*, pages 317–331, 2006.
- John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. 1947.

A Missing Proofs from Subsection 7.2

Proof of Claim 2

Assume towards a contradiction that there exists an alternative $z^* \in A_1$ such that the mechanism \mathcal{M} asks at most $\frac{\lambda}{\lambda+1} \cdot m^{\frac{1}{\lambda+1}}$ agents of $T_1(z^*)$ at the first position, and \mathcal{M} has distortion $\mathcal{D}(\mathcal{M}) \notin \Omega\left(\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}\right)$. Let S be the set of the at least $\frac{1}{\lambda+1} \cdot m^{\frac{1}{\lambda+1}}$ agents of $T_1(z^*)$ that are *not* queried by \mathcal{M} at the first position. Hence, we have that $\delta_{i1} = 1$ for every agent $i \notin S$.

We now define two valuation profiles \mathbf{v}_1 and \mathbf{v}_2 , which are consistent to the ordinal profile and any information revealed by the queries of \mathcal{M} , but differ on the value that the agents of S have for alternative z^* :

- In both \mathbf{v}_1 and \mathbf{v}_2 , every agent $i \notin S$ has value $\delta_{ij} \cdot m^{-\frac{j+1/2}{\lambda+1}} + (1 - \delta_{ij}) \cdot m^{-\frac{j}{\lambda+1}}$ for the alternative $a_{ij} \in A_j$ that she ranks at position $j \in [\lambda]$, value m^{-1} for alternatives x_1 and x_2 , and zero value for her tail alternatives.
- In both \mathbf{v}_1 and \mathbf{v}_2 , every agent $i \in S$ has value $\delta_{ij} \cdot m^{-\frac{j+1/2}{\lambda+1}} + (1 - \delta_{ij}) \cdot m^{-\frac{j}{\lambda+1}}$ for the alternative $a_{ij} \in A_j$ that she ranks at position $j \in [\lambda] \setminus \{1\}$, value m^{-1} for alternatives x_1 and x_2 , and zero value for her tail alternatives.
- In \mathbf{v}_1 , every agent $i \in S$ has value $m^{-\frac{3/2}{\lambda+1}}$ for z^* .
- In \mathbf{v}_2 , every agent $i \in S$ has value 1 for z^* .

Given the definition of these two valuation profiles, it is easy to compute the social welfare of the alternatives:

- The social welfare of every alternative $z \in Y$ is

$$\text{SW}(z|\mathbf{v}_1) = \text{SW}(z|\mathbf{v}_2) = 0.$$

- The social welfare of alternatives x_1 and x_2 is

$$\text{SW}(x_1|\mathbf{v}_1) = \text{SW}(x_1|\mathbf{v}_2) = \text{SW}(x_2|\mathbf{v}_1) = \text{SW}(x_2|\mathbf{v}_2) = m \cdot m^{-1} = 1.$$

- The social welfare of any alternative $z \in \cup_{j \in [\lambda]} A_j \setminus \{z^*\}$ is

$$\begin{aligned} \text{SW}(z|\mathbf{v}_1) = \text{SW}(z|\mathbf{v}_2) &= \frac{\lambda}{\lambda+1} \cdot m^{\frac{j}{\lambda+1}} \cdot m^{-\frac{j+1/2}{\lambda+1}} + \frac{1}{\lambda+1} \cdot m^{\frac{j}{\lambda+1}} \cdot m^{-\frac{j}{\lambda+1}} \\ &= \frac{\lambda}{\lambda+1} \cdot m^{-\frac{1}{2(\lambda+1)}} + \frac{1}{\lambda+1} \leq 1. \end{aligned}$$

- The social welfare of alternative z^* is

$$\begin{aligned} \text{SW}(z^*|\mathbf{v}_1) &= \frac{\lambda}{\lambda+1} \cdot m^{\frac{1}{\lambda+1}} \cdot m^{-\frac{1+1/2}{\lambda+1}} + \frac{1}{\lambda+1} \cdot m^{\frac{1}{\lambda+1}} \cdot m^{-\frac{3/2}{\lambda+1}} = m^{-\frac{1}{2(\lambda+1)}}, \\ \text{SW}(z^*|\mathbf{v}_2) &= \frac{\lambda}{\lambda+1} \cdot m^{\frac{1}{\lambda+1}} \cdot m^{-\frac{1+1/2}{\lambda+1}} + \frac{1}{\lambda+1} \cdot m^{\frac{1}{\lambda+1}} \cdot 1 = \frac{\lambda}{\lambda+1} \cdot m^{-\frac{1}{2(\lambda+1)}} + \frac{1}{\lambda+1} \cdot m^{\frac{1}{\lambda+1}}. \end{aligned}$$

Depending on the choices of the mechanism \mathcal{M} , we set the valuation profile to be either \mathbf{v}_1 or \mathbf{v}_2 so that the distortion is as high as possible. In particular, we have:

- If \mathcal{M} selects alternative z^* , we set the valuation profile to be \mathbf{v}_1 . Hence, the social welfare of the winner z^* is $m^{-\frac{1}{2} \cdot \frac{1}{\lambda+1}}$, while any alternative of X is optimal with social welfare 1, yielding distortion $m^{\frac{1}{2(\lambda+1)}}$.
- If \mathcal{M} selects some alternative $z \in X \cup_{j \in [\lambda]} A_j \setminus \{z^*\}$, we set the valuation profile to be \mathbf{v}_2 . Hence, the social welfare of the winner z is at most 1, while alternative z^* is optimal with social welfare at least $\frac{1}{\lambda+1} \cdot m^{\frac{1}{\lambda+1}}$, yielding distortion at least $\frac{1}{\lambda+1} \cdot m^{\frac{1}{\lambda+1}}$.

In any case, the distortion is $\Omega\left(\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}\right)$ and the proof of the claim follows. \square

Proof of Claim 3

For every alternative $w \in A_{j+1}$, let $S_w \subseteq T_{j+1}(w)$ be the set of agents that rank w at position $(j+1)$ and are queried by \mathcal{M} at the first j positions. By the definition of the ordinal profile, the set $T_{j+1}(w)$ consists of $m^{\frac{j+1}{\lambda+1}}$ agents. These agents are partitioned into $m^{\frac{1}{\lambda+1}}$ sets of size $m^{\frac{j}{\lambda+1}}$ so that the agents of each such set all rank the same alternative of A_j at position j . Therefore, by the assumption of the claim, we have that $|S_w| > m^{\frac{1}{\lambda+1}} \cdot \left(1 - \frac{j}{\lambda+1}\right) \cdot m^{\frac{j}{\lambda+1}} = \left(1 - \frac{j}{\lambda+1}\right) \cdot m^{\frac{j+1}{\lambda+1}}$.

Now, assume towards a contradiction that there exists an alternative $w^* \in A_{j+1}$ such that \mathcal{M} queries at most $\left(1 - \frac{j+1}{\lambda+1}\right) \cdot m^{\frac{j+1}{\lambda+1}}$ of the agents in S_{w^*} at the first $(j+1)$ positions, and $\mathcal{D}(\mathcal{M}) \notin \Omega\left(\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}\right)$. Let S be the set of the agents in S_{w^*} that are not queried by \mathcal{M} at position $(j+1)$. By our discussion so far, we have that $|S| \geq |S_{w^*}| - \left(1 - \frac{j+1}{\lambda+1}\right) \cdot m^{\frac{j+1}{\lambda+1}} > \frac{1}{\lambda+1} \cdot m^{\frac{j+1}{\lambda+1}}$, and therefore $\delta_{i,j+1} = 1$ for every agent $i \notin S$.

We now define two valuation profiles \mathbf{v}_1 and \mathbf{v}_2 , which are consistent to the ordinal profile and any information revealed by the queries of \mathcal{M} , but differ on the value that the agents of S have for alternative w^* :

- In both \mathbf{v}_1 and \mathbf{v}_2 , every agent $i \notin S$ has value $\delta_{i\ell} \cdot m^{-\frac{\ell+1/2}{\lambda+1}} + (1 - \delta_{i\ell}) \cdot m^{-\frac{\ell}{\lambda+1}}$ for the alternative $a_{i\ell} \in A_\ell$ that she ranks at position $\ell \in [\lambda]$, value m^{-1} for alternatives x_1 and x_2 , and zero value for her tail alternatives.
- In both \mathbf{v}_1 and \mathbf{v}_2 , every agent $i \in S$ has value $\delta_{i\ell} \cdot m^{-\frac{\ell+1/2}{\lambda+1}} + (1 - \delta_{i\ell}) \cdot m^{-\frac{\ell}{\lambda+1}}$ for the alternative $a_{i\ell} \in A_\ell$ that she ranks at position $\ell \in [\lambda] \setminus \{j+1\}$, value m^{-1} for alternatives x_1 and x_2 , and zero value for her tail alternatives.
- In \mathbf{v}_1 , every agent $i \in S$ has value $m^{-\frac{j+3/2}{\lambda+1}}$ for w^* .
- In \mathbf{v}_2 , every agent $i \in S$ has value $m^{-\frac{j+1/2}{\lambda+1}}$ for w^* .

Given the definition of these two valuation profiles, it is easy to compute the social welfare of the alternatives:

- The social welfare of every alternative $z \in Y$ is

$$\text{SW}(z|\mathbf{v}_1) = \text{SW}(z|\mathbf{v}_2) = 0.$$

- The social welfare of alternatives x_1 and x_2 is

$$\text{SW}(x_1|\mathbf{v}_1) = \text{SW}(x_1|\mathbf{v}_2) = \text{SW}(x_2|\mathbf{v}_1) = \text{SW}(x_2|\mathbf{v}_2) = m \cdot m^{-1} = 1.$$

- The social welfare of any alternative $z \in \cup_{\ell \in [\lambda]} A_\ell \setminus \{w^*\}$ is

$$\begin{aligned} \text{SW}(z|\mathbf{v}_1) &= \text{SW}(z|\mathbf{v}_2) = \frac{\lambda}{\lambda+1} \cdot m^{\frac{\ell}{\lambda+1}} \cdot m^{-\frac{\ell+1/2}{\lambda+1}} + \frac{1}{\lambda+1} \cdot m^{\frac{\ell}{\lambda+1}} \cdot m^{-\frac{\ell}{\lambda+1}} \\ &= \frac{\lambda}{\lambda+1} \cdot m^{-\frac{1}{2(\lambda+1)}} + \frac{1}{\lambda+1} \leq 1. \end{aligned}$$

- The social welfare of alternative w^* is

$$\begin{aligned} \text{SW}(w^*|\mathbf{v}_1) &= \frac{\lambda}{\lambda+1} \cdot m^{\frac{j+1}{\lambda+1}} \cdot m^{-\frac{j+1+1/2}{\lambda+1}} + \frac{1}{\lambda+1} \cdot m^{\frac{j+1}{\lambda+1}} \cdot m^{-\frac{j+3/2}{\lambda+1}} = m^{-\frac{1}{2(\lambda+1)}}, \\ \text{SW}(w^*|\mathbf{v}_2) &= \frac{\lambda}{\lambda+1} \cdot m^{\frac{j+1}{\lambda+1}} \cdot m^{-\frac{j+1+1/2}{\lambda+1}} + \frac{1}{\lambda+1} \cdot m^{\frac{j+1}{\lambda+1}} \cdot m^{-\frac{j+1/2}{\lambda+1}} \\ &= \frac{\lambda}{\lambda+1} \cdot m^{-\frac{1}{2(\lambda+1)}} + \frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}. \end{aligned}$$

Depending on the choices of the mechanism \mathcal{M} , we set the valuation profile to be either \mathbf{v}_1 or \mathbf{v}_2 so that the distortion is as high as possible. In particular, we have:

- If \mathcal{M} selects alternative w^* , we set the valuation profile to be \mathbf{v}_1 . Hence, the social welfare of the winner w^* is $m^{-\frac{1}{2(\lambda+1)}}$, while any alternative of X is optimal with social welfare 1, yielding distortion $m^{\frac{1}{2(\lambda+1)}}$.
- If \mathcal{M} selects some alternative $z \in X \cup_{\ell \in [\lambda]} A_\ell \setminus \{w^*\}$, we set the valuation profile to be \mathbf{v}_2 . Hence, the social welfare of the winner z is at most 1, while alternative w^* is optimal with social welfare at least $\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}$, yielding distortion at least $\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}$.

In any case, the distortion is $\Omega\left(\frac{1}{\lambda+1} \cdot m^{\frac{1}{2(\lambda+1)}}\right)$ and the proof of the claim follows. \square