

Fair and Efficient Resource Allocation with Partial Information

Daniel Halpern¹, Nisarg Shah²

¹Harvard University

²University of Toronto

dhalpern@g.harvard.edu, nisarg@cs.toronto.edu

Abstract

We study the fundamental problem of allocating indivisible goods to agents with additive preferences. We consider eliciting from each agent only a ranking of her k most preferred goods instead of her full cardinal valuations. We characterize the value of k needed to achieve envy-freeness up to one good and approximate maximin share guarantee, two widely studied fairness notions. We also analyze the multiplicative loss in social welfare incurred due to the lack of full information with and without the fairness requirements.

1 Introduction

The theory of fair division studies how goods (or bads) should be fairly divided between individuals (a.k.a. agents) with different preferences over them. While the pioneering fair division research in economics, starting with the work of Steinhaus [1948], focused on *divisible* goods which can be split between the agents, a significant body of recent research within computer science has focused on allocation of *indivisible* goods [Bouveret *et al.*, 2016].

Suppose we wish to partition a set of indivisible goods M among a set of agents N . In doing so, we would like to take the agents' preferences into account; thus, the first step is to decide how to represent these subjective opinions over the possible bundle of goods the agent could receive. Some of the early work on fair division uses complete ordinal rankings: agents can have a nearly-arbitrary ordering over all $2^{|M|}$ subsets of the goods. Although of theoretical interest, as the number of goods grows, this domain quickly becomes too expressive and often leads to methods that are computationally infeasible (see, e.g. [Herreiner and Puppe, 2002]). On the other end, another common approach is to allow agents to express ordinal preferences over the $|M|$ singleton subsets and extend these to ordinal preferences over all possible bundles (see, e.g., [Brams and King, 2005; Brams *et al.*, 2003; Aziz *et al.*, 2015]). However, this suffers from the opposite problem and can often be too restrictive.

Recent work has thus focused on a different option, *additive cardinal preferences*. This preference domain is popular as it offers a sweet spot between simplicity and expressiveness. Here, each agent i places a non-negative value

$v_i(g)$ on each good g and her value for a bundle of goods $S \subseteq M$ is assumed to be the sum of her values for the individual goods in S , i.e., $\sum_{g \in S} v_i(g)$. Theoretically, this valuation class gives way to algorithms achieving strong fairness guarantees [Amanatidis *et al.*, 2016; Caragiannis *et al.*, 2019; Chaudhury *et al.*, 2020; Ghodsi *et al.*, 2018; Garg and Taki, 2020]. Practically, additive valuations are much simpler to elicit than fully combinatorial valuations, which has led to their adoption by popular fair division tools such as Spliddit and Adjusted Winner.¹

However, expressing additive valuations still requires placing an exact numerical value on each good, which can sometimes be difficult or infeasible. An interesting tradeoff can be achieved by eliciting ordinal preferences from the agents, but viewing them as partial information regarding underlying cardinal preferences. This idea originates from the related field of voting theory, where a growing body of work on the *distortion* framework uses ordinal preferences of voters over candidates as means to pick a candidate approximately maximizing social welfare according to the underlying cardinal preferences [Procaccia and Rosenschein, 2006; Boutilier *et al.*, 2015; Caragiannis *et al.*, 2017; Mandal *et al.*, 2019; Mandal *et al.*, 2020; Kempe, 2020; Amanatidis *et al.*, 2020].

In this paper, we focus on eliciting from each agent a ranking of her k most preferred goods (i.e., a prefix of her preference ranking over the goods). A system designer deliberating on whether to use such partial information over traditional cardinal valuations may immediately be interested in the *price* of the missing information. In line with the aforementioned work, we analyze distortion in the context of fair division, i.e., the worst-case (multiplicative) loss in social welfare — the sum of the values that agents place on their own bundles — incurred due to the missing information.

In addition, we are also interested in achieving qualitative fairness guarantees; it is, after all, *fair* division. Two popular guarantees for allocation of indivisible goods are envy-freeness up to one good (EF1) [Lipton *et al.*, 2004; Budish, 2011] and approximate maximin share guarantee (MMS) [Kurokawa *et al.*, 2018], which we define in Section 2. With access to agents' full preference rankings over the goods, it is known that EF1 can be achieved via the round robin algorithm [Lipton *et al.*, 2004; Caragiannis *et al.*, 2019],

¹www.spliddit.org, www.nyu.edu/projects/adjustedwinner/

under which agents take turns picking goods in a cyclic fashion. For MMS, Amanatidis *et al.* [2016] show that, using just ordinal preferences over the goods, it is impossible to guarantee better than a $1/H_n$ approximation of MMS, where $H_n = \Theta(\log n)$ is the n^{th} harmonic number and n is the number of agents; in contrast, given additive cardinal preferences, even $3/4$ -MMS can be achieved [Ghodsi *et al.*, 2018; Garg and Taki, 2020]. What is the best MMS approximation that can be achieved given agents’ preference rankings over all the goods? More generally, if we are only given agents’ preference rankings over their k most preferred goods, for what values of k can we achieve EF1 and approximate MMS? What distortion do we incur if, in addition to the missing cardinal information, we also impose these fairness requirements? We answer these questions in our work.

1.1 Our Contribution

A bit more formally, a deterministic (resp. randomized) ordinal allocation rule takes as input the partial preference rankings and returns an allocation (resp. a distribution over allocations) of the goods to the agents. The distortion of the rule is the ratio of the maximum social welfare of any allocation to the (expected) social welfare of the allocation returned, in the worst case over all problem instances in a family. As is common in the literature on distortion, we assume normalized valuations: the total value each agent places on all goods is normalized to 1. We are interested in two questions. First, how much information is needed to achieve certain fairness guarantees? Second, what is the best distortion of any deterministic or randomized ordinal allocation rule with or without a fairness constraint?²

Our results answer these questions for all values of k , but for simplicity, we summarize the results for when complete rankings are given ($k = m$) in Figure 1. Without any fairness constraint, the simple deterministic rule that simply allocates all the goods to a single agent achieves distortion n . We show that not even a randomized rule with access to complete rankings can achieve distortion better than n .

Next, we consider two fairness requirements: envy-freeness up to one good (EF1) and approximate maximin share (MMS). EF1 is known to be achievable given complete rankings ($k = m$). We characterize the exact value of k needed to achieve EF1. For MMS, we derive almost tight bounds the best possible approximation as a function of k . For the case of complete rankings ($k = m$), our results show that $1/(2H_n)$ -MMS is achievable, almost matching the asymptotic upper bound of $1/H_n$ due to Amanatidis *et al.* [2016]. Thus, we establish, for the first time, that the best approximation to MMS given ordinal preference information scales logarithmically in the number of agents.

We also show that when ordinal allocation rules are required to guarantee EF1 or α -MMS for $\alpha > 0$, deterministic rules face $\Omega(n^2)$ distortion while randomized rules face $\Omega(n)$ distortion, and matching upper bounds can be derived (in case of MMS, along with best-known α -MMS approximation). Our distortion upper bounds are achieved through effi-

²For a randomized rule, we require that the fairness constraint be met by *all* allocations in the support of the distribution returned.




Figure 1: The table on the left summarizes the optimal distortion for deterministic and randomized rules with access to the complete rankings ($k = m$). Note that α -MMS is achievable for $\alpha = 1/2H_n$, but not for $\alpha > 1/H_n$. However, the distortion lower bounds hold for any $\alpha > 0$. The diagram on the right shows the worst-case ratio of social welfare between pairs of settings from the following three: cardinal valuations given, cardinal valuations given but property P required, ordinal preferences given but property P required. The diagram holds for both $P \in \{\text{EF1}, \alpha\text{-MMS}\}$ and the top-right arrow, the price of fairness P , is due to Barman *et al.* [2020].

cient algorithms. In the full version,³ we also show that various other fairness guarantees studied in the literature cannot be achieved given just ordinal preference information, even with complete rankings.

1.2 Related Work

There has been a substantial amount of work on using ordinal preferences in fair allocation of indivisible goods. For example, Aziz *et al.* [2015] consider the question of checking the existence of allocations that possibly or necessarily satisfy certain fairness guarantees such as envy-freeness given only ordinal preferences of the agents over the goods. Bouveret *et al.* [2010] study similar questions, but given partial ordinal preferences of the agents over *bundles of goods*.

Some of the work does not assume any underlying cardinal preferences; instead, it aims to obtain guarantees defined directly in terms of the ordinal preferences. For example, Baumeister *et al.* [2017] and Nguyen *et al.* [2017] use the so-called scoring vectors to convert agents’ ordinal preferences into numerical proxies for their utility and then consider maximizing various notions of social welfare or guaranteeing various fairness properties in terms of such utilities.

Another related line of work uses ordinal allocation rules (such as picking sequence rules) in settings with cardinal valuations. For example, Aziz *et al.* [2016] focus on the complexity of checking what social welfare such rules can possibly or necessarily achieve. Amanatidis *et al.* [2016] seek to use picking sequence rules to obtain approximation of the maximin fair share guarantee; indeed, as mentioned earlier, we settle a question left open in their work. However, their main focus is on ensuring truthfulness, i.e., preventing agents from manipulating their preferences. Manipulations under picking sequence rules have received significant attention [Aziz *et al.*, 2017b; Aziz *et al.*, 2017a].

2 Model

For $j \in \mathbb{N}$, let $[j] = \{1, \dots, j\}$. Let $N = [n]$ be a set of *agents* and $M = [m]$ be a set of *goods*. Each agent i is en-

³www.cs.toronto.edu/nisarg/papers/distortion-fair-division.pdf

dowed with a *valuation* function $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$, which is *additive*: $v_i(S) = \sum_{g \in S} v_i(\{g\})$ for all $i \in N, S \subseteq M$; and *unit-sum*: $v_i(M) = 1$ for all $i \in N$. To simplify notation, we write $v_i(g)$ instead of $v_i(\{g\})$ for a good $g \in M$. We refer to $v = (v_1, \dots, v_n)$ as the *valuation profile*.

For $k \in [m]$, a *top- k ranking* σ_i of agent i is a ranking of agent i 's k most valuable goods (ties broken arbitrarily). We say that a good is *ranked* by an agent if it appears in their top- k ranking and *unranked* otherwise. We refer to $\sigma = (\sigma_1, \dots, \sigma_n)$ as the *top- k preference profile* (or, simply, preference profile). Note that the value of k is the same for all agents. When $k = m$, we refer to these as complete rankings. We say that v_i is consistent with σ_i , denoted $v_i \triangleright \sigma_i$, if $v_i(g) \geq v_i(g')$ for all $g, g' \in M$ such that either $g \succ_{\sigma_i} g'$ if both g and g' are ranked or g is ranked and g' is unranked. We say that v is consistent with σ , denoted $v \triangleright \sigma$, if $v_i \triangleright \sigma_i$ for each $i \in N$.

We are interested in taking as input (N, M, k, σ) , which we refer to as an *instance*, and finding an allocation of the goods to the agents. For a set of goods $S \subseteq M$ and $\ell \in \mathbb{N}$, let $\Pi_\ell(S)$ denote the set of ordered partitions of S into ℓ bundles. An allocation $A = (A_1, \dots, A_n) \in \Pi_n(M)$ is a partition of the goods into n bundles, where A_i is the bundle allocated to agent i . Under this allocation, the *utility* to agent i is $v_i(A_i)$. Given a valuation profile v , the *social welfare* of an allocation A is $\text{sw}(A, v) = \sum_{i \in N} v_i(A_i)$; we simply write $\text{sw}(A)$ when the valuation profile v is clear from the context.

We will use \mathcal{I} to denote a family of instances. We will use \mathcal{I}^R to denote the family of instances in which relation R over k, n , and m is satisfied. For example, $\mathcal{I}^{k=m}$ is the family of instances with complete rankings and $\mathcal{I}^{k \geq n-1}$ is the family of instances with rankings of at least $n-1$ goods.

A (randomized) *ordinal allocation rule* (hereinafter, simply a rule) f for a family of instances \mathcal{I} takes an instance (N, M, k, σ) from \mathcal{I} — for simplicity, we refer to σ as the sole input to f — and returns a distribution over the set of allocations $\Pi_n(M)$. We say that f is *deterministic* if it always returns a distribution with singleton support. We will sometimes refer to a distribution over allocations as a *randomized allocation*. The *distortion* of an ordinal allocation rule f with respect to a family of instances \mathcal{I} , denoted $\text{dist}^{\mathcal{I}}(f)$, is the worst-case approximation ratio it provides to the social welfare over all instances of \mathcal{I} :

$$\text{dist}^{\mathcal{I}}(f) = \sup_{(N, M, k, \sigma) \in \mathcal{I}} \sup_{v: v \triangleright \sigma} \frac{\max_{A \in \Pi_n(M)} \text{sw}(A, v)}{\mathbb{E}[\text{sw}(f(\sigma), v)]},$$

where the expectation is over possible randomization in f . When \mathcal{I} is clear from the context, we may drop it from the notation. Note that if $\mathcal{I}_1 \subseteq \mathcal{I}_2$, then $\text{dist}^{\mathcal{I}_1}(f) \leq \text{dist}^{\mathcal{I}_2}(f)$. Following prior work and to help compare our distortion bounds to the known price of fairness bounds (see Figure 1), we provide distortion bounds parametrized by the number of agents n . We are interested in the lowest distortion that deterministic and randomized ordinal allocation rules can achieve.

A *fairness property* P maps every instance $I = (N, M, k, \sigma)$ to a (possibly empty) set of allocations $P(I)$; every allocation in $P(I)$ is said to satisfy P in instance I . Often, fairness properties are defined in terms of agent valuations rather than rankings. In this case, an allocation is said

to satisfy P in instance I only if it is satisfied by P for all valuations consistent with σ . We say that a rule f satisfies property P if for all σ , every allocation in the support of $f(\sigma)$ satisfies P . We are also interested in determining whether ordinal allocation rules can satisfy prominent fairness properties, and when they can, determining the lowest possible distortion they can achieve subject to such properties.

Given an instance $I = (N, M, k, \sigma)$ with valuations v , we are interested in the following fairness properties.

Definition 1 (EF1). An allocation A is called *envy-free up to one good (EF1)* if for every pair of agents i, j , either $v_i(A_i) \geq v_i(A_j)$ or there exists a good $g \in A_j$ such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$.

Definition 2 (Balancedness). An allocation A is called *balanced* if $|A_i| - |A_j| \leq 1$ for all $i, j \in N$, i.e., if the agents receive approximately an equal number of goods.

Definition 3 (MMS). The *maximin share* of agent i is

$$\text{MMS}_i = \max_{A \in \Pi_n(M)} \min_{A_j \in A} v_i(A_j).$$

Given $\alpha \in [0, 1]$, an allocation A is called *α -maximin share fair (α -MMS)* if $v_i(A_i) \geq \alpha \cdot \text{MMS}_i$ for all agents $i \in N$. When $\alpha = 1$, we simply say that A is an MMS allocation.

3 Distortion of Ordinal Allocation Rules

We begin by analyzing the lowest distortion that deterministic and randomized ordinal allocation rules can achieve in the absence of any fairness requirement. This precisely captures value of cardinal preference information, or the loss incurred in social welfare due to having only ordinal preferences.

Even without any preference information ($k = 0$), a trivial *deterministic* rule that allocates all the goods to an arbitrary single agent achieves distortion n : indeed, the social welfare of such an allocation is 1, while the maximum social welfare cannot be larger than n since valuations are unit-sum ($v_i(M) = 1$ for all $i \in M$). We show that not even *randomized* ordinal allocation rules with access to complete rankings ($k = m$) can achieve lower distortion. The proof can be found in the full version.

Theorem 1. *There exists a deterministic ordinal allocation rule with distortion n for the family $\mathcal{I}^{k \geq 0}$. On the other hand, no randomized ordinal allocation rule achieves distortion lower than n even for the restricted family of $\mathcal{I}^{k=m}$.*

3.1 Fairness Lower Bounds

In this section, we analyze the lowest distortion that ordinal allocation rules can achieve when they are required to satisfy some fairness constraints. This captures the *combined price* of the lack of cardinal preference information and the imposition of fairness constraints. Figure 1 contrasts this with the sole price of the former analyzed in Section 3 and the sole price of the latter from known results in the literature. Perhaps not surprisingly, it turns out that the two together lead to a much greater loss in social welfare than each individually.

Another consequence of our results is that while randomized ordinal rules are no more powerful than deterministic ones in the absence of any fairness requirements (Theorem 1), imposing fairness requirements makes their powers diverge.

Keeping aside the question of distortion, we are also interested in determining which fairness properties ordinal allocation rules can satisfy. A negative answer can be interpreted as a qualitative price of the lack of cardinal preferences.

We begin by establishing a lower bound on the distortion of deterministic ordinal allocation rules that holds when any fairness property from a broad class is imposed, even with access to complete rankings; later, we argue that the fairness properties of our interest belong to this class. Recall that we require the allocation returned by the rule to satisfy the fairness property, regardless of the unobserved cardinal valuations (consistent with the observed ordinal preferences).

Theorem 2. *Let P be a fairness property such that when the number of goods equals the number of agents, for every preference profile σ , an allocation satisfies P for all valuations consistent with σ if and only if each agent receives a single good. Then, the distortion of every deterministic ordinal allocation rule satisfying P is $\Omega(n^2)$ for the family $\mathcal{I}^{k=m}$.*

Proof. Fix such a fairness property P , a number of agents n , and a deterministic ordinal allocation rule f on $\mathcal{I}^{k=m}$ (i.e., only taking complete rankings) satisfying P . First, let us suppose that n is even. We construct an instance with n goods, that is, with $m = n$. We split the goods into three different categories and construct a preference profile σ as follows. The first category consists of a single good g^* that is ranked highest by all agents. The next category consists of $n/2$ goods labeled $g_{\{1,2\}}, g_{\{3,4\}}, \dots, g_{\{n-1,n\}}$. For each $\ell \in [n/2]$, good $g_{\{2\ell-1,2\ell\}}$ is ranked second by both agents $2\ell-1$ and 2ℓ . The final category consists of the remaining $n/2-1$ goods. The construction above identifies the two most preferred goods for all agents; their preference rankings from the third rank onward can be arbitrary.

Let A be the allocation returned by f given σ . By the assumption of the theorem statement, each agent must receive exactly one good in A . Without loss of generality, let us assume that agent 1 receives g^* . In addition, for each $\ell \in \{2, \dots, n/2\}$, at least one of agents $2\ell-1$ and 2ℓ does not receive good $g_{\{2\ell-1,2\ell\}}$; without loss of generality, assume that agent $2\ell-1$ does not receive it. Let us construct a consistent valuation profile as follows:

- Agent 1 has value $1/n$ for each good.
- Agent 2 has value 1 for g^* and 0 for all other goods.
- For $\ell \in \{2, \dots, n/2\}$, agent $2\ell-1$ has value $1/2$ for g^* , $1/2$ for $g_{\{2\ell-1,2\ell\}}$, and 0 for all other goods; and agent 2ℓ has value 1 for g^* and 0 for all other goods.

Under A , the only agent receiving positive utility is agent 1, who receives utility $1/n$. Therefore, the social welfare is $1/n$. In contrast, consider the allocation that gives g^* to agent 2, $g_{\{2\ell-1,2\ell\}}$ to agent $2\ell-1$ for each $\ell \in \{2, \dots, n/2\}$, and the remaining goods arbitrarily such that each agent receives a single good. It is easy to check that its social welfare is at least $1 + (n/2-1) \cdot 1/2 = n/4 + 1/2$. Therefore, the distortion of f is at least $(n/4 + 1/2)/(1/n) \in \Omega(n^2)$.

If n is odd, we can construct the described instance with $n-1$ agents and $n-1$ goods, add a good ranked last by all agents, and add an agent whose preference ranking matches

that of one of the other agents. Using similar arguments as above, regardless of the allocation A chosen by f , we can construct a consistent valuation profile in which the social welfare of A is $1/n$, while the optimal social welfare is at least $(n-1)/4 + 1/2$, resulting in $\Omega(n^2)$ distortion. \square

Notice that in the proof of Theorem 2, we contrast the social welfare achieved by the rule satisfying P against that of an allocation that assigns each agent a single good, thus also satisfying P . That is, the $\Omega(n^2)$ lower bound continues to hold even when comparing to the optimal social welfare *subject to* P . On the other hand, our matching upper bounds presented later hold even when comparing to the optimal social welfare *without any fairness constraints*.

4 EF1

We now turn our attention to EF1. We begin by fully characterizing the values of k , in relation to n and m , for which we can achieve EF1. The rules we construct are based on *picking sequence rules*. A picking sequence is simply a sequence of agents p_1, \dots, p_ℓ , where $\ell \leq m$. It is a deterministic rule that works as follows: it first gives agent p_1 their favorite good, then gives agent p_2 their favorite good among the ones remaining, and so on, for ℓ steps. If $\ell < m$, we design a way — different for each case — for allocating the remaining goods. A well-known picking sequence rule is *round robin*, which has the cyclic picking sequence $1, \dots, n, 1, \dots, n, \dots$ repeated for a total of m steps.

Theorem 3. *With n agents and m goods, it is possible to guarantee EF1 using top- k rankings if and only if*

$$k \geq \begin{cases} m - n, & \text{if } m \bmod n = 0; \\ m - 2, & \text{if } m \bmod n = 1; \\ m - (m \bmod n), & \text{if } m \bmod n > 1. \end{cases}$$

Proof. Fix arbitrary n , m , and k . We begin with the lower bounds, showing that EF1 can only be achieved if k is sufficiently large. All of our constructions have the same preference profile: all agents agree on which goods are in the top k , that is, they rank goods g_1, \dots, g_k in some order and do not rank the remaining $m-k$ goods. For a given allocation A , let $s_i = |A_i \cap \{g_1, \dots, g_k\}|$ be the number of the top- k goods received by agent i . Note that $k \geq \sum_{i \in N} s_i$. We use the following lemma.

Lemma 1. *If agents agree on which goods are in the top k and an allocation A is EF1 for all consistent valuations, then $s_i \geq |A_j| - 1$ for all distinct agents $i, j \in N$.*

Proof. Consider a consistent valuation profile in which agent i has zero value for the goods $A_i \setminus \{g_1, \dots, g_k\}$ but equal value for all other goods (including all of the ones in A_j). If $s_i < |A_j| - 1$, EF1 would be violated for agent i . \square

Suppose there exists an allocation A guaranteed to be EF1. We show that this implies k is sufficiently large as per the theorem statement. Let $q \in \mathbb{N}$ and $r \in [n-1]$ be such that $m = qn + r$. Since A is guaranteed to be EF1 given ordinal preferences, it must be balanced, that is, $|A_j| - |A_i| \leq 1$ for all agents i and j : using Lemma 1, we can see that $|A_i| \geq s_i \geq$

$|A_j| - 1$. In our case, this means that r agents have bundles with size $q + 1$ and the remaining $n - r$ have bundles with size q . In the following, we use $k \geq \sum_{i \in N} s_i$.

- Suppose $r = 0$, so $|A_i| = q$ for all agents i . By Lemma 1, $s_i \geq q - 1$ all agents i , so $k \geq (q - 1)n = m - n$.
- Suppose $r = 1$. Therefore, one bundle, without loss of generality A_1 , has size $q + 1$, and all the others have size q . We have that $s_1 \geq q - 1$ and $s_i \geq q$ for all $i \neq 1$ by Lemma 1. This implies $k \geq qn - 1 = m - 2$.
- Suppose $r > 1$. As at least two agents have bundles of size $q + 1$, by Lemma 1, $s_i \geq q$ for all agents i . This implies $k \geq qn$.

Next, we prove the upper bounds, showing that if k is sufficiently large, EF1 can be guaranteed. To do this, we make modifications to the aforementioned round robin rule, which (with a picking sequence of length m) is known to guarantee EF1. Note that to run this rule, only ordinal information is needed but complete rankings (or at least $k \geq m - 1$) are needed. To work with round robin, we label the goods in the order they are chosen as follows $g_{1,1}, g_{2,1}, \dots, g_{n,1}, g_{1,2}, \dots, g_{\ell,t}$, where good $g_{i,j}$ is the j^{th} good picked by agent i . We refer to t , the largest number of goods picked by any agent, as the number of “rounds” and goods $g_{i,t}$ for $i \in N$ as goods picked in the last round. Note that not all agents will necessarily pick a good in the last round. We make use of the following lemma. The proof can be found in the full version.

Lemma 2. *Suppose we run round robin to m steps with access to complete rankings but reassign the goods received by agents in the last round such that no agent receives more than one good. Then the resulting allocation remains EF1.*

All the rules we design will run round robin to at most k steps and then assign the remaining goods in a way that can be accomplished without access to the remaining ordinal preferences. We will then argue that the remaining goods were assigned in a way such that they only permuted the last round of goods in a hypothetical allocation given by running round robin to m steps with access to complete rankings. By Lemma 2, this implies EF1.

First, suppose $m \bmod n = 0$ and $k \geq m - n$. The rule works as follows: it runs round robin for $m - n$ steps and assigns the remaining n goods so that each agent receives exactly one. This is EF1 by Lemma 2 as the resulting allocation could also have been computed by running round robin using complete rankings and reassigning the goods from the last round in the way that was arbitrarily chosen.

Next, suppose $m \bmod n = 1$ and $k \geq m - 2$. The rule works in a very similar way. It runs round robin for $m - 2$ steps and assigns the remaining two goods to the last agent in the order (or if $m = 1$, just assigns the good to an arbitrary agent). This is EF1 by Lemma 2 as the resulting allocation could have been computed by running round robin and giving the singular good of the final round to the n^{th} agent in the order.

Finally, suppose $m \bmod n > 2$ and $k \geq m - (m \bmod n)$. As before, we run round robin for $m - (m \bmod n)$ steps and

assign the remaining $m \bmod n$ goods such that each agent receives at most one. This is again EF1 by Lemma 2. \square

Let \mathcal{I}^{EF1} be the family of instances with k , n , and m values satisfying the relation specified in Theorem 3 (i.e. for which it is possible to achieve EF1). Then, the best possible distortion subject to the requirement of achieving EF1 is as follows. The proof can be found in the full version.

Theorem 4. *Among deterministic ordinal rules, all EF1 rules have unbounded distortion on $\mathcal{I}^{EF1} \cap \mathcal{I}^{k=0}$ and the lowest possible distortion of an EF1 rule with respect to $\mathcal{I}^{EF1} \cap \mathcal{I}^{k \geq 1}$ is $\Theta(n^2)$. Among randomized ordinal rules, the lowest possible distortion of an EF1 rule on \mathcal{I}^{EF1} is $\Theta(n)$.*

5 MMS

We now turn to our most technical results, which are regarding approximate maximin share (MMS) guarantee. Before we consider distortion subject to approximate MMS, we need to know what approximation to MMS is possible to achieve. Given full cardinal information, it is known that exact MMS cannot be achieved [Kurokawa *et al.*, 2018], but $3/4$ -MMS can [Ghods *et al.*, 2018; Garg and Taki, 2020].

Given complete preference rankings ($k = m$), Amanatidis *et al.* [2016] show that it is not possible to achieve α -MMS for $\alpha > 1/H_n$, where $H_n = \Theta(\log n)$ is the n^{th} Harmonic number. On the opposite end, they only establish a weaker $\Omega(1/\sqrt{n})$ lower bound, leaving open the question of what the best possible MMS approximation is given complete preference rankings. We settle this question by showing that the best possible MMS approximation for $k = m$ is $\Theta(1/H_n)$ (specifically, we derive a lower bound of $1/2H_n$). We also extend the lower and upper bounds to the case of $k < m$.

Our algorithm is similar to the one provided by Amanatidis *et al.* [2016] to achieve $\Omega(1/\sqrt{n})$, but our improvement crucially relies on Lemma 3, which requires an intricate proof (given in ??) to achieve the desired bounds. Note that for $m \leq n$, MMS can trivially be satisfied by giving each agent at most one good; thus, we focus on $m > n$.

Theorem 5. *When $m > n$, the following hold.*

- If $k < n - 1$, we cannot achieve α -MMS for any $\alpha > 0$.
- If $k = n - 1$, we can achieve $\lfloor \frac{1}{\frac{m-n+2}{2}} \rfloor$ -MMS, but no higher.
- If $k \geq n$, we can achieve α -MMS for $\alpha = \frac{k-n+1}{m-n+1} \cdot \frac{1}{2H_n}$, but not for $\alpha > \frac{k}{H_n(m-n) - (m-k)}$.

Proof. Most of the proof is located in ??. Here, we show the most interesting case: the lower bound for $k \geq n$. We borrow and build upon ideas from the proof of the $\Omega(1/\sqrt{n})$ lower bound due to Amanatidis *et al.* [2016].

We construct a picking sequence rule achieving the desired MMS approximation. Fix an agent. For now, suppose we are working with complete rankings, with $k = m$. Suppose the first time this agent appears in the picking sequence is at the ℓ^{th} position (we call this the agent’s 0^{th} appearance) for some $\ell \leq n$, and then, the agent’s j^{th} appearance occurs at or before position $(\ell + \lfloor j \cdot 2H_n(n - \ell + 1) \rfloor)$ in the picking

sequence for every j (as long as this quantity does not exceed m). Then, we claim that the agent must be guaranteed at least $1/(2H_n)$ fraction of her MMS value. To see this, note that the agent picks a good at least as valuable as her ℓ^{th} most favorite good in her 0^{th} appearance, and then an additional good at least as valuable as her $(\ell + \lfloor j \cdot 2H_n(n - \ell + 1) \rfloor)^{\text{th}}$ favorite good in her j^{th} appearance for each j . Let S denote the total value the agent places on her $\ell - 1$ most valuable goods. Then, this picking sequence guarantees the agent utility at least $\frac{1-S}{2H_n(n-\ell+1)}$. On the other hand, note that the MMS value of the agent is at most $\frac{1-S}{n-\ell+1}$; this is because regardless of how the agent partitions the goods into n bundles, ignoring the (at most) $\ell - 1$ bundles containing her $\ell - 1$ most valuable goods, even the average value across the remaining (at least $n - \ell + 1$) bundles is at most $\frac{1-S}{n-\ell+1}$. Hence, it follows that the agent is guaranteed at least a $1/2H_n$ fraction of her MMS.

Now, instead of assuming we have complete rankings, suppose we just have top- k for some $k \geq n$. In this case, suppose we run the above picking sequence for just k steps. Then, rather than being guaranteed $\frac{1-S}{2H_n(n-\ell+1)}$, the agent is only guaranteed $\frac{k-\ell+1}{m-\ell+1} \cdot \frac{1-S}{2H_n(n-\ell+1)}$. As their MMS value remains the same and $\ell \leq n$, this agent is guaranteed $\frac{k-n+1}{m-n+1} \cdot \frac{1}{2H_n}$ -MMS as needed. The remaining $m - k$ goods can be allocated arbitrarily.

Our picking sequence gives a guarantee of this style to each agent, albeit for different values of ℓ . In particular, for each agent $i \in [n]$, the picking sequence provides this guarantee with $\ell = i$. The construction is very simple.

1. For $1 \leq i \leq n$ and $0 \leq j \leq \lfloor \frac{m-i}{2H_n(n-i+1)} \rfloor$, we create the pair $(i, i + \lfloor j \cdot 2H_n(n - i + 1) \rfloor)$, indicating that agent i 's j^{th} appearance must occur at or before the position indicated in the second component — we refer to this as the deadline.
2. We sort the pairs with respect to their second coordinate.
3. The first coordinates with respect to the above sorting are a prefix of the picking sequence.
4. If the length of the above sequence is m , we are done; otherwise we arbitrarily assign the remaining picks.
5. If $k < m$, we truncate the sequence to length k .

The idea of Steps 2–4 is to produce a picking sequence that meets all the deadlines by using earliest-deadline-first scheduling. It is known that if all the deadlines can be met, then this greedy scheduling procedure is guaranteed to return a sequence meeting them. To show that all deadlines are met, we want to show that there are at most d pairs introduced in Step 1 with the second coordinate (deadline) at most d , for all $d \leq m$. Note that in particular, this implies that there are at most m pairs in total, so Step 3 would not produce a sequence of length more than m .

To prove this, let us first consider $d \leq n$. Observe that the 1^{st} appearance deadline of any agent is at or after position $n + 1$: this is because $i + \lfloor 2H_n(n - i + 1) \rfloor \geq 1 + \lfloor n - i + 1 \rfloor = n + 1$ for all $i \in [n]$. This implies that the only pairs with deadline at most n are the n pairs of the form (i, i) for $i \in [n]$

corresponding to the 0^{th} appearances of all the agents, which immediately implies the desired goal holds for all $d \leq n$.

Next, consider $d \geq n + 1$. The number of pairs for agent i with the second coordinate at most d is at most $1 + \lfloor \frac{d-i}{2H_n(n-i+1)} \rfloor$. Therefore, the number of total pairs with second coordinate at most d is at most $\sum_{i=1}^n 1 + \lfloor \frac{d-i}{2H_n(n-i+1)} \rfloor = n + \sum_{i=1}^n \lfloor \frac{d-i}{2H_n(n-i+1)} \rfloor$. Our goal is to show that this value is at most d , which is equivalent to the following lemma. The proof is deferred to ??.

Lemma 3. *For all $n \in \mathbb{N}$ and for all $d \geq n + 1$, $\sum_{i=1}^n \lfloor \frac{d-i}{2H_n(n-i+1)} \rfloor \leq d - n$.*

This completes the proof of the theorem. \square

Strikingly, while $\Omega(n^2)$ distortion is unbeatable subject to α -MMS for any $\alpha > 0$, for $k \geq n$, we can achieve a matching $O(n^2)$ distortion even while simultaneously achieving the best-known MMS approximations for all k , introduced in Theorem 5. The proof is deferred to ??.

Theorem 6. *The best possible distortion is as follows.*

- On $\mathcal{I}^{k=n-1}$, any deterministic ordinal rule satisfying α -MMS for $\alpha > 0$ must have unbounded distortion. However, there is a randomized ordinal rule achieving the best-possible $1/\lfloor \frac{m-n+2}{2} \rfloor$ -MMS and distortion n .
- On $\mathcal{I}^{k \geq n}$, any deterministic ordinal rule satisfying α -MMS for $\alpha > 0$ must have distortion $\Omega(n^2)$, and there is a deterministic ordinal rule achieving $O(n^2)$ distortion with α -MMS for the best-known $\alpha = \frac{k-n+1}{m-n+1} \cdot \frac{1}{2H_n}$. Further, there is a randomized ordinal rule that achieves α -MMS for $\alpha = \frac{k-n+1}{m-n+1} \cdot \frac{1}{2H_n}$ with distortion n .

6 Discussion

In this paper, we analyze which fairness properties can be achieved and what loss in social welfare must be incurred (distortion) when only ordinal preference information is provided in the form of top- k rankings.

This is inspired by a growing literature on distortion in voting [Procaccia and Rosenschein, 2006]. A recent line of work has focused on imposing additional structure on the underlying cardinal preferences [Anshelevich *et al.*, 2018; Gkatzelis *et al.*, 2020]. In fair division, one can also study natural restrictions on the underlying cardinal preferences such as a limit on the number of goods an agent can derive positive utility from or on the maximum difference between the values two agents can derive from the same good.

Another thread of research on distortion in voting has focused on the tradeoff between distortion and the amount of preference information elicited [Mandal *et al.*, 2019; Mandal *et al.*, 2020; Kempe, 2020; Amanatidis *et al.*, 2020]. Our results already offer one such tradeoff by allowing the designer to pick the value of k . An interesting direction for the future would be to study such a tradeoff while allowing arbitrary — not necessarily ordinal — elicitation and measuring the number of bits of information elicited.

References

- [Amanatidis *et al.*, 2016] G. Amanatidis, G. Birmpas, and E. Markakis. On truthful mechanisms for maximin share allocations. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 31–37, 2016.
- [Amanatidis *et al.*, 2020] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas, and A. A. Voudouris. Peeking behind the ordinal curtain: Improving distortion via cardinal queries. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1782–1789, 2020.
- [Anshelevich *et al.*, 2018] E. Anshelevich, O. Bhardwaj, E. Elkind, J. Postl, and P. Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.
- [Aziz *et al.*, 2015] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.
- [Aziz *et al.*, 2016] H. Aziz, T. Kalinowski, T. Walsh, and L. Xia. Welfare of sequential allocation mechanisms for indivisible goods. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*, pages 787–794, 2016.
- [Aziz *et al.*, 2017a] H. Aziz, S. Bouveret, J. Lang, and S. Mackenzie. Complexity of manipulating sequential allocation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 328–334, 2017.
- [Aziz *et al.*, 2017b] H. Aziz, P. Goldberg, and T. Walsh. Equilibria in sequential allocation. In *Proceedings of the 5th International Conference on Algorithmic Decision Theory (ADT)*, pages 270–283, 2017.
- [Barman *et al.*, 2020] S. Barman, U. Bhaskar, and N. Shah. Optimal bounds on the price of fairness for indivisible goods. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*, pages 356–369, 2020.
- [Baumeister *et al.*, 2017] D. Baumeister, S. Bouveret, J. Lang, N. T. Nguyen, T. T. Nguyen, J. Rothe, and A. Saffidine. Positional scoring-based allocation of indivisible goods. *Autonomous Agents and Multi-Agent Systems*, 31(3):628–655, 2017.
- [Boutilier *et al.*, 2015] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- [Bouveret *et al.*, 2010] S. Bouveret, U. Endriss, and J. Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI)*, pages 387–392, 2010.
- [Bouveret *et al.*, 2016] S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair allocation of indivisible goods. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, pages 284–310. Cambridge University Press, 2016.
- [Brams and King, 2005] S. J. Brams and D. L. King. Efficient fair division: Help the worst off or avoid envy? *Rationality and Society*, 17(4):387–421, 2005.
- [Brams *et al.*, 2003] S. J. Brams, P. H. Edelman, and P. C. Fishburn. Fair division of indivisible items. *Theory and Decision*, 55(2):147–180, 2003.
- [Budish, 2011] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2017] I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research*, 58:123–152, 2017.
- [Caragiannis *et al.*, 2019] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.
- [Chaudhury *et al.*, 2020] B. R. Chaudhury, J. Garg, and K. Mehlhorn. EFX exists for three agents. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 1–19, 2020.
- [Garg and Taki, 2020] J. Garg and S. Taki. An improved approximation algorithm for maximin shares. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 379–380, 2020.
- [Ghodsí *et al.*, 2018] M. Ghodsí, M. HajiAghayi, M. Seddighin, S. Seddighin, and H. Yami. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 539–556, 2018.
- [Gkatzelis *et al.*, 2020] V. Gkatzelis, D. Halpern, and N. Shah. Resolving the optimal metric distortion conjecture. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*, pages 1427–1438, 2020.
- [Herreiner and Puppe, 2002] D. Herreiner and C. Puppe. A simple procedure for finding equitable allocations of indivisible goods. *Social Choice and Welfare*, 19(2):415–430, 2002.
- [Kempe, 2020] D. Kempe. Communication, distortion, and randomness in metric voting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2087–2094, 2020.
- [Kurokawa *et al.*, 2018] D. Kurokawa, A. D. Procaccia, and J. Wang. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM*, 65(2):1–27, 2018.
- [Lipton *et al.*, 2004] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, pages 125–131, 2004.
- [Mandal *et al.*, 2019] D. Mandal, A. D. Procaccia, N. Shah, and D. P. Woodruff. Efficient and thrifty voting by any means necessary. In *Proceedings of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 7180–7191, 2019.
- [Mandal *et al.*, 2020] 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)*, pages 795–813, 2020.
- [Nguyen *et al.*, 2017] N. T. Nguyen, T. T. Nguyen, and J. Rothe. Approximate solutions to max-min fair and proportionally fair allocations of indivisible goods. In *Proceedings of the 16th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 262–271, 2017.
- [Procaccia and Rosenschein, 2006] A. D. Procaccia and J. S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, pages 317–331, 2006.
- [Steinhaus, 1948] H. Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.