

DON'T ROLL THE DICE, ASK TWICE: THE TWO-QUERY DISTORTION OF MATCHING PROBLEMS AND BEYOND*

GEORGIOS AMANATIDIS[†], GEORGIOS BIRMPAS[‡], ARIS FILOS-RATSIKAS[§], AND ALEXANDROS A. VOUDOURIS[¶]

Abstract. In most social choice settings, the participating agents express their preferences over the different alternatives in the form of linear orderings. While this clearly simplifies preference elicitation, it inevitably leads to poor performance with respect to optimizing a cardinal objective, such as the social welfare, since the values of the agents remain virtually unknown. This loss in performance because of lack of information is measured by the notion of *distortion*. A recent array of works put forward the agenda of designing mechanisms that learn the values of the agents for a small number of alternatives via *queries*, and use this limited extra information to make better-informed decisions, thus improving distortion. Following this agenda, in this work we focus on a class of combinatorial problems that includes most well-known matching problems and several of their generalizations. For problems such as One-Sided Matching, Two-Sided Matching, General Graph Matching, and Short Cycle Packing, we design *two-query* mechanisms that achieve the best-possible worst-case distortion in terms of social welfare, and outperform the best-possible expected distortion achieved by randomized ordinal mechanisms. Our results extend to problems like k -Constrained Resource Allocation, General Graph k -Matching, and k -Clique Packing, when k is restricted to be any constant.

Key words. distortion, queries, matching, social choice, graph problems

MSC codes. 68Q25, 68W25, 68W40, 91B12, 91B14, 91B68

1. Introduction. The notion of *distortion* in social choice settings was defined to capture the loss in aggregate objectives due to the lack of precise information about the preferences of the participants [38]. More concretely, the distortion was originally defined as a measure of the deterioration of the total happiness of the agents when access is given only to the (ordinal) preference rankings of the agents, rather than to the complete numerical (cardinal) information about their preferences. This research agenda has successfully been applied to a plethora of different settings, giving rise to a rich and vibrant line of work in major venues at the intersection of computer science and economics. For a comprehensive overview, see the survey of Anshelevich et al. [8].

Out of all of these scenarios, some of the most fundamental are *matching* problems, in which agents are matched to items or other agents, aiming to maximize the *social welfare* of the matching (the total value of the agents). An example is the classic *One-Sided Matching* setting [29], where the goal is to match n items to n agents based on the preferences of the agents over the items. For this setting, Filos-Ratsikas et al. [24] showed that the best achievable distortion is $\Theta(\sqrt{n})$. Importantly, this guarantee is only possible if one is allowed to use *randomization* and the values of the agents are *normalized*, since if any of these assumptions is relaxed, it is impossible to achieve sub-linear distortion using only ordinal information.

Moving on from merely preference rankings, Amanatidis et al. [4] recently put forward

*A conference version of this work has been included in the Proceedings of the 2022 Conference on Neural Information Processing Systems (NeurIPS 2022) [5].

Funding: This work was supported by the ERC Advanced Grant 788893 AMDROMA “Algorithmic and Mechanism Design Research in Online Markets”, the MIUR PRIN project ALGADIMAR “Algorithms, Games, and Digital Markets”, the NWO Veni project No. VI.Veni.192.153, and the project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program.

[†]School of Mathematics, Statistics and Actuarial Science, University of Essex, United Kingdom, and Archimedes/Athena RC, Greece (georgios.amanatidis@essex.ac.uk).

[‡]Department of Computer Science, University of Liverpool, United Kingdom (g.birmpas@liverpool.ac.uk).

[§]School of Informatics, University of Edinburgh, United Kingdom (aris.filos-ratsikas@ed.ac.uk).

[¶]School of Computer Science and Electronic Engineering, University of Essex, United Kingdom (alexandros.voudouris@essex.ac.uk).

the agenda of studying the tradeoffs between information and efficiency, when the employed mechanisms are equipped with the capability of learning the values of the agents via *queries*. The rationale is that asking the agents for more detailed information about only *a few* options is still cognitively not too burdensome, and could result in notable improvements on the distortion. This was indeed confirmed in that work for general social choice, and in a follow-up work for several matching problems [6]. Specifically, the latter work shows that it is possible to obtain distortion $O(n^{1/k})$ with $O(\log n)$ queries per agent for any constant integer k , and distortion $O(1)$ with $O(\log^2 n)$ queries per agent. Crucially, the mechanisms achieving these bounds do not use randomization nor demand the values to be normalized.

While these works make a significant first step, they leave some important questions unanswered. The mechanisms they propose require a logarithmic number of queries to achieve *any* significant improvement. Answering that many queries might still be cognitively too demanding for the agents, especially when there is a large number of possible options. The main high-level motivation of this research agenda is that a small amount of information can be more valuable than randomization. But really constitutes a “small amount”? Ideally, we would like to design mechanisms that make only a few queries per agent, independently of the size of the input parameters. Since with a single query, sub-linear distortion bounds are not possible [6, 4], the first fundamental question that we would like to answer is the following:

What is the best achievable distortion when we can only ask two queries per agent?

1.1. Results and Technical Overview. We settle the aforementioned question for several matching problems, such as *One-Sided Matching*, *Two-Sided Matching*, and *General Graph Matching*; our approach extends to other more general graph-theoretic problems as well. For all matching problems considered, we show that there is a deterministic mechanism that makes two queries per agent, runs in polynomial time, and achieves a distortion of $O(\sqrt{n})$. This upper bound is based on a novel mechanism, which we call MATCH-TWOQUERIES in the case of One-Sided Matching (see Mechanism 3.2). The mechanism asks two queries per agent and computes a maximum-weight matching based on the revealed values due to these queries. It starts by querying the agents at the first position of their preference rankings. For the second query, it computes a certain type of assignment A of agents to items (or agents to agents in more general matching problems), to which we refer as a *sufficiently representative assignment*, and queries the agents about the items they are assigned to in A . The existence of such an assignment for all instances is far from trivial, and one of our main technical contributions is to show its existence and efficient computation for the wide range of problems we consider.

We also show that no deterministic mechanism for these settings that makes two queries per agent can achieve a distortion better than $\Omega(\sqrt{n})$. This lower bound follows by a more general construction yielding a lower bound of $\Omega(n^{1/\lambda})$ on the distortion of any mechanism that makes a constant number λ of queries per agent. This mirrors the corresponding lower bounds of Amanatidis et al. [6] for One-Sided Matching.

These results extend to graph-theoretic problems where the maximum degree of a feasible solution is bounded by a parameter k , such as k -Constrained Resource Allocation, General Graph k -Matching, and k -Clique Packing, although the corresponding mechanisms do not always run in polynomial time. The distortion in this case is $O(k^2\sqrt{n})$ and the corresponding lower bound is $\Omega(n^{1/\lambda}/k^3)$ for mechanisms that make λ queries per agent. That is, when k is restricted to be any constant, our mechanisms still achieve the (asymptotically) best possible distortion with two queries per agent.

Despite the generality of our results, their most impressive implications are for One-Sided Matching: We show that by using only *two* cardinal queries per agent, we can match

the bound of $\Theta(\sqrt{n})$ for purely ordinal mechanisms, *without requiring randomization or any normalization*. MATCH-TWOQUERIES clearly also outperforms another mechanism of Amanatidis et al. [6], which uses two queries and achieves a distortion of $O(n^{2/3}\sqrt{\log n})$ assuming that the values of each agent sum up to 1. In contrast, our mechanism works for unrestricted values, and achieves the best possible distortion of $O(\sqrt{n})$ based on conceptually much simpler ideas.

Results for general social choice. Given that our approach works for a wide variety of matching problems, one might be curious as to whether similar arguments could be used to show bounds for the general social choice setting, where n agents have preferences over m alternatives, and the goal is to select an alternative with high social welfare; this was, after all, the original setting that Amanatidis et al. [4] studied in the introduction of the information-distortion tradeoff research agenda.¹ In this setting, the situation is quite similar: the upper bounds follow by mechanisms that ask $O(\log m)$ queries, and nothing positive is known for smaller numbers of queries.

For the upper bounds, we show a conditional result. We show that a mechanism with structure similar to that of MATCH-TWOQUERIES can indeed achieve a distortion of $O(\sqrt{m})$ using only two queries, subject to being able to compute a sufficiently representative set of alternatives, which is analogous to the sufficiently representative assignment in matching problems. It turns out that this property is very closely connected to the notion of an (*approximately*) *stable committee* [30, 21], and it follows that it exists when $m = \Omega(n)$, thus allowing us to obtain the desired bound of $O(\sqrt{m})$ when this is true. This case is quite natural, as it captures instances where a group of people need to decide over a large set of possible options (e.g., shortlisting candidates for a job, deciding the best paper for a conference, etc.). Interestingly, in contrast to the matching setting for which we show that sufficiently representative assignments can be found via a simple greedy algorithm, computing sufficiently representative sets of alternatives in general social choice requires rather involved techniques [30, 21]. An obvious open question here is whether the $O(\sqrt{m})$ bound can also be achieved by asking only two queries when $m = o(n)$. This seems to be a more challenging task to prove; we discuss it further in Section 6.

We also show that the bound of $O(\sqrt{m})$ is the best possible, as part of a more general distortion lower bound of $\Omega(m^{1/\lambda})$ for mechanisms that make a constant number λ of queries per agent; the latter result significantly improves the previously known lower bound of $\Omega(m^{1/2(\lambda+1)})$ [4].

Roadmap. For the sake of presentation, we fully demonstrate how our methodology works for the *One-Sided Matching* problem in Section 3. Before doing so, we start with some necessary notation and terminology in Section 2. In Section 4, we briefly discuss other graph-theoretic problems for which our methodology can be applied. Our results for general social choice are presented in Section 5. We conclude with some interesting open problems in Section 6.

1.2. Additional Related Work. The literature on the distortion of ordinal mechanisms in social choice is long and extensive, focusing primarily on settings with normalized utilities (e.g., [16, 22, 18, 25]), or with metric preferences (e.g., [7, 10, 9, 19, 20, 28, 37, 31]); see the survey of Anshelevich et al. [8] for a detailed exposition. The distortion of mechanisms for One-Sided Matching and more general graph-theoretic problems has been studied in a series of works for a variety of preference models, but solely with ordinal information [12, 2, 13, 11, 24, 17].

¹ It is easy to see that the general social choice setting generalizes the matching setting, as any matching can be seen as an alternative, but agents have the same value for alternatives/matchings that assign them to the same items. See the proof of Theorem 6 of Filos-Ratsikas et al. [26] for details.

Besides the papers of Amanatidis et al. [4, 6], the effect of limited cardinal information on the distortion has also been studied in other works [3, 34, 35, 15]. The One-Sided Matching problem has also been studied by Ma et al. [33] using a different type of cardinal queries and focusing on Pareto optimality (rather than social welfare).

Our upper bound for the general social choice setting makes use of the results of Cheng et al. [21] and Jiang et al. [30] for (approximately) stable committees (see also Aziz et al. [14]); a stable committee is very similar to a representative set of alternatives in our terminology. Cheng et al. [21] showed that, while *exactly* stable committees do not always exist [30], finding a random version of such committees, coined *stable lotteries*, is always possible and can be done in polynomial time. Later on, Jiang et al. [30] showed that, via an intricate derandomization process, stable lotteries can yield approximately stable committees, where the approximation is a small multiplicative constant; for our purposes, this is sufficient. Interestingly, very recently, Ebadian et al. [22] used stable lotteries to construct a purely ordinal randomized social choice mechanism that achieves the best possible distortion under unit-sum normalized values.

2. Preliminaries on One-Sided Matching, Mechanisms, and Distortion. In One-Sided Matching, there is a set \mathcal{N} of n agents and a set \mathcal{A} of n items. Each agent $i \in \mathcal{N}$ has a value $v_{i,j}$ for each item $j \in \mathcal{A}$; we refer to the matrix $\mathbf{v} = (v_{i,j})_{i \in \mathcal{N}, j \in \mathcal{A}}$ as the *valuation profile*. A (*one-sided*) *matching* $X : \mathcal{N} \rightarrow \mathcal{A}$ is a bijection from the set of agents to the set of items, i.e., each agent is *matched* to a different single item. Our goal is to choose a matching X to maximize the *social welfare*, defined as the total value of the agents for the items they have been matched to according to X : $\text{SW}(X|\mathbf{v}) = \sum_{i \in \mathcal{N}} v_{i,X(i)}$. Usually \mathbf{v} is clear from the context, so we then simplify our notation to $\text{SW}(X)$ for the social welfare of matching X .

As in most of the related literature, we assume that we do not have access to the valuation profile of the agents. Instead, we have access to the *ordinal preference* \succ_i of each agent i , which is derived from the values of the agent for the items, such that $a \succ_i b$ if $v_{i,a} \geq v_{i,b}$; we refer to the vector $\succ_{\mathbf{v}} = (\succ_i)_{i \in \mathcal{N}}$ as the *ordinal profile* of the agents.

A *mechanism* \mathcal{M} in our setting operates as follows: It takes as input the ordinal profile $\succ_{\mathbf{v}}$ of the agents. It then makes a number $\lambda \geq 1$ of *queries* per agent to learn part of the valuation profile. In particular, each agent is asked her value for at most λ items. Given the answers to the queries, and also using the ordinal profile, \mathcal{M} computes a feasible solution (here a matching) $\mathcal{M}(\succ_{\mathbf{v}})$.

In this paper we focus on mechanisms that make two queries per agent, i.e., $\lambda = 2$, and compute a solution of high social welfare. However, pinpointing an (approximately) optimal solution without having full access to the valuation profile of the agents can be quite challenging; the ordinal profile may be consistent with a huge number of different valuation profiles, even after the queries. Nevertheless, we aim to achieve the best asymptotic performance possible, as quantified by the notion of *distortion*.

DEFINITION 2.1. *The distortion of a mechanism \mathcal{M} is the worst-case ratio (over the set \mathcal{V} of all valuation profiles in instances with n agents and n items) between the optimal social welfare and the social welfare of the solution chosen by \mathcal{M} :*

$$\text{dist}(\mathcal{M}) = \sup_{\mathbf{v} \in \mathcal{V}, |\mathcal{N}|=n, |\mathcal{A}|=n} \frac{\max_{X \in \mathcal{X}} \text{SW}(X|\mathbf{v})}{\text{SW}(\mathcal{M}(\succ_{\mathbf{v}})|\mathbf{v})},$$

where \mathcal{X} is the set of all matchings between \mathcal{N} and \mathcal{A} .

3. An Optimal Two-Query Mechanism. In this section, we present a mechanism for One-Sided Matching that makes two queries per agent and achieves a distortion of $O(\sqrt{n})$.

Due to the lower bound of $\Omega(n^{1/\lambda})$ on the distortion of any mechanism that can make up to λ queries per agent shown by Amanatidis et al. [6], our mechanism is asymptotically best possible when $\lambda = 2$.

Without any normalization assumptions about the valuation functions, it is easy to see that a mechanism cannot have *any* guarantee unless it queries every agent about her favorite item. However, there are no obvious criteria suggesting how to use the *second* query. Before we present the details of our mechanism, we define a particular type of assignment of agents to items that will be critical for deciding where to make the second queries.

DEFINITION 3.1. *A many-to-one assignment A of agents to items (i.e., each agent i is assigned to one item $A(i)$, but multiple agents may be assigned to the same item) is a sufficiently representative assignment if*

- (a) *For every item $j \in \mathcal{A}$, $|\{i \in \mathcal{N} : A(i) = j\}| \leq \sqrt{n}$, i.e., there are at most \sqrt{n} agents assigned to j ;*
- (b) *For any matching X , $|\{i \in \mathcal{N} : X(i) \succ_i A(i)\}| \leq \sqrt{n}$, i.e., there are at most \sqrt{n} agents that prefer the item they are matched to in X to the item they are assigned to in A .*

We show that a sufficiently representative assignment exists for any instance and can be efficiently computed. We do so via a simple polynomial-time algorithm, which we call \sqrt{n} -SERIAL DICTATORSHIP (Mechanism 3.1). This algorithm creates \sqrt{n} copies of each item and then runs a serial dictatorship algorithm, which first fixes an ordering of the agents and then assigns each agent to her most preferred available item according to her ordinal preference. It is easy to see that the running time of \sqrt{n} -SERIAL DICTATORSHIP is polynomial (in particular, it is $O(n^{1.5})$).

Mechanism 3.1 \sqrt{n} -SERIAL DICTATORSHIP($\mathcal{N}, \mathcal{A}, \succ_{\mathbf{v}}$)

- 1: Let \mathcal{B} be a multiset containing \sqrt{n} copies of each $j \in \mathcal{A}$
 - 2: **for** every agent $i \in \mathcal{N}$ **do**
 - 3: Let α_i be a most preferred item of agent i in \mathcal{B}
 - 4: Remove α_i from \mathcal{B}
 - 5: **end for**
 - 6: **return** $A = (\alpha_i)_{i \in \mathcal{N}}$
-

THEOREM 3.2. *For any instance, the output of \sqrt{n} -SERIAL DICTATORSHIP is a sufficiently representative assignment.*

Proof. Let A be the output of the algorithm. During the execution of the algorithm, whenever every copy of an item has been assigned, we say that such an item is *exhausted*. Assume, towards a contradiction, that A is not a sufficiently representative assignment. By construction, every item is assigned to at most \sqrt{n} agents, so there must be a matching violating the second condition of Definition 3.1. That is, there is a subset of items \mathcal{A}' and a subset of agents \mathcal{N}' , such that $|\mathcal{A}'| = |\mathcal{N}'| > \sqrt{n}$, and each agent $i \in \mathcal{N}'$ prefers to be assigned to a distinct item $\beta_i \in \mathcal{A}'$ (i.e., $\beta_i \neq \beta_j$ for $i \neq j$) instead of the item she is assigned to in A .

Consider any agent $i \in \mathcal{N}'$. The fact that this agent was not assigned to β_i by the algorithm implies that when the agent was picked, item β_i was exhausted. Since this is true for all agents in \mathcal{N}' , at the end of the algorithm all items of \mathcal{A}' must be exhausted. However, an item is exhausted when all its \sqrt{n} copies have been assigned and there are n agents in total, so we can only have as many as $n/\sqrt{n} = \sqrt{n}$ exhausted items. This means that $|\mathcal{A}'| \leq \sqrt{n}$, a contradiction. \square

Now, our mechanism MATCH-TWOQUERIES (Mechanism 3.2) first queries every agent

about her favorite item. Next, it computes a sufficiently representative assignment A and queries each agent about the item she is assigned to in A . Finally, it outputs a matching that maximizes the social welfare based *only* on the revealed values (all other values are set to 0). Although computational efficiency is not our primary focus here, if we use a polynomial-time algorithm for computing a maximum weight matching (e.g., the Hungarian method [32]), MATCH-TWOQUERIES runs in polynomial time as well.

Mechanism 3.2 MATCH-TWOQUERIES($\mathcal{N}, \mathcal{A}, \succ_{\mathbf{v}}$)

- 1: Query each $i \in \mathcal{N}$ about her favorite item w.r.t. \succ_i
 - 2: Compute a *sufficiently representative assignment* A
 - 3: Query each agent about the item she is assigned to in A
 - 4: Set all non-revealed values to 0
 - 5: **return** a maximum-weight perfect matching Y
-

Of course, if we compare the mechanism’s behaviour to an actual optimal matching X , we expect to see that we asked agents about the “wrong” items most of the time: for many agents the second query is about better items than what they are matched to in X , and for many agents it is about worse items. The desired bound of $O(\sqrt{n})$ on the distortion of MATCH-TWOQUERIES is established by balancing the loss due to each of these two cases.

THEOREM 3.3. MATCH-TWOQUERIES has distortion $O(\sqrt{n})$.

Proof. Consider any instance with valuation profile \mathbf{v} . Let Y be the matching computed by the MATCH-TWOQUERIES mechanism when given as input the ordinal profile $\succ_{\mathbf{v}}$, and let X be an optimal matching. Let $\text{SW}_R(Y)$ be the *revealed* social welfare of Y , i.e., the total value of the agents for the items they are matched to in Y and for which they *were queried* about. We will show that $\text{SW}(X) \leq (1 + 2\sqrt{n}) \cdot \text{SW}_R(Y)$, and then use the fact that $\text{SW}(Y) \geq \text{SW}_R(Y)$ to directly get the desired bound on the distortion.

We can write the optimal social welfare as

$$\text{SW}(X) = \text{SW}_R(X) + \text{SW}_C(X),$$

where $\text{SW}_R(X)$ is the revealed social welfare of X that takes into consideration only the values revealed by the queries, whereas $\text{SW}_C(X)$ is the *concealed* social welfare of X that takes into consideration only the values not revealed by any queries. Since Y is the matching that maximizes the social welfare based only on the revealed values, we have that

$$(3.1) \quad \text{SW}_R(X) \leq \text{SW}_R(Y).$$

To bound the quantity $\text{SW}_C(X)$, let S be the set of agents who are not queried about the items they are matched to in X . We partition S into the following two subsets consisting of agents for whom the second query of the mechanism is used to ask about items that the agents consider *better* or *worse* than the items they are matched to in X . Recall that an agent i is queried about the item $A(i)$ she is assigned to according to the sufficiently representative assignment A . So, S is partitioned into

$$S^{\geq} = \{i \in S : v_{i,A(i)} \geq v_{i,X(i)}\}, \quad \text{and} \quad S^{<} = \{i \in S : v_{i,A(i)} < v_{i,X(i)}\}.$$

Given these sets, we can now write

$$\text{SW}_C(X) = \text{SW}_C^{\geq}(X) + \text{SW}_C^{<}(X),$$

where

$$\text{SW}_C^{\geq}(X) = \sum_{i \in S^{\geq}} v_{i,X(i)}$$

and

$$\text{SW}_C^{\leq}(X) = \sum_{i \in S^{\leq}} v_{i,X(i)}.$$

For every item j , let $S_j^{\geq} = \{i \in S^{\geq} : A(i) = j\}$ be the set of all agents in S^{\geq} that are queried about j by the mechanism using the second query. Thus, $S^{\geq} = \bigcup_{j \in \mathcal{A}} S_j^{\geq}$. Since A is a sufficiently representative assignment, $|S_j^{\geq}| \leq \sqrt{n}$ for every item j . Therefore,

$$\begin{aligned} \text{SW}_C^{\geq}(X) &= \sum_{j \in \mathcal{A}} \sum_{i \in S_j^{\geq}} v_{i,X(i)} \leq \sum_{j \in \mathcal{A}} \sum_{i \in S_j^{\geq}} v_{i,j} \\ &\leq \sum_{j \in \mathcal{A}} |S_j^{\geq}| \cdot \max_{i \in S_j^{\geq}} v_{i,j} \leq \sqrt{n} \sum_{j \in \mathcal{A}} \max_{i \in S_j^{\geq}} v_{i,j} \\ (3.2) \quad &\leq \sqrt{n} \cdot \text{SW}_R(Y). \end{aligned}$$

For the last inequality, recall that A assigns every agent to a single item, and thus the sets S_j^{\geq} are disjoint. In addition, the values of all the agents for the items they are matched to according to A are revealed by the second query of the mechanism. Notice that there exists a feasible matching in which the agent in S_j^{\geq} that has the maximum value for j is matched to j . Since Y is a matching that maximizes the social welfare based on the revealed values, from the above, we obtain that $\text{SW}_R(Y) \geq \sum_{j \in \mathcal{A}} \max_{i \in S_j^{\geq}} v_{i,j}$.

Next consider the quantity $\text{SW}_C^{\leq}(X)$. By the fact that A is a sufficiently representative assignment, it follows that $|S^{\leq}| \leq \sqrt{n}$; otherwise X would constitute a matching for which there are strictly more than \sqrt{n} agents that prefer the item they are matched to in X to the item they are assigned to by A . Combined with the fact that all agents are queried at the first position of their ordinal preferences, we obtain

$$\begin{aligned} \text{SW}_C^{\leq}(X) &= \sum_{i \in S^{\leq}} v_{i,X(i)} \leq \sum_{i \in S^{\leq}} \max_{j \in \mathcal{A}} v_{i,j} \\ &\leq |S^{\leq}| \cdot \max_{i \in S^{\leq}} \max_{j \in \mathcal{A}} v_{i,j} \\ (3.3) \quad &\leq \sqrt{n} \cdot \text{SW}_R(Y). \end{aligned}$$

The bound follows directly by (3.1), (3.2) and (3.3). \square

4. Further Combinatorial Optimization Problems. The approach of Section 3 can be extended to a much broader class of graph-theoretic problems. Informally, our approach works when the objective is to maximize an additive function over subgraphs of a given graph which contain all “small” matchings and have bounded maximum degree. Although technically this is not necessary for our results, we consider this maximum degree to be a constant in all of our definitions below, as this is the case when our results are asymptotically tight. Nevertheless, we include it as a parameter in the statements of Theorems 4.2 and 4.5 for full generality. We make the space of feasible solutions more precise in the following definition.

DEFINITION 4.1. *Given a constant $k \in \mathbb{N}$ and a weighted graph G on n nodes, we say that a family \mathcal{F} of subgraphs of G is a matching extending k -family if:*

- *Graphs in \mathcal{F} have maximum degree at most k ;*

- For any matching M of G of size at most $\lfloor n/3k \rfloor$, there is a graph in \mathcal{F} containing M .²

Clearly, the set of matchings of a graph (viewed as subgraphs rather than subsets of edges) is a matching extending 1-family, but it is not hard to see that Definition 4.1 captures other constraints, like subgraphs that are unions of disjoint paths and cycles (matching extending 3-family) or unions of disjoint cliques of size k (matching extending $(k-1)$ -family).

We are ready to introduce the general full information optimization problem that we tackle here; we then move on to its social choice analog. As this is a special case of the class of problems captured by Max-on-Graphs (introduced by Amanatidis et al. [6]), we use a similar formulation and name. Note that, in the above definition, the family \mathcal{F} is independent of the weights w . This is necessary as w will be unknown in general.

k -Max-on-Graphs: Given some $k \in \mathbb{N}$, a weighted graph $G = (U, E, w)$, and a concise description of a matching extending k -family \mathcal{F} , find a solution

$$H^* \in \arg \max_{H \in \mathcal{F}} \sum_{e \in E(H)} w(e).$$

One-Sided Matching, as studied in Section 3, is the special case of k -Max-On-Graphs, where G is the complete bipartite graph on the set of agents \mathcal{N} and the set of items \mathcal{A} , the weight of an edge $\{i, j\}$ is the value $v_{i,j}$ of agent i for item j , and \mathcal{F} contains all the 1-factors (matchings) of G . Note that the weights of the graph in this case are defined in terms of the valuation functions of the agents. Moreover, recall that in our setting only the ordinal preferences of the agents are given and their cardinal values can be accessed only via queries; so, we do not know the weights that have not been revealed by a query. This is the case for all the problems we are interested in, and is captured by the next definition. To avoid unnecessary notation, items are modeled as dummy agents with all their cardinal values equal to 0. In addition, we write $w(H) := \sum_{e \in E(H)} w(e)$.

Ordinal- k -Max-on-Graphs: Fix a $k \in \mathbb{N}$ and let \mathcal{N} be a set of n agents. A weighted graph $G = (\mathcal{N}, E, w)$ is given *without* its weights. Every agent $i \in \mathcal{N}$ has a (private) valuation function $v_i : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$, such that, for every $e = \{i, j\} \in E$,

$$w(e) = v_i(j) + v_j(i).$$

We are also given an *ordinal profile* $\succ_{\mathbf{v}} = (\succ_i)_{i \in \mathcal{N}}$ that is consistent to $\mathbf{v} = (v_i)_{i \in \mathcal{N}}$, and a concise description of a matching extending k -family \mathcal{F} . The goal is to find $H^* \in \arg \max_{H \in \mathcal{F}} w(H)$.

Besides One-Sided Matching, a large number of problems that are relevant to computational social choice are captured by Ordinal- k -Max-on-Graphs as special cases where the space of feasible input graphs is restricted. We give a few examples:

General Graph Matching: Given a weighted graph $G = (U, E, w)$, find a matching of maximum weight, i.e., \mathcal{F} contains the matchings of G and clearly is a matching extending 1-family. In the social choice analog of the problem, $U = \mathcal{N}$.

Two-Sided Matching: This is a special case of General Graph Matching in which $G = (U_1 \cup U_2, E, w)$ is a bipartite graph. It is an extensively studied problem in economics and computational social choice [27, 39].

² The constant 3 used here is arbitrary, subject to allowing for our proofs to work. A definition that would fully capture the intuition of what is done here would have two constant parameters $k, \ell \in \mathbb{N}$ and would refer to matchings of size at most $\lfloor n/\ell k \rfloor$.

***k*-Clique Packing:** Given a weighted complete graph $G = (U, E, w)$, where $|U| = n$ is a multiple of $k \geq 2$, the goal is to partition U into $|U|/k$ clusters of size k to maximize the total weight of the edges in the clusters. That is, \mathcal{F} contains all spanning subgraphs of G that are the union of cliques of size k . As claimed above, \mathcal{F} is a matching extending $(k - 1)$ -family: clearly every graph in \mathcal{F} has maximum degree $k - 1$, and any matching of size $\lfloor n/(3(k - 1)) \rfloor$ (which is less than n/k) can be extended to a graph in \mathcal{F} by arbitrarily grouping each pair of matched nodes with $k - 2$ unmatched nodes, and then arbitrarily grouping the remaining nodes k at a time.

This problem generalizes General Graph Matching (for which $k = 2$) and often is referred to as Max n/k -Sum Clustering in the literature; see [11]. In its social choice analog, $U = \mathcal{N}$.

General Graph *k*-Matching: Given a weighted graph $G = (U, E, w)$, find a k -matching of maximum weight, i.e., \mathcal{F} contains all the subgraphs of G where each node has degree at most k . As \mathcal{F} already contains all matchings of G of any size, it is straightforward that it is a matching extending k -family. In the social choice analog of the problem, $U = \mathcal{N}$.

***k*-Constrained Resource Allocation:** Given a bipartite weighted graph $G = (U_1 \cup U_2, E, w)$, the goal is to assign at most k nodes of U_2 to each node in U_1 so that the total weight of the corresponding edges is maximized. That is, \mathcal{F} contains the subgraphs of G where each node in U_1 has degree at most k and each node in U_2 has degree at most 1. Again, \mathcal{F} already contains all matchings of G of any size, so it is a matching extending k -family.

This problem generalizes One-Sided Matching. In its social choice analog, $\mathcal{N} = U_1 \cup U_2$ is partitioned into the “actual agents” $\mathcal{N}_1 = U_1$ and the “items” $\mathcal{N}_2 = U_2$, and $v_i(j)$ can be strictly positive only for $i \in \mathcal{N}_1, j \in \mathcal{N}_2$.

Short Cycle Packing: Given an integer ℓ and a weighted complete graph $G = (U, E, w)$, the goal is to find a collection of node-disjoint cycles of length at most ℓ so that their total weight is maximized. Here, \mathcal{F} contains any such collection of short cycles. Arguing as in k -Clique Packing, it is straightforward to see that \mathcal{F} is a matching extending $(\ell - 1)$ -family (although it is not hard to show that it is actually a matching extending 3-family). The social choice analog of the problem has $U = \mathcal{N}$, and is closely related to Clearing Kidney ℓ -Exchanges [1].

It is straightforward to extend the notion of distortion (Definition 2.1) for Ordinal- k -Max-on-Graphs by taking the supremum over all instances of a certain size n and letting \mathcal{X} be the set of feasible solutions of each instance.

As already discussed in the Introduction, for One-Sided Matching, Amanatidis et al. [6] showed a lower bound of $\Omega(n^{1/\lambda})$ on the distortion of all deterministic mechanisms that can make up to $\lambda \geq 1$ queries per agent. Using this, we can get the analogous result for all the aforementioned problems. Although for some of them, like Two-Sided Matching and General Graph Matching, the lower bound is immediate, here we show it for any problem that is captured by Ordinal- k -Max-on-Graphs and is general enough. To make the latter precise, we only consider special cases of Ordinal- k -Max-on-Graphs where the space of feasible input graphs contains all complete bipartite graphs. Note that this is true for all the problems we explicitly defined herein. For the statement of the theorem, we assume that for every (feasible) graph G a matching extending k -family $\mathcal{F}(G)$ is specified.

THEOREM 4.2. *Consider any problem captured by Ordinal- k -Max-on-Graphs for which any complete bipartite graph is a feasible input and its feasible solutions are given by $\mathcal{F}(\cdot)$. No deterministic mechanism using at most $\lambda \geq 1$ queries per agent can achieve a distortion better than $\Omega(n^{1/\lambda}/k^3)$ for the problem.*

Proof. We are going to show that if we had a deterministic mechanism \mathcal{M} for Ordinal- k -Max-on-Graphs that makes at most $\lambda \geq 1$ queries per agent and achieves distortion $o(n^{1/\lambda})$, then we could design a deterministic mechanism for One-Sided Matching that also makes at most λ queries per agent and has distortion $o(n^{1/\lambda})$. As the latter is impossible [6], that would imply that the lower bound applies to Ordinal- k -Max-on-Graphs as well.

Let $(\mathcal{N}, \mathcal{A}, \succ_{\mathbf{v}})$ be an arbitrary instance of One-Sided Matching with $|\mathcal{N}| = |\mathcal{A}| = n$ and underlying weights defined by $\mathbf{v} = (v_{i,j})_{i \in \mathcal{N}, j \in \mathcal{A}}$. We essentially use the same instance for Ordinal- k -Max-on-Graphs: A complete bipartite weighted graph $G = (U_1 \cup U_2, E, w)$ with $U_1 = \mathcal{N}$, $U_2 = \mathcal{A}$, and valuation functions defined as $u_i(j) = v_{i,j}$ and $u_j(i) = 0$ for every $i \in \mathcal{N}, j \in \mathcal{A}$; the induced ordinal profile is also well-defined. Clearly, the matchings in the two instances are exactly the same and have the same weight (although they may not be feasible with respect to $\mathcal{F}(G)$). However, the feasible solutions for Ordinal- k -Max-on-Graphs include subgraphs where the nodes may have degree up to k instead of 1. We need to relate the weight of an (approximately) optimal solution for the Ordinal- k -Max-on-Graphs instance to the value of an optimal matching for the One-Sided Matching instance.

Let M be a maximum weight matching in G (and thus a maximum-value matching for the original One-Sided Matching instance) and $H \in \mathcal{F}(G)$ be an optimal solution. Consider the submatching \hat{M} of M that uses the $\lfloor 2n/3k \rfloor$ heaviest edges of M . Using the fact that $\lfloor x \rfloor \geq x/2$ for $x \geq 1$, we get

$$(4.1) \quad w(\hat{M}) \geq \frac{\lfloor 2n/3k \rfloor}{n} w(M) \geq \frac{1}{3k} w(M).$$

Since $\mathcal{F}(G)$ is a matching extending k -family and \hat{M} is sufficiently small (since $|V(G)| = 2n$), there is some $\hat{H} \in \mathcal{F}(G)$ such that \hat{M} is a subgraph of \hat{H} . As H is a maximum-weight element of $\mathcal{F}(G)$, we directly get $w(H) \geq w(\hat{H}) \geq w(\hat{M})$ and, combining with (4.1), we have

$$(4.2) \quad w(H) \geq \frac{1}{3k} w(M).$$

Now, if $H' \in \mathcal{F}(G)$ is an α -approximate solution to the same Ordinal- k -Max-on-Graphs instance, then (4.2) implies

$$(4.3) \quad w(H') \geq \frac{1}{\alpha} w(H) \geq \frac{1}{3\alpha k} w(M).$$

We can construct a matching from H' using only ordinal information. In particular, for each $i \in U_1$, among the edges in H' that are incident to i , we keep the best one with respect to \succ_i . Of course, the resulting graph H'' may not be a matching, as each node in U_2 may still have degree up to k . However, note that this process also keeps at least a $1/k$ fraction of the weight incident to each $i \in U_1$, and thus of the total weight. So, (4.3) implies

$$w(H'') \geq \frac{1}{k} w(H') \geq \frac{1}{3\alpha k^2} w(M).$$

We repeat the process for the remaining nodes: for each $j \in U_2$, we keep the best of its edges in H'' with respect to \succ_j . Now the resulting graph M' is a matching and has at least a $1/k$ fraction of the total weight of H'' and, thus,

$$w(M') \geq \frac{1}{k} w(H'') \geq \frac{1}{3\alpha k^3} w(M).$$

If needed, we can extend M' to a perfect matching M'' by arbitrarily matching the unmatched nodes of U_1 and U_2 , and consider its analog in the original instance. Clearly,

$w(M'') \geq w(M')$, and thus M'' is a $(3\alpha k^3)$ -approximate solution for the original One-Sided Matching instance. Therefore, if there existed a mechanism \mathcal{M} with distortion $\alpha = o(n^{1/\lambda}/k^3)$ for Ordinal- k -Max-on-Graphs with feasible solutions given by \mathcal{F} , we could use it for the above instance to get H' and then M'' , which would have weight within a factor of $o(3n^{1/\lambda})$ from a maximum weight matching. This would imply a distortion of $o(n^{1/\lambda})$ for One-Sided Matching, a contradiction. \square

We are particularly interested in the case of $\lambda = 2$. In the next two sections we are going to present a mechanism for this case, which is asymptotically optimal for constant k , namely it achieves distortion $O(k^2\sqrt{n})$, matching the lower bound we just derived (up to k^5).

4.1. Sufficiently Representative Assignments. We now revisit the notion of a *sufficiently representative assignment*. We will appropriately adjust it to refer to a single set of agents (which, for the case of One-Sided Matching, includes both the actual agents and the items), and also incorporates the parameter k from the definition of Ordinal- k -Max-on-Graphs.

DEFINITION 4.3. *Given $\mathcal{N}_1, \mathcal{N}_2 \subseteq \mathcal{N}$ and $k \in \mathbb{N}$, a many-to-one assignment A of agents in \mathcal{N}_1 to agents in \mathcal{N}_2 is an $(\mathcal{N}_1, \mathcal{N}_2, k)$ -sufficiently representative assignment if:*

(a) *For every agent $j \in \mathcal{N}_2$, $|\{i \in \mathcal{N}_1 : A(i) = j\}| \leq \sqrt{n}$, i.e., there are at most \sqrt{n} agents from \mathcal{N}_1 assigned to j ;*

(b) *For any bipartite graph H with node set $\mathcal{N}_1 \cup \mathcal{N}_2$ and maximum degree k , $|\{i \in \mathcal{N}_1 : \exists j \in \mathcal{N}_2 \text{ such that } \{i, j\} \in E(H) \text{ and } j \succ_i A(i)\}| \leq k\sqrt{n}$, i.e., there are at most $k\sqrt{n}$ agents in \mathcal{N}_1 that prefer some of their neighbors in H to the agent they are assigned to according to A .*

For One-Sided Matching and k -Constrained Resource Allocation, \mathcal{N}_1 is the set of actual agents and \mathcal{N}_2 is the set of items. In contrast, for all other problems considered here we have $\mathcal{N}_1 = \mathcal{N}_2 = \mathcal{N}$.

Similarly to what we did in Section 3, we need to show that an $(\mathcal{N}_1, \mathcal{N}_2, k)$ -sufficiently representative assignment exists for any instance of Ordinal- k -Max-on-Graphs and for any $\mathcal{N}_1, \mathcal{N}_2 \subseteq \mathcal{N}$. We rely on the same high-level idea for the construction: \sqrt{n} copies of each agent in \mathcal{N}_2 are created and then a SERIAL DICTATORSHIP algorithm is run with respect to the agents in \mathcal{N}_1 . The running time of \sqrt{n} -SERIAL DICTATORSHIP remains $O(n^{1.5})$.

Mechanism 4.1 \sqrt{n} -SERIAL DICTATORSHIP($\mathcal{N}_1, \mathcal{N}_2, \succ_v$)

- 1: Let \mathcal{B} be a multiset with \sqrt{n} copies of each $j \in \mathcal{N}_2$
 - 2: **for** every agent $i \in \mathcal{N}_1$ **do**
 - 3: Let α_i be i 's most preferred agent in \mathcal{B} w.r.t. \succ_i
 - 4: Remove α_i from \mathcal{B}
 - 5: **end for**
 - 6: **return** $A = (\alpha_i)_{i \in \mathcal{N}_1}$
-

The following theorem is the analog of Theorem 3.2. While the proof is very similar, the counting argument here is somewhat less intuitive compared to the case of One-Sided Matching due to the differences between Definitions 3.1 and 4.3.

THEOREM 4.4. *The assignment computed by \sqrt{n} -SERIAL DICTATORSHIP is an $(\mathcal{N}_1, \mathcal{N}_2, k)$ -sufficiently representative assignment.*

Proof. Let A be the assignment produced by the algorithm. During the execution of the algorithm, whenever all the copies of an agent $j \in \mathcal{N}_2$ have been matched, we will say that j is *exhausted*. Assume towards a contradiction that A is not an $(\mathcal{N}_1, \mathcal{N}_2, k)$ -sufficiently representative assignment. By construction, the first condition of Definition 4.3 is obviously

satisfied. So, there must be a graph H , as described in the second condition of Definition 4.3, with respect to which there exists a subset $\mathcal{S}_1 \subseteq \mathcal{N}_1$ with $|\mathcal{S}_1| > k\sqrt{n}$ such that every $i \in \mathcal{S}_1$ prefers her best neighbor in H , say β_i , to agent α_i she has been assigned to in A . Let $\mathcal{S}_2 \subseteq \mathcal{N}_2$ be the set that contains all these β_i s. Because H has maximum degree at most k we have $|\mathcal{S}_2| \geq |\mathcal{S}_1|/k > \sqrt{n}$.

Consider any agent $i \in \mathcal{S}_1$. The fact that this agent was not assigned to β_i by \sqrt{n} -SERIAL DICTATORSHIP implies that when it was i 's turn to pick, agent β_i was exhausted. Therefore, at the end of the algorithm, all agents of \mathcal{S}_2 must be exhausted. Since an agent in \mathcal{N}_2 is exhausted when all its \sqrt{n} copies have been assigned and there are at most n agents in \mathcal{N}_1 , we can only have as many as $\frac{n}{\sqrt{n}} = \sqrt{n}$ exhausted agents. The latter means that $|\mathcal{S}_2| \leq \sqrt{n}$, a contradiction. \square

4.2. The General Mechanism. We are now ready to show that it is possible to achieve distortion $O(k^2\sqrt{n})$ for any problem that can be modeled as a special case of Ordinal- k -Max-on-Graphs. Our mechanism generalizes the main idea of An Optimal Two-Query Mechanism of querying each agent about for her overall favorite alternative, as well as the alternative suggested by an appropriate sufficiently representative assignment. For the latter, we need to specify \mathcal{N}_1 and \mathcal{N}_2 : These are typically both equal to the whole \mathcal{N} , unless the problem distinguishes between actual agents and items, in which case these two groups are captured by \mathcal{N}_1 and \mathcal{N}_2 , respectively. In any case, all the edges in a feasible solution have at least one endpoint in each of \mathcal{N}_1 and \mathcal{N}_2 .

Mechanism 4.2 GENERAL-TWOQUERIES($G, \mathcal{F}, \mathcal{N}_1, \mathcal{N}_2, \succ_{\mathbf{v}}$)

- 1: Query each agent $i \in \mathcal{N}_1$ for her favorite alternative in \mathcal{N}_2 w.r.t. \succ_i
 - 2: Compute $A = \sqrt{n}$ -SERIAL DICTATORSHIP($\mathcal{N}_1, \mathcal{N}_2, \succ_{\mathbf{v}}$)
 - 3: Query each agent $i \in \mathcal{N}_1$ about the agent $\alpha_i \in \mathcal{N}_2$ she is assigned to in A
 - 4: Set all non-revealed values to 0
 - 5: **return** a maximum-weight member of \mathcal{F}
-

Note that the final step of the algorithm involves computing a solution that is optimal according to the revealed values. There are computational issues to consider here. However, we first tackle the question of whether it is even possible to match the lower bounds of Theorem 4.2 for $\lambda = 2$, despite the lack of information. We briefly discuss how to transform GENERAL-TWOQUERIES into a polynomial-time mechanism after the proof of Theorem 4.5 below. Again, for the statement of the theorem we assume that $k \in \mathbb{N}$ is a constant and that, for every G , a matching extending k -family $\mathcal{F}(G)$ is specified.

THEOREM 4.5. *For Ordinal- k -Max-on-Graphs with feasible solutions given by $\mathcal{F}(\cdot)$, GENERAL-TWOQUERIES has distortion $O(k^2\sqrt{n})$.*

Proof. Consider any instance with valuation profile \mathbf{v} and relevant sets of agents \mathcal{N}_1 and \mathcal{N}_2 . Let Y be the solution computed by the GENERAL-TWOQUERIES mechanism when given as input $(G, \mathcal{F}, \mathcal{N}_1, \mathcal{N}_2, \succ_{\mathbf{v}})$, and let X denote an optimal solution. Let $w_R(Y)$ be the *revealed* weight of Y as seen by the mechanism, that is, the weight of Y taking into account only the values that have been revealed by the queries. We will show that $w(X) \leq (1 + 10k^2\sqrt{n}) \cdot w_R(Y)$, and the bound on the distortion will then follow by the obvious fact that $w(Y) \geq w_R(Y)$.

We can write the optimal weight as

$$(4.4) \quad w(X) = w_R(X) + w_C(X),$$

where $w_R(X)$ is the revealed weight of X that takes into account only the values that have been revealed by the queries, whereas $w_C(X)$ is the *concealed* weight of X that takes into account only the values that have not been revealed by the queries of the mechanism. Since Y is the solution that maximizes the social welfare based only on the revealed values, we have that

$$(4.5) \quad w_R(X) \leq w_R(Y).$$

Thus, it suffices to bound $w_C(X)$.

Let S be the set of agents in \mathcal{N}_1 who are not queried about all their neighbors in X . We partition S into two subsets S^\geq and $S^<$ consisting of agents for whom the second query of the mechanism is used to ask about someone they consider *better* or *worse* than their best neighbor in X , respectively. For an agent $i \in \mathcal{N}_1$, let χ_i be i 's favorite neighbor in X and recall that i is queried about agent $\alpha_i \in \mathcal{N}_2$ to whom she is assigned according to the $(\mathcal{N}_1, \mathcal{N}_2, k)$ -sufficiently representative assignment A . So,

$$\begin{aligned} S^\geq &= \{i \in S : v_{i, \alpha_i} \geq v_{i, \chi_i}\}, \\ S^< &= \{i \in S : v_{i, \alpha_i} < v_{i, \chi_i}\}. \end{aligned}$$

Let $N_X(i)$ be the set of agents who are neighbors of i in X and for which i was not queried about. We now define

$$\begin{aligned} w_C^\geq(X) &= \sum_{i \in S^\geq} \sum_{j \in N_X(i)} v_{i,j} \\ w_C^<(X) &= \sum_{i \in S^<} \sum_{j \in N_X(i)} v_{i,j} \end{aligned}$$

Clearly, $w_C(X) = w_C^\geq(X) + w_C^<(X)$.

For every agent $j \in \mathcal{N}_2$, let $S_j^\geq = \{i \in S^\geq : \alpha_i = j\}$ be the set of all agents in S^\geq that are queried about j by the mechanism using the second query. So, $S^\geq = \bigcup_{j \in \mathcal{N}_2} S_j^\geq$. Since A is an $(\mathcal{N}_1, \mathcal{N}_2, k)$ -sufficiently representative assignment, the first condition of Definition 4.3 implies that $|S_j^\geq| \leq \sqrt{n}$ for every $j \in \mathcal{N}_2$. Therefore,

$$\begin{aligned} w_C^\geq(X) &= \sum_{j \in \mathcal{N}_2} \sum_{i \in S_j^\geq} \sum_{\ell \in N_X(i)} v_{i,\ell} \\ &\leq \sum_{j \in \mathcal{N}_2} \sum_{i \in S_j^\geq} \sum_{\ell \in N_X(i)} v_{i,j} \\ &\leq \sum_{j \in \mathcal{N}_2} \sum_{i \in S_j^\geq} k \cdot v_{i,j} \\ &\leq k \sum_{j \in \mathcal{N}_2} |S_j^\geq| \max_{i \in S_j^\geq} v_{i,j} \\ (4.6) \quad &\leq k\sqrt{n} \sum_{j \in \mathcal{N}_2} \max_{i \in S_j^\geq} v_{i,j} \end{aligned}$$

where the first inequality holds by the definition of the sets S_j^\geq , for every $j \in \mathcal{N}_2$. To complete our bound on $w_C^\geq(X)$, we need the following claim.

CLAIM 1. *For all the problems of interest,*

$$\sum_{j \in \mathcal{N}_2} \max_{i \in \mathcal{S}_j^{\geq}} v_{i,j} \leq 9k \cdot w_R(Y).$$

Proof. For each $j \in \mathcal{N}_2$, let $i_j \in \arg \max_{i \in \mathcal{S}_j^{\geq}} v_{i,j}$. Consider the subgraph H of the input graph G with edge set $E(H) = \{\{i_j, j\} \mid j \in \mathcal{N}_2\}$, i.e., H contains exactly the edges that define the sum of interest. In particular, we have

$$\sum_{j \in \mathcal{N}_2} \max_{i \in \mathcal{S}_j^{\geq}} v_{i,j} = \sum_{j \in \mathcal{N}_2} v_{i_j, j} = w_R(H).$$

We now claim that each node in H has degree at most 2. To see this, consider an agent $\ell \in \mathcal{N}$. There is at most one $j \in \mathcal{N}_2$ such that $\ell \in \mathcal{S}_j^{\geq}$ (since these sets are disjoint), and thus we may have $\ell = i_j$ for at most one $j \in \mathcal{N}_2$, resulting in the edge $\{j, \ell\}$ in H . Additionally, ℓ may itself be in \mathcal{N}_2 , resulting in a second edge $\{\ell, i_\ell\}$ in H . Other than these two, there can be no other edges of H adjacent to ℓ .

Since H has maximum degree at most 2, it must contain a matching M of comparable weight. Specifically, H must be the union of node-disjoint paths and cycles. We construct a (possibly empty) matching M_1 on H by arbitrarily picking one edge from each odd cycle and one edge from the beginning of each odd path. If we remove M_1 from H , then the remaining graph is the union of node-disjoint even paths and even cycles, and thus can be decomposed into two disjoint matchings M_2, M_3 in a straightforward way. Since M_1, M_2 , and M_3 cover all the edges of H , the best of them, say M , must have weight at least $w_R(H)/3$, i.e.,

$$w_R(M|\mathbf{v}) = \frac{1}{3}w_R(H).$$

Now we can work with M like in the proof of Theorem 4.2. Consider the submatching \hat{M} of M containing the $\lfloor 2n/3k \rfloor$ heaviest edges of M to get

$$(4.7) \quad w_R(\hat{M}) \geq \frac{1}{3k}w_R(M) \geq \frac{1}{9k}w_R(H).$$

Since \mathcal{F} is a matching extending k -family and \hat{M} is sufficiently small, there is some $\hat{Y} \in \mathcal{F}(G)$ such that \hat{M} is a subgraph of \hat{Y} . As Y is a maximum-weight element of $\mathcal{F}(G)$ with respect to the revealed weights, we directly get $w_R(Y) \geq w_R(\hat{Y}) \geq w_R(\hat{M})$ and, combining with (4.7), we have

$$w_R(Y) \geq \frac{1}{9k}w(H),$$

as desired. \square

By combining (4.6) with Claim 1, we get

$$(4.8) \quad w_C^{\geq}(X) \leq 9k^2\sqrt{n} \cdot w_R(Y).$$

We next consider the quantity $w_C^{\leq}(X)$. By the fact that A is an $(\mathcal{N}_1, \mathcal{N}_2, k)$ -sufficiently representative assignment, it follows that $|S^{\leq}| \leq k\sqrt{n}$; otherwise X would be a graph that violates the second condition of Definition 4.3. Combined with the fact that all agents in \mathcal{N}_1 are queried about their favorite alternative, we can obtain the following upper bound on $w_C^{\leq}(X)$. Recall that for $i \in \mathcal{N}_1$, we have $N_X(i) \subseteq \mathcal{N}_2$ and $|N_X(i)| \leq k$.

$$w_C^{\leq}(X) = \sum_{i \in S^{\leq}} \sum_{j \in N_X(i)} v_{i,j}$$

$$\begin{aligned}
&\leq \sum_{i \in S^c} k \cdot \max_{j \in \mathcal{N}} v_{i,j} \\
&\leq k |S^c| \max_{i \in S^c} \max_{j \in \mathcal{N}} v_{i,j} \\
(4.9) \quad &\leq k^2 \sqrt{n} \cdot w_R(Y).
\end{aligned}$$

The bound now follows by (4.4), (4.5), (4.8), (4.9). \square

A subtle point here is that of computational efficiency. Although designing polynomial time mechanisms is not our primary goal, it is clear that the only possible bottleneck is the last step of GENERAL-TWOQUERIES. Indeed, the mechanism runs in polynomial time whenever there is a polynomial-time algorithm (exact or $O(1)$ -approximation) for the full information version of the corresponding optimization problem. The good news is that all variants of matching problems we presented can be solved efficiently by Edmond's algorithm [23] or its extensions [36]. For the next corollary, recall that $k = 1$ for General Graph Matching and Two-Sided Matching.

COROLLARY 4.6. *There exist deterministic polynomial-time mechanisms for General Graph Matching, Two-Sided Matching, General Graph k -Matching, and k -Constrained Resource Allocation which all use at most two queries per agent and have distortion $O(k^2 \sqrt{n})$.*

5. Towards Tight Bounds for General Social Choice. Here we consider the general social choice setting where a set \mathcal{N} of n agents have preferences over a set \mathcal{A} of m alternatives. As in the One-Sided Matching problem, there is a valuation profile $\mathbf{v} = (v_{i,j})_{i \in \mathcal{N}, j \in \mathcal{A}}$ specifying the non-negative value that each agent i has for every alternative j . The goal is to choose a single alternative $x \in \mathcal{A}$ to maximize the social welfare, that is, the total value of the agents for x : $\text{SW}(x|\mathbf{v}) = \sum_{i \in \mathcal{N}} v_{i,x}$. Again, when \mathbf{v} is clear from context, we will drop it from notation. Similarly to One-Sided Matching and the problems discussed in the previous section, \mathbf{v} is unknown, and we are only given access to the ordinal profile $\succ_{\mathbf{v}}$ that is induced by \mathbf{v} . Social choice mechanisms must decide a single alternative based only on $\succ_{\mathbf{v}}$ and the values they can learn by making a small number of queries. The notion of distortion (Definition 2.1) can be extended for this setting as well, by taking the supremum over all instances with n agents and m alternatives, and letting \mathcal{X} be the set \mathcal{A} of alternatives.

For this general social choice setting, Amanatidis et al. [4] showed a lower bound of $\Omega(m^{1/(2(\lambda+1))})$ on the distortion of mechanisms that make at most $\lambda \geq 1$ queries per agent. We improve this result by showing a lower bound of $\Omega(m^{1/\lambda})$ for any constant λ .

THEOREM 5.1. *In the social choice setting, the distortion of any deterministic mechanism that makes at most a constant number $\lambda \geq 1$ of queries per agent is $\Omega(m^{1/\lambda})$.*

Proof. Let \mathcal{M} be an arbitrary mechanism that makes at most $\lambda \geq 1$ queries per agent. Consider the following instance with n agents and $m = n$ alternatives. We assume that m is superconstant (otherwise the theorem holds trivially) and satisfies the condition $m \geq \frac{1}{2} \sum_{\ell=1}^{\lambda} m^{(\lambda-\ell+1)/\lambda} + 2$. We partition the set of alternatives \mathcal{A} into $\lambda + 2$ sets $A_1, A_2, \dots, A_{\lambda+1}, A_{\lambda+2}$, such that

- $|A_\ell| = \frac{1}{2} m^{(\lambda-\ell+1)/\lambda}$ for $\ell \in [\lambda]$;
- $|A_{\lambda+1}| = 2$;
- $|A_{\lambda+2}| = m - \frac{1}{2} \sum_{\ell=1}^{\lambda} m^{(\lambda-\ell+1)/\lambda} - 2$.

The ordinal profile has the following properties:

- For every $\ell \in [\lambda + 1]$, each alternative $j \in A_\ell$ is ranked at position ℓ by a set $T_{j,\ell}$ of $\frac{m}{|A_\ell|} = \Theta(m^{(\ell-1)/\lambda})$ agents.
- For every $\ell \in [\lambda]$, every pair of agents that rank the same alternative in A_ℓ at position ℓ , rank the same alternative in $A_{\ell+1}$ at position $\ell + 1$.

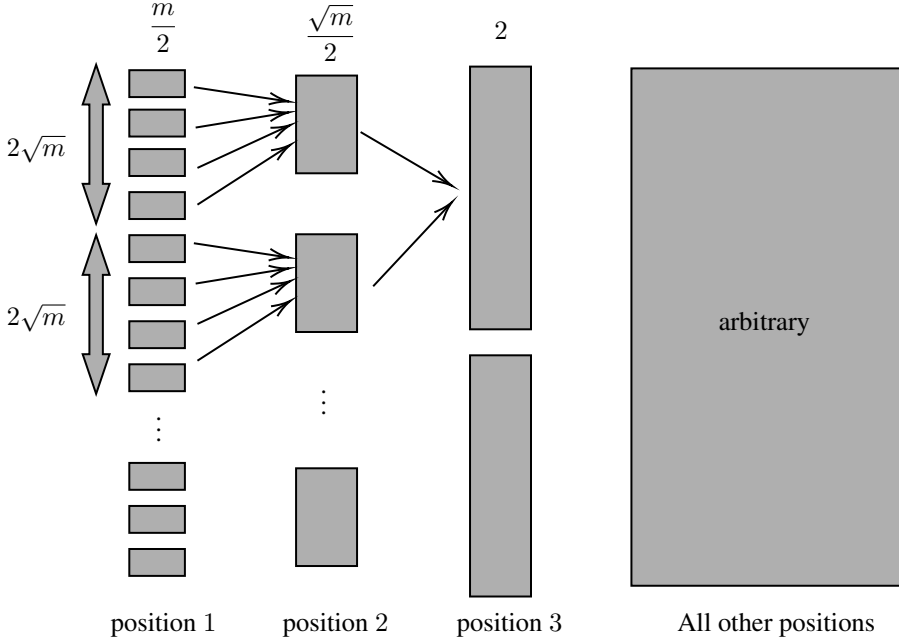


FIG. 1. An overview of the instance used in the proof of Theorem 5.1 two queries ($\lambda = 2$). Each rectangle in the first three positions corresponds to an alternative. Each rectangle at position 1 contains two agents. Each rectangle at position 2 contains the agents from $2\sqrt{m}$ rectangles at position 1, as indicated by the arrows, meaning that those agents rank the same alternative second. The rectangles at position 3 contain $m/2$ agents each, corresponding to $\sqrt{m}/4$ rectangles at position 2. That is, the agents that rank second one of the first $\sqrt{m}/4$ alternatives at position 2, rank third the alternative corresponding to the first rectangle at position 3; similarly, the agents that rank second one of the last $\sqrt{m}/4$ alternatives at position 2, rank third the alternative corresponding to the second rectangle at position 3. The ranking of the alternatives in the remaining positions is consistent but otherwise arbitrary.

- For every agent, the alternatives that she does not rank in the first $\lambda + 1$ positions are ranked arbitrarily from position $\lambda + 2$ to m .

An example of the ordinal profile when $\lambda = 2$ is depicted in Figure 1. For every agent i , a query of \mathcal{M} for alternative j reveals a value of

- $m^{-\ell/\lambda}$ if i ranks j at position $\ell \in [\lambda + 1]$, and
- and a value of 0 if i ranks j at any other position.

Given this instance as input, let y be the alternative that \mathcal{M} chooses as the winner. No matter the choice of y , we will define the cardinal profile so that it is consistent to the information revealed by the queries of \mathcal{M} , and the values of the agents for alternative y are also consistent to the information that would have been revealed, irrespective of whether those values have actually been revealed. That is, any agent has a value of $m^{-\ell/\lambda}$ for y if she ranks y at position $\ell \in [\lambda + 1]$, and a value of 0 if she ranks y at any other position. Hence, the social welfare of y is

- $\Theta(m^{(\ell-1)/\lambda}) \cdot m^{-\ell/\lambda} = \Theta(m^{-1/\lambda})$ if $y \in A_\ell$ for $\ell \in [\lambda + 1]$, or
- 0 if $y \in A_{\lambda+2}$.

Consequently, to show the desired bound of $\Omega(m^{1/\lambda})$ on the distortion of \mathcal{M} , it suffices to assume that $y \in A_\ell$ for some $\ell \in [\lambda + 1]$, and prove that the values of the agents that have not been revealed and do not correspond to alternative y can always be defined such that there exists an alternative $x \neq y$ with social welfare $\Omega(1)$.

Suppose towards a contradiction that the cardinal profile cannot be defined in a way

so that there exists an alternative x with social welfare $\Omega(1)$. We make the following two observations:

- (O1) If there exists an alternative $x \in A_1 \setminus \{y\}$ for which at least one agent in $T_{x,1}$ is *not* queried by \mathcal{M} for x , then we can set the value of this agent for x to be constant. Consequently, all the agents in $\bigcup_{j \in A_1 \setminus \{y\}} T_{j,1}$ that rank alternatives different that y at position 1, must be queried at position 1.
- (O2) Let $\varepsilon \in (0, 1)$ be a constant and $\ell \in \{2, \dots, \lambda + 1\}$. Consider any alternative $x \in A_\ell \setminus \{y\}$ and any set of agents $S \subseteq T_{x,\ell}$ such that $|S| \geq \varepsilon \cdot \frac{m}{|A_\ell|} = \Theta(m^{(\ell-1)/\lambda})$. If at least $\frac{1}{2}|S|$ agents in S are *not* queried by \mathcal{M} for x , then we could set the value of all these agents for x to be $m^{-(\ell-1)/\lambda}$ (which is the revealed value when \mathcal{M} queries for alternatives ranked at position $\ell - 1$), and the social welfare of x would be at least

$$\frac{1}{2}|S| \cdot m^{-(\ell-1)/\lambda} = \Theta(m^{(\ell-1)/\lambda}) \cdot m^{-(\ell-1)/\lambda} = \Theta(1).$$

Consequently, for every alternative $x \in A_\ell \setminus \{y\}$ and set $S \subseteq T_{x,\ell}$ such that $|S| \geq \varepsilon \cdot \frac{m}{|A_\ell|}$, at least $\frac{1}{2}|S|$ agents in S must be queried at position ℓ for x .

Given these two observations, we are now ready to show by induction that the mechanism must make $\lambda + 1$ queries for a high proportion of the agents, contradicting that \mathcal{M} makes at most λ queries per agent.

For the base case, consider an alternative $x \in A_2 \setminus \{y\}$. By the definition of the ordinal profile, the agents in $T_{x,2}$ who rank x at position 2 are partitioned into $\frac{|A_1|}{|A_2|}$ subsets such that all $\frac{m}{|A_1|}$ agents in each subset rank first the same alternative of A_1 . By (O1) we have that, besides the agents that rank alternative y at position 1, all other agents must be queried at position 1. Hence, there exists a set $S \subseteq T_{x,2}$ consisting of $|S| \geq \left(\frac{|A_1|}{|A_2|} - 1\right) \cdot \frac{m}{|A_1|}$ agents that are queried at position 1. By the definitions of A_1 and A_2 , and since m is superconstant, we have that $|S| \geq \frac{1}{2} \cdot \frac{m}{|A_2|}$. By (O2) for $\varepsilon = \frac{1}{2}$ and $\ell = 2$, we have that at least $\frac{1}{2}|S| \geq \frac{1}{4} \cdot \frac{m}{|A_2|}$ of the agents in S must also be queried at position 2 for x .

Let $\ell \in \{3, \dots, \lambda + 1\}$ and assume as induction hypothesis that for every alternative $z \in A_{\ell-1} \setminus \{y\}$ there is a set of agents $S_z \subseteq T_{z,\ell-1}$ such that $|S_z| \geq \frac{1}{2^{2(\ell-2)}} \cdot \frac{m}{|A_{\ell-1}|}$ who are queried by \mathcal{M} at the first $\ell - 1$ positions. Consider an alternative $x \in A_\ell \setminus \{y\}$. By the definition of the ordinal profile, the agents in $T_{x,\ell}$ who rank alternative x at position ℓ are partitioned into $\frac{|A_{\ell-1}|}{|A_\ell|}$ subsets such that all $\frac{m}{|A_{\ell-1}|}$ agents in each subset rank the same alternative in $A_{\ell-1}$ at position $\ell - 1$. So, by our induction hypothesis, there is a set $S \subseteq T_{x,\ell}$ consisting of

$$|S| \geq \left(\frac{|A_{\ell-1}|}{|A_\ell|} - 1\right) \cdot \frac{1}{2^{2(\ell-2)}} \cdot \frac{m}{|A_{\ell-1}|}$$

agents that are queried at the first $\ell - 1$ positions. By the definition of $A_{\ell-1}$ and A_ℓ , and since m is superconstant, we have that

$$|S| \geq \frac{1}{2^{2(\ell-2)+1}} \cdot \frac{m}{|A_\ell|}.$$

Since $\ell \leq \lambda + 1$ and λ is a constant, by observation (O2) for $\varepsilon = \frac{1}{2^{2(\ell-2)+1}}$, we have that at least

$$\frac{1}{2}|S| \geq \frac{1}{2^{2(\ell-1)}} \cdot \frac{m}{|A_\ell|}$$

agents in S must also be queried at position ℓ for x .

Now, let $x \in A_{k+1} \setminus \{y\}$. The above induction shows that there are at least $\frac{1}{2^{2\lambda}} \cdot \frac{m}{|A_{\lambda+1}|}$ agents in $T_{x,\lambda+1}$ who must be queried by \mathcal{M} at the first $\lambda + 1$ positions. This contradicts the fact that \mathcal{M} can make at most λ queries per agent, and the theorem follows. \square

Our approach for all the problems discussed in the previous sections can also be applied to the much more general social choice setting, *subject to* being able to compute a particular set of alternatives.

DEFINITION 5.2. *Let $c \geq 1$ be any constant. A subset of alternatives $B \subseteq \mathcal{A}$ with $|B| \leq c \cdot \sqrt{m}$ is a sufficiently representative set if, for every alternative $j \in \mathcal{A}$, at most \sqrt{m} agents prefer j over their favorite alternative in B .*

An interesting question is whether a sufficiently representative set of alternatives always exists. Unfortunately, it is not difficult to construct examples when this is not the case, when m is much smaller than n . For example, for any $k > \sqrt{m}$, consider an instance with $n = k \cdot m!$ agents, such that for each possible ordering of the m alternatives there are exactly k agents that have it as their preference. Then, for any subset B of at most \sqrt{m} alternatives and any alternative $j \in \mathcal{A} \setminus B$, there are at least $k > \sqrt{m}$ agents that prefer j over any alternative in B .

On the other hand, a sufficiently representative set of alternatives trivially exists when m is much larger than n (namely, when $m = \Omega(n^2)$). We can actually obtain something stronger than this, using the following result of Jiang et al. [30].

THEOREM 5.3 ([30]). *For any $\xi \in [n]$, there exists a set S of alternatives with $|S| \leq 16 \cdot n/\xi$ such that for every $j \in \mathcal{A}$, there are at most ξ agents that prefer j over their favorite alternative in S .*

A set S as in the theorem above is called an *approximately stable committee* [21, 30]. Clearly, when $m = \Omega(n)$ and $\xi = \sqrt{n}$, an approximately stable committee is also a sufficiently representative set with $c = 16$.

We now present a mechanism that works *under the assumption* that sufficiently representative sets of alternatives can be (efficiently) computed. Given the discussion above, we state an unconditional bound as a corollary of Theorem 5.4 at the end of the section.

Mechanism 5.1 SC-TWOQUERIES($\mathcal{N}, \mathcal{A}, \succ_{\mathbf{v}}$)

- 1: Query each agent about her favorite alternative
 - 2: Compute a sufficiently representative set B
 - 3: Query each agent for her favorite alternative in B
 - 4: For every $j \in \mathcal{A}$, compute the revealed welfare $\text{SW}_R(j)$
 - 5: **return** $y \in \arg \max_{j \in \mathcal{A}} \text{SW}_R(j)$
-

In particular, SC-TWOQUERIES (Mechanism 5.1) first queries each agent about her overall favorite alternative (the one ranked first). Then, given a sufficiently representative set of alternatives B , it queries each agent for her favorite alternative in B . Given the answers to these two queries per agent, the mechanism outputs an alternative that maximizes the *revealed* social welfare which is based only on the values learned from the queries.

THEOREM 5.4. *The mechanism SC-TWOQUERIES has distortion $O(\sqrt{m})$, when restricted to the social choice instances for which a sufficiently representative set of alternatives exists.*

Proof. Consider any social choice instance with valuation profile \mathbf{v} that induces the ordinal preference profile $\succ_{\mathbf{v}}$. Let y be the alternative chosen by the mechanism when given as input this instance, and denote by x the optimal alternative. We will show that

$\text{SW}(x) \leq (1 + (1 + c) \cdot \sqrt{m})\text{SW}_R(y)$. The bound on the distortion will then follow by the obvious fact that $\text{SW}(y) \geq \text{SW}_R(y)$.

We can write the optimal welfare as

$$(5.1) \quad \begin{aligned} \text{SW}(x) &= \text{SW}_R(x) + \text{SW}_C(x) \\ &\leq \text{SW}_R(y) + \text{SW}_C(x), \end{aligned}$$

where $\text{SW}_C(x)$ is the concealed welfare of x , consisting of the values of agents for x that were not revealed by the queries of the mechanism, and the inequality follows by the fact that y is the alternative that maximizes the revealed welfare. Let S be the set of agents who were *not* queried about their value for x , and partition S into the following two subsets:

- S^{\geq} consists of the agents in S for whom the second query is about an alternative that the agent considers *better* than x ;
- $S^{<}$ consists of the agents in S for whom the second query is about an alternative that the agent considers *worse* than x .

Given these sets, now let

$$\text{SW}_C^{\geq}(x) = \sum_{i \in S^{\geq}} v_{i,x} \quad \text{and} \quad \text{SW}_C^{<}(x) = \sum_{i \in S^{<}} v_{i,x}.$$

be the contribution of the agents in S^{\geq} and of the agents in $S^{<}$ to the concealed welfare of x , respectively. That is,

$$\text{SW}_C(x) = \text{SW}_C^{\geq}(x) + \text{SW}_C^{<}(x).$$

By the definition of the mechanism, each agent is queried about her favorite alternative in the sufficiently representative set B . For every $j \in B \setminus \{x\}$, let $S_j^{\geq} \subseteq S^{\geq}$ be the set of agents in S^{\geq} who are queried for alternative j instead of x . Thus, $S^{\geq} = \bigcup_{j \in B \setminus \{x\}} S_j^{\geq}$. By the definition of S^{\geq} , the fact that y maximizes the revealed welfare, and since $|B| \leq c \cdot \sqrt{m}$, we obtain

$$(5.2) \quad \begin{aligned} \text{SW}_C^{\geq}(x) &= \sum_{j \in B \setminus \{x\}} \sum_{i \in S_j^{\geq}} v_{i,x} \\ &\leq \sum_{j \in B \setminus \{x\}} \sum_{i \in S_j^{\geq}} v_{i,j} \\ &\leq \sum_{j \in B \setminus \{x\}} \text{SW}_R(j|\mathbf{v}) \\ &\leq |B| \cdot \text{SW}_R(y|\mathbf{v}) \\ &\leq c \cdot \sqrt{m} \cdot \text{SW}_R(y). \end{aligned}$$

Since all the agents in $S^{<}$ are queried for alternatives in the sufficiently representative set B that they consider worse than x and B , it must be the case that $|S^{<}| \leq \sqrt{m}$. Since all agents are queried at the first position for their favorite alternative, we obtain

$$\begin{aligned} \text{SW}_C^{<}(x) &= \sum_{i \in S^{<}} v_{i,x} \\ &\leq \sum_{i \in S^{<}} \max_{j \in \mathcal{A}} v_{i,j} \\ &\leq |S^{<}| \cdot \max_{i \in S^{<}} \max_{j \in \mathcal{A}} v_{i,j} \end{aligned}$$

$$(5.3) \quad \leq \sqrt{m} \cdot \text{SW}_R(y).$$

The bound now follows by (5.1), (5.2) and (5.3). \square

Finally, combining Theorems 5.3 and 5.4, we obtain the following result.

COROLLARY 5.5. *When $m = \Omega(n)$, SC-TWOQUERIES has distortion $O(\sqrt{m})$.*

6. Conclusion and Open Problems. In this paper, we showed that for a large class of problems, which includes One-Sided Matching and many other well-studied graph-theoretic problems, it is possible to achieve a distortion of $O(\sqrt{n})$ using a deterministic mechanism that makes at most two queries per agent, and that this is best possible asymptotically. Our whole methodology is based on computing assignments of agents to items or other agents that exhibit a very particular structure. In addition, in the social choice setting, when $m = \Omega(n)$, sets of alternatives with analogous properties can be computed, and our methodology yields a two-query mechanism with best possible distortion for this setting as well.

A question that is left unresolved is the best possible distortion for the problems with solutions bounded by a parameter k that is superconstant in general, such as k -Constrained Resource Allocation, General Graph k -Matching, and k -Clique Packing. We suspect that approaches which are problem-specific will be required to make progress in this direction. Of course, an interesting open problem is to design a mechanism for the social choice setting that makes two queries and achieves the best possible distortion of $O(\sqrt{m})$ when $m = o(n)$, or show that this is impossible. We believe that to obtain a positive result one would need to come up with an adaptive mechanism, which decides where to ask each query based not only on the ordinal information, but also on the answers to all previous ones. Another question, about any of the problems we considered, is whether one can design mechanisms that make at most a constant $\lambda \geq 3$ queries per agent and their distortion matches the lower bound of $\Omega(n^{1/\lambda})$ (or, in the case of social choice, $\Omega(m^{1/\lambda})$). Again, we strongly suspect that the same type of adaptivity will be required for this task as well.

REFERENCES

- [1] D. J. ABRAHAM, A. BLUM, AND T. SANDHOLM, *Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges*, in Proceedings of the 8th ACM Conference on Electronic commerce (EC), 2007, pp. 295–304.
- [2] B. ABRAMOWITZ AND E. 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), 2018, pp. 894–901.
- [3] B. ABRAMOWITZ, E. ANSHELEVICH, AND W. ZHU, *Awareness of voter passion greatly improves the distortion of metric social choice*, in Proceedings of the The 15th Conference on Web and Internet Economics (WINE), 2019, pp. 3–16.
- [4] G. AMANATIDIS, G. BIRMPAS, A. FILOS-RATSIKAS, AND A. A. VOUDOURIS, *Peeking behind the ordinal curtain: Improving distortion via cardinal queries*, Artificial Intelligence, 296 (2021), p. 103488.
- [5] G. AMANATIDIS, G. BIRMPAS, A. FILOS-RATSIKAS, AND A. A. VOUDOURIS, *Don't roll the dice, ask twice: The two-query distortion of matching problems and beyond*, in Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems (NeurIPS), 2022.
- [6] G. AMANATIDIS, G. BIRMPAS, A. FILOS-RATSIKAS, AND A. A. VOUDOURIS, *A few queries go a long way: Information-distortion tradeoffs in matching*, Journal of Artificial Intelligence Research, 74 (2022).
- [7] E. ANSHELEVICH, O. BHARDWAJ, E. ELKIND, J. POSTL, AND P. SKOWRON, *Approximating optimal social choice under metric preferences*, Artificial Intelligence, 264 (2018), pp. 27–51.
- [8] E. ANSHELEVICH, A. FILOS-RATSIKAS, N. SHAH, AND A. A. VOUDOURIS, *Distortion in social choice problems: The first 15 years and beyond*, in Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), 2021, pp. 4294–4301.
- [9] E. ANSHELEVICH, A. FILOS-RATSIKAS, AND A. A. VOUDOURIS, *The distortion of distributed metric social choice*, Artificial Intelligence, 308 (2022), p. 103713.
- [10] E. ANSHELEVICH AND J. POSTL, *Randomized social choice functions under metric preferences*, Journal of Artificial Intelligence Research, 58 (2017), pp. 797–827.

- [11] E. ANSHELEVICH AND S. 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), 2016, pp. 390–396.
- [12] E. ANSHELEVICH AND W. 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), 2018, pp. 3–20.
- [13] E. ANSHELEVICH AND W. ZHU, *Tradeoffs between information and ordinal approximation for bipartite matching*, Theory of Computing Systems, 63 (2019), pp. 1499–1530.
- [14] H. AZIZ, E. ELKIND, P. FALISZEWSKI, M. LACKNER, AND P. SKOWRON, *The Condorcet principle for multiwinner elections: from shortlisting to proportionality*, in Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI), 2017, pp. 84–90.
- [15] G. BENADÈ, S. NATH, A. D. PROCACCIA, AND N. SHAH, *Preference elicitation for participatory budgeting*, Management Science, 67 (2021), pp. 2813–2827.
- [16] C. BOUTILIER, I. CARAGIANNIS, S. HABER, T. LU, A. D. PROCACCIA, AND O. SHEFFET, *Optimal social choice functions: A utilitarian view*, Artificial Intelligence, 227 (2015), pp. 190–213.
- [17] I. CARAGIANNIS, A. FILOS-RATSIKAS, S. K. S. FREDERIKSEN, K. A. HANSEN, AND Z. TAN, *Truthful facility assignment with resource augmentation: An exact analysis of serial dictatorship*, in International Conference on Web and Internet Economics (WINE), 2016, pp. 236–250.
- [18] I. CARAGIANNIS, S. NATH, A. D. PROCACCIA, AND N. SHAH, *Subset selection via implicit utilitarian voting*, Journal of Artificial Intelligence Research, 58 (2017), pp. 123–152.
- [19] I. CARAGIANNIS, N. SHAH, AND A. A. VOUDOURIS, *The metric distortion of multiwinner voting*, in Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI), 2022, pp. 4900–4907.
- [20] M. CHARIKAR AND P. RAMAKRISHNAN, *Metric distortion bounds for randomized social choice*, in Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA), 2022, pp. 2986–3004.
- [21] Y. CHENG, Z. JIANG, K. MUNAGALA, AND K. WANG, *Group fairness in committee selection*, ACM Transactions on Economics and Computation (TEAC), 8 (2020), pp. 1–18.
- [22] S. EBADIAN, A. KAHNG, D. PETERS, AND N. SHAH, *Optimized distortion and proportional fairness in voting*, in Proceedings of the 23rd ACM Conference on Economics and Computation (EC), 2022, pp. 523–600.
- [23] J. EDMONDS, *Maximum matching and a polyhedron with 0, 1-vertices*, Journal of research of the National Bureau of Standards B, 69 (1965), pp. 55–56.
- [24] A. FILOS-RATSIKAS, S. K. S. FREDERIKSEN, AND J. ZHANG, *Social welfare in one-sided matchings: Random priority and beyond*, in Proceedings of the 7th Symposium of Algorithmic Game Theory (SAGT), 2014, pp. 1–12.
- [25] A. FILOS-RATSIKAS, E. MICHA, AND A. A. VOUDOURIS, *The distortion of distributed voting*, Artificial Intelligence, 286 (2020), p. 103343.
- [26] A. FILOS-RATSIKAS AND P. B. MILTERSEN, *Truthful approximations to range voting*, in Proceedings of the 10th International Conference on Web and Internet Economics (WINE), 2014, pp. 175–188.
- [27] D. GALE AND L. S. SHAPLEY, *College admissions and the stability of marriage*, The American Mathematical Monthly, 69 (1962), pp. 9–15.
- [28] V. GKATZELIS, D. HALPERN, AND N. SHAH, *Resolving the optimal metric distortion conjecture*, in Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2020, pp. 1427–1438.
- [29] A. HYLLAND AND R. ZECKHAUSER, *The efficient allocation of individuals to positions*, Journal of Political economy, 87 (1979), pp. 293–314.
- [30] Z. JIANG, K. MUNAGALA, AND K. WANG, *Approximately stable committee selection*, in Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC), 2020, pp. 463–472.
- [31] F. E. KIZILKAYA AND D. KEMPE, *Plurality veto: A simple voting rule achieving optimal metric distortion*, in Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI), 2022, pp. 349–355.
- [32] H. W. KUHN, *Variants of the Hungarian method for assignment problems*, Naval Research Logistics Quarterly, 3 (1956), pp. 253–258.
- [33] T. MA, V. MENON, AND K. LARSON, *Improving welfare in one-sided matchings using simple threshold queries*, in Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), Z. Zhou, ed., 2021, pp. 321–327.
- [34] D. MANDAL, A. D. PROCACCIA, N. SHAH, AND D. P. WOODRUFF, *Efficient and thrifty voting by any means necessary*, in Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS), 2019, pp. 7178–7189.
- [35] D. MANDAL, N. SHAH, AND D. P. WOODRUFF, *Optimal communication-distortion tradeoff in voting*, in Proceedings of the 21st ACM Conference on Economics and Computation (EC), 2020, pp. 795–813.
- [36] A. B. MARSH III, *Matching algorithms.*, The Johns Hopkins University, 1979.
- [37] K. MUNAGALA AND K. WANG, *Improved metric distortion for deterministic social choice rules*, in Proceed-

- ings of the 2019 ACM Conference on Economics and Computation (EC), 2019, pp. 245–262.
- [38] A. D. PROCACCIA AND J. S. ROSENSCHEIN, *The distortion of cardinal preferences in voting*, in International Workshop on Cooperative Information Agents (CIA), 2006, pp. 317–331.
- [39] A. E. ROTH AND M. SOTOMAYOR, *Two-sided matching*, in Handbook of Game Theory with Economic Applications, vol. 1, Elsevier, 1992, ch. 16, pp. 485 – 541.