
FAIR AND EFFICIENT RESOURCE ALLOCATION WITH PARTIAL INFORMATION*

Daniel Halpern
Harvard University
dhalpern@g.harvard.edu

Nisarg Shah
University of Toronto
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.²

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

*A preliminary version appeared in the proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI-21).

²www.spliddit.org, www.nyu.edu/projects/adjustedwinner/

underlying cardinal preferences Procaccia and Rosenschein [2006]; Boutilier *et al.* [2015]; Caragiannis *et al.* [2017]; Mandal *et al.* [2019, 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.* [2016], 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 (where each agent provides a ranking of her k most preferred goods) 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?³

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 efficient algorithms. In Appendix A, 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*.

³For a randomized rule, we require that the fairness constraint be met by *all* allocations in the support of the distribution returned.

Fairness	Det	Rand
None	n	n
EF1	$\Theta(n^2)$	$\Theta(n)$
α -MMS	$\Theta(n^2)$	$\Theta(n)$




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].

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,a].

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 endowed 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.

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 $\mathcal{I}^{k=m}$.*

Proof. Given the observation prior to the theorem statement, we only need to show that the distortion of an arbitrary randomized ordinal allocation rule f , even with access to the complete rankings, is at least n .

Fix an integer x , and let $\varepsilon = 1/x$. Let us consider a preference profile σ over x^n goods, under which every agent has the same preference ranking: lower indexed goods are preferred to higher indexed goods. We construct a family of consistent valuations \mathcal{V} , and show that regardless of the randomized allocation returned by f given σ , the worst-case loss in social welfare across just this family of valuations already approaches n as x grows large.

We note that our construction below is similar to the one used by Bei *et al.* [2019] for lower bounding the social welfare loss incurred by a specific deterministic ordinal allocation rule, but our analysis significantly more intricate as it applies to all randomized ordinal allocation rules.

For $\ell \in [n]$, we say that an agent is of type t_ℓ if they like goods 1 through x^ℓ equally, at value $1/x^\ell$ each, and have value 0 for the remaining goods. Let $T = \{t_1, \dots, t_n\}$ be the set of all types. The family \mathcal{V} consists of valuation profiles under which there is exactly one agent of each type. Such valuations can be represented by a bijection $\tau : [n] \mapsto T$ mapping agents to types.

Consider the partition of goods $W = (W_1, \dots, W_n)$ such that $W_1 = \{1, \dots, x\}$, and $W_\ell = \{x^{\ell-1} + 1, \dots, x^\ell\}$ for $2 \leq \ell \leq n$. First, note that under every valuation in \mathcal{V} , there exists an allocation with social welfare at least $(1 - \varepsilon)n$. This is achieved by assigning each W_ℓ to the agent of type t_ℓ , which gives them value at least $(1/x^\ell) \cdot |W_\ell| \geq (1/x^\ell) \cdot (x^\ell - x^{\ell-1}) \geq 1 - 1/x = 1 - \varepsilon$, resulting in social welfare at least $(1 - \varepsilon)n$. In contrast, we show that under some valuation in \mathcal{V} , the randomized allocation returned by f generates poor social welfare.

First, note that under any valuation in \mathcal{V} , all agents are indifferent between the goods in W_ℓ , for each $\ell \in [n]$. Hence, we can succinctly describe a randomized allocation by the expected fraction of goods from W_ℓ assigned to the different agents, for each $\ell \in [n]$. In other words, we can describe a randomized allocation \mathcal{A} by an $n \times n$ matrix X such that $X_{i,\ell}$ is the expected fraction of goods from W_ℓ assigned to agent i . Note that for the randomized allocation to be feasible, it must be the case that $\sum_{i=1}^n X_{i,\ell} = 1$ for all $\ell \in [n]$.

Let X be the matrix corresponding to the randomized allocation \mathcal{A} returned by f on input σ constructed above. Our goal is to find a valuation $v \in \mathcal{V}$, characterized by a bijection $\tau : [n] \rightarrow T$, under which the expected social welfare of \mathcal{A} is not much more than 1. First, we claim that if agent i is of type t_ℓ , then their expected utility under \mathcal{A} is at most $X_{i,\ell} + \varepsilon$. To see this, note that agent i is in expectation receiving an $X_{i,\ell}$ fraction of the goods from W_ℓ , and W_ℓ consists of at least a $1 - \varepsilon$ fraction of the goods they like, as argued above. Hence, the expected utility to agent i under such an allocation is at most $X_{i,\ell}$ from the goods in W_ℓ and at most ε from the goods outside of W_ℓ , which is at most $X_{i,\ell} + \varepsilon$ in total. Thus, the expected social welfare of \mathcal{A} is at most $\sum_{i,\ell:\tau(i)=t_\ell} X_{i,\ell} + n\varepsilon$. Our goal is to show that there exists a bijection τ for which this quantity is at most $1 + n\varepsilon$.

Note that if we choose τ uniformly at random, then the expected value of $\sum_{i,\ell:\tau(i)=t_\ell} X_{i,\ell}$ is equal to $1/n \cdot \sum_{i \in [n], \ell \in [n]} X_{i,\ell} = 1$. Hence, there must exist a bijection τ for which $\sum_{i,\ell:\tau(i)=t_\ell} X_{i,\ell} \leq 1$. Under the corresponding valuation profile, the expected social welfare of \mathcal{A} is at most $1 + n\varepsilon$. This shows that the distortion of f is at least $\frac{(1-\varepsilon)n}{1+n\varepsilon}$. As x approaches ∞ , ε approaches 0, which establishes that the distortion of f is at least n , as desired. \square

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.

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.*

Proof. The proof is a slight modification of the classic proof that round robin is EF1. First, we will show the classic proof for completeness and then show how to modify it to achieve our desired result. Without loss of generality, we will add dummy goods $g_{\ell+1,t}, \dots, g_{n,t}$, which all agents value at zero, so that all agents receive the same number of goods. Note that the addition of dummy goods does not affect whether an allocation is EF1. We have that $A_i = \{g_{i,j} | 1 \leq j \leq t\}$. Fix two agents i and i' ; we will show that i does not envy i' up to one good. First, note that $v_i(g_{i,j}) \geq v_i(g_{i',j+1})$ for all $1 \leq j \leq t - 1$. This is because $g_{i',j+1}$ was available on i 's j^{th} pick, but i preferred good $g_{i,j}$. Therefore, we have that

$$v_i(A_i) \geq v_i(A_i \setminus \{g_{i,t}\}) = \sum_{j=1}^{t-1} v_i(g_{i,j}) \geq \sum_{j=1}^{t-1} v_i(g_{i',j+1}) = \sum_{j=2}^t v_i(g_{i',j}) = v_i(A_j \setminus \{g_{j,1}\})$$

as needed.

Now for the modification. We have that $g_{i,t}$ may be replaced by any good received in the t^{th} round and the above inequalities still hold as $g_{i,t}$ is simply ignored from A_i . Further, $g_{i',t}$ may be replaced by any good g in the t^{th} round and $v_i(g_{i,t-1}) \geq v_i(g)$ will still hold. This implies EF1 will still hold even after the reassignment. \square

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

Next, we will consider the best possible distortion of these rules. However, when analyzing the distortion of randomized rules, we will need the following observations. Given a deterministic picking rule, one way to turn this into a randomized rule, known as its uniform variant, is as follows: pick a permutation of the agents τ uniformly at random, and then run the rule using this agent labeling. Earlier, we observed that distortion n can be achieved trivially by allocating all the goods to an arbitrary single agent. The following result shows that a more interesting class of rules — *uniform picking sequence rules* — also achieves distortion n . This result is a well-known folklore in the literature, but we present a proof for completeness. A proof for the special case of round robin method was given by Freeman *et al.* [2020].

Proposition 1. *If a rule begins with a picking sequence of at least cm steps for some constant $0 < c \leq 1$, then the uniform variant has distortion in $[n, n/c]$. In particular, it has distortion $\Theta(n)$, and if $c = 1$, it has distortion n .*

Proof. Consider some rule f that always runs a picking sequence of length cm . Fix an arbitrary $I = (N, M, \sigma, k)$ with m goods and underlying valuations v . Let the picking sequence be $p_1 p_2 \dots p_\ell$ for $\ell \geq cm$. We show that the

expected social welfare under the rule is at least c . Since the optimal social welfare is upper bounded by n , this gives the desired upper bound of cn on distortion.

To show the expected social welfare is at least c , we show that the expected utility to each agent i is at least c/n . Suppose agent i has preference ranking $g_1 \succ \dots \succ g_m$. Because utilities are additive, we can decompose the expected utility to agent i as the sum of the expected utilities received by agent i from all ℓ picks.

Note that agent i only receives a nonzero utility from the j^{th} pick if $\tau(p_j) = i$, where τ is the random permutation chosen by the rule. This happens with probability $1/n$, and with this probability, agent i picks a good at least as valuable to her as g_j (since one of g_1, \dots, g_j must be available during j^{th} pick), thus receiving utility at least $v_i(g_j)$. Therefore, the expected utility of agent i from the j^{th} pick is at least $1/n \cdot v_i(g_j)$ for each $j \leq \ell$. Using linearity of expectation, we get that the expected utility to agent i is at least $\sum_{j=1}^{\ell} 1/n \cdot v_i(g_j) \geq 1/n \cdot c \cdot v_i(M) = c/n$, where the first transition holds because g_1, \dots, g_ℓ account for at least a c fraction of their total value. \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.

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 $\Theta(n)$.*

Proof. First, let us consider $\mathcal{I}^{EF1} \cap \mathcal{I}^{k=0}$. We will show it has unbounded distortion. This family includes an instance with two agents, 1 and 2, two items g_1 and g_2 , and empty rankings. As shown in the proof of Theorem 3, EF1 implies balanced, so any EF1 rule must give one item to each agent. Without loss of generality, suppose g_1 goes to agent 1 and g_2 goes to agent 2. It's consistent with the rankings that agent 1 has value one for g_2 and zero for g_1 , while agent 2 has value one for g_1 and zero for g_2 . This shows the distortion of the rule is unbounded.

Next, let us consider $\mathcal{I}^{EF1} \cap \mathcal{I}^{k \geq 1}$. By Theorem 2, the distortion of any EF1 rule on this family must be $\Omega(n^2)$. We will show that $O(n^2)$ distortion is possible. To do this, we will find an EF1 rule that always gives agent 1 value at least $1/3n$. As the total welfare of any allocation is at most n , this will imply the rule has the desired distortion. We will split the rule into two cases, $m \leq n$ and $m > n$. If $m \leq n$, then any allocation such that each agent receives at most one good is EF1. As $k \geq 1$, we know which good is agent 1's favorite, so we will choose such an allocation where agent 1 receives this good. The value that an agent has for the most valuable good is at least $1/m \geq 1/n > 1/3n$ as needed. If $m > n$, we will run the associated EF1 algorithm described in the proof of Theorem 3, and ensure agent 1 is first in the order of the round robin portion of the rule. We claim that if round robin is run for ℓ steps, then agent 1 receives at least $\frac{\ell}{mn}$ value. Since in all the EF1 rules, whenever $m > n$, round robin is run for at least $m/3$ steps, this shows the desired result. To show the above claim, note that agent 1's top ℓ goods account for at least ℓ/m value. Further, in each round of the round robin step, agent 1 receives an item at least as valuable to them as the other $n - 1$ items, so receives at least $1/n$ of their value of the top ℓ goods.

For randomized rules, note that the $\Omega(n)$ distortion follows from Theorem 1. For $O(n)$, we will again split it into two cases, $m \leq n$ and $m > n$. When $m \leq n$, we will pick an assignment of the goods such that each agent receives at most one uniformly at random. This is EF1 as each agent receives at most one good. Further, each agent receives each good with probability $1/n$, so the expected value of an arbitrary agent i is $\sum_{g \in M} 1/n \cdot v_i(g) = 1/n \cdot \sum_{g \in M} v_i(g) = 1/n$. Hence, the expected social welfare is 1, so distortion is at most $O(n)$.

When $m > n$, we will run the associated EF1 rule from Theorem 3 in a uniform way, picking a permutation τ of the agents uniformly at random, and then running round robin with this agent order for the designated number of steps, finally distributing the remaining goods in the specified way. The key here is that, as $m > n$, we will always be running uniform round robin for at least $m/3$ steps. As a straightforward consequence of Proposition 1, this randomized rule will have distortion at most $3n \in O(n)$. \square

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 Ghodsi *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 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 $\frac{1}{\lfloor \frac{m-n+2}{2} \rfloor}$ -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. We begin with the most interesting case, a 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 1st 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 0th 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 $1 + \left\lfloor \frac{d-i}{2H_n(n-i+1)} \right\rfloor$. Therefore, the number of total pairs with second coordinate at most d is at most $\sum_{i=1}^n 1 + \left\lfloor \frac{d-i}{2H_n(n-i+1)} \right\rfloor = n + \sum_{i=1}^n \left\lfloor \frac{d-i}{2H_n(n-i+1)} \right\rfloor$. Our goal is to show that this value is at most d , which is equivalent to the following lemma.

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

Proof. There are two cases, either $d \geq 2n$ or $d < 2n$. First suppose that $d \geq 2n$. Then,

$$\begin{aligned} \sum_{i=1}^n \left\lfloor \frac{d-i}{2H_n(n-i+1)} \right\rfloor &\leq \sum_{i=1}^n \frac{d-i}{2H_n(n-i+1)} \\ &\leq \frac{d}{2H_n} \sum_{i=1}^n \frac{1}{n-i+1} \\ &= \frac{d}{2} \leq d - n. \end{aligned}$$

Next, assume that $n + 1 \leq d < 2n$. Let $r = d - n$. Note that $1 \leq r < n$. The inequality in the lemma is therefore equivalent to (note the cancelling $+1$ and -1 in the numerator) $\sum_{i=1}^n \left\lfloor \frac{r+(n-i+1)-1}{2H_n(n-i+1)} \right\rfloor \leq r$. Note that summing over i ranging from 1 through n is equivalent to summing over $n - i + 1$ ranging from 1 through n ; substituting i to denote $n - i + 1$, we need to prove $\sum_{i=1}^n \left\lfloor \frac{r+i-1}{2H_n \cdot i} \right\rfloor \leq r$.

Call the summand value $f(i) = \left\lfloor \frac{r+i-1}{2H_n \cdot i} \right\rfloor$. First, we show that $f(i) \leq r$ for all $i \in [n]$. For this, it is sufficient to show $r + i - 1 \leq 2H_n \cdot i \cdot r$. We have that $r + i - 1 \leq 2 \max(i, r) \leq 2H_n \cdot i \cdot r$ because H_n, i , and r are all at least 1.

Let $g(j)$ be the number of values of i such that $f(i) \geq j$; more formally, $g(j) = \sum_{i=1}^n \mathbb{1}[f(i) \geq j]$. Now since $f(i) \leq r \leq n$, we have that $\sum_{i=1}^n f(i) = \sum_{j=1}^n g(j)$ as each agent i is counted precisely $f(i)$ times in the RHS sum. Next, we bound $g(j)$ for each $j \in [n]$. To do this, we show that if i is such that $f(i) \geq j$, then i must be bounded by some value, B . This would imply that at most B values of i have $f(i) \geq j$, implying $g(j) \leq B$.

$$\begin{aligned} f(i) = \left\lfloor \frac{r+i-1}{2H_n \cdot i} \right\rfloor &\geq j \\ \implies r+i-1 &\geq 2H_n \cdot i \cdot j && (2H_n \cdot i > 0) \\ \implies r-1 &\geq (2H_n j - 1) \cdot i \\ \implies \frac{r-1}{2H_n j - 1} &\geq i. && (2H_n j - 1 > 0 \text{ as } H_n \cdot j \geq 1) \end{aligned}$$

Hence, $g(j) \leq \frac{r-1}{2H_n j - 1}$ for all j . In addition, since $H_n \cdot j \geq 1$, $H_n \cdot j \leq 2H_n \cdot j - 1$, so $\frac{r-1}{2H_n \cdot j - 1} \leq \frac{r-1}{H_n \cdot j} \leq \frac{r}{H_n \cdot j}$.

Plugging this into our earlier equation:

$$\sum_{i=1}^n f(i) = \sum_{j=1}^n g(j) \leq \sum_{j=1}^n \frac{r}{H_n \cdot j} = \frac{r}{H_n} \sum_{j=1}^n \frac{1}{j} = r,$$

as desired. \square

We will now show the corresponding upper-bound for $k \geq n$, showing it is impossible to achieve α -MMS for $\alpha > \frac{k}{H_n(m-n)-(m-k)}$. This is a generalization of a result by Amanatidis *et al.* [2016]. They show when rankings are complete, so $k = m$, no mechanism can guarantee α -MMS for $\alpha > \frac{1}{H_n} + \varepsilon$ for any $\varepsilon > 0$. More specifically, in their

model, such an approximation must hold for all m , so they in fact show it is impossible to achieve $\alpha > \frac{m}{(m-n)H_n}$ (which coincides with our bound when $m = k$), and take the limit as m gets large. Our proof follows a similar structure to theirs with extra complexity due to the generalization.

Suppose all agents agree on the top k rankings, $g_1 \succ \dots \succ g_k$ and do not rank the remaining $m - k$. First, we will show that if a mechanism achieves a positive MMS approximation on this allocation, the top n goods g_1, \dots, g_n must be assigned such that each agent receives exactly one. Suppose this is not the case, which implies there is some agent i who received none of these goods. In this case, it's possible that agent i valued g_1, \dots, g_n at $1/n$ and the remaining goods at zero, having a positive MMS, but receiving zero value. Therefore, without loss of generality, we can relabel the agents such that g_i goes to agent i .

Fix the chosen allocation A . Let r_i be the number of goods received by agent i of the first k , $r_i = |A_i \cap \{g_1, \dots, g_k\}|$, and let q_i be the number of goods received by agent i of the remaining $m - k$, $q_i = |A_i \setminus \{g_1, \dots, g_k\}|$. We have that $\sum_{i=1}^n r_i = k$ and $\sum_{i=1}^n q_i = m - k$. It's consistent with the rankings that agent i has the following valuations: $1/n$ for goods g_1, \dots, g_{i-1} , 0 for the q_i non-top k goods they received, $A_i \setminus \{g_1, \dots, g_k\}$, and equal value $\frac{1}{m-q_i-i+1}$ for the remaining goods. One possible partition for maximin share is to create $i - 1$ singleton bundles with goods g_1, \dots, g_{i-1} and split the remaining $m - q_i - i + 1$ positively valued goods into $n - i + 1$ bundles as equally as possible. One of these bundles must have at most $\left\lfloor \frac{m-q_i-i+1}{n-i+1} \right\rfloor$ goods, however, agent i received only r_i of these items. Hence, if this allocation is α -MMS, it must be the case that $p_i \geq \alpha \cdot \left\lfloor \frac{m-i-q_i+1}{n-i+1} \right\rfloor$. As $\sum_{i=1}^n p_i = k$, we have that $\sum_{i=1}^n \alpha \cdot \left\lfloor \frac{m-i-q_i+1}{n-i+1} \right\rfloor \leq k$. From this, we get

$$\begin{aligned} \alpha &\leq \frac{k}{\sum_{i=1}^n \left\lfloor \frac{m-i-q_i+1}{n-i+1} \right\rfloor} \\ &\leq \frac{k}{\sum_{i=1}^n \frac{m-i-q_i+1}{n-i+1} - 1} \\ &= \frac{k}{\sum_{i=1}^n \frac{m-n-q_i}{n-i+1}} \\ &= \frac{k}{(m-n) \sum_{i=1}^n \frac{1}{n-i+1} - \sum_{i=1}^n \frac{q_i}{n-i+1}} \\ &\leq \frac{k}{(m-n)H_n - \sum_{i=1}^n \frac{q_i}{1}} \\ &= \frac{k}{(m-n)H_n - (m-k)} \end{aligned}$$

as we wanted to show.

Next, let us consider the case of $k < n - 1$, where we will show that no allocations are α -MMS for $\alpha > 0$. Consider a preference profile in which all agents agree on the top k goods, that is, they rank g_1, \dots, g_k , and do not rank the remaining $m - k$. Fix some arbitrary allocation A , we will show A does not satisfy α -MMS for any $\alpha > 0$. As $k < n - 1$, it must be the case that at least $n - k$ agents do not receive any of g_1, \dots, g_k . Let i be one of these agents that among them receives fewest goods, that is, $i \in \operatorname{argmin}_{j: A_j \cap \{g_1, \dots, g_k\} = \emptyset} |A_j|$. We claim that $|A_i| \leq m - n$. To see this, note that all of the at least $n - k$ agents without any of the first k goods g_1, \dots, g_k have bundles of size at least $|A_i|$, so we have that $(n - k)|A_i| + k \leq m$. Using the fact that $k = n - (n - k)$, after moving terms around, we get $|A_i| \leq \frac{m-n}{n-k} + 1$. As $n - k > 1$ and $m - n > 0$, we get that $|A_i| < m - n + 1$, so since $|A_i|$ is an integer, it follows that $|A_i| \leq m - n$. Therefore, it's consistent with the preference profile that agent i has value $1/n$ for n of the goods not in A_i including g_1, \dots, g_k , and zero for all other goods. In this case, MMS_i is $1/n$, but i is receiving zero value from A , so A cannot be α -MMS for any positive α .

Finally, let us now consider the case where $k = n - 1$. We will begin with upper-bounds, showing that no allocation can achieve higher than $1/\left\lfloor \frac{m-n+2}{2} \right\rfloor$ -MMS. We will begin with a lemma.

Lemma 4. *When $k = n - 1$, if agents agree on the top k goods, the only allocations that can satisfy α -MMS for positive α are those where $n - 1$ of the agents get exactly one of g_1, \dots, g_k , and the last agent receives all the remaining goods.*

Proof. Consider an arbitrary allocation A where this is not the case. We will show A cannot satisfy α -MMS for positive α . Since A does not satisfy this, there must be some agent i who receives none of g_1, \dots, g_k and also does

not receive at least one good $g \in M \setminus \{g_1, \dots, g_k\}$. It's consistent with the preference profile that i has value $1/n$ for g_1, \dots, g_k and g , and zero value for the remaining goods. Therefore, MMS_i is $1/n$, but under A , i receive zero value, so A cannot achieve a positive approximation. \square

To show it is impossible to satisfy α -MMS for α larger than $1/\lfloor \frac{m-n+2}{2} \rfloor$, suppose all agents agree on the rankings, so they report $g_1 \succ \dots \succ g_k$ for the first k goods and do not report the remaining $m - k$. Fix an arbitrary allocation A , we will show it cannot satisfy α -MMS for $\alpha > 1/\lfloor \frac{m-n+2}{2} \rfloor$. By Lemma 4, if A does not give $n - 1$ agents exactly one of g_1, \dots, g_k and give the last agent all the remaining goods, it cannot have a positive MMS approximation, so suppose it does. Without loss of generality, suppose agent i receives good g_i for $i \in [n - 1]$ and agent n receives the remaining goods. In particular, consider agent $n - 1$, who received just good g_k . Suppose their valuations were as follows: they have value $1/n$ for items g_1, \dots, g_{k-1} and equal value of $\frac{1}{m-n+2}$ for all the remaining goods. We have that the agent's MMS is at least $\lfloor \frac{m-n+2}{2} \rfloor \cdot \frac{1}{m-n+2}$ as they could create the partition with the first $n - 2$ sets being $\{g_1\}, \dots, \{g_{k-1}\}$ with the last two bundles created by splitting up the remaining goods as equally as possible, with sizes $\lfloor \frac{m-n+2}{2} \rfloor$ and $\lceil \frac{m-n+2}{2} \rceil$. However, the agent only receives value, $\frac{1}{m-n+2}$, so the MMS approximation is at most $1/\lfloor \frac{m-n+2}{2} \rfloor$ -MMS.

Next, we will show it is in fact possible to achieve $1/\lfloor \frac{m-n+2}{2} \rfloor$ -MMS. The mechanism to achieve this will be the following picking sequence, the agents 1 through n followed by all ns , $1, 2, \dots, n, n, \dots, n$. We will show that every agent here receives a $1/\lfloor \frac{m-n+2}{2} \rfloor$ approximation to MMS. First, we will show the n^{th} agent in fact receives an MMS allocation. Indeed, as she receives all but $n - 1$ of the goods, in any partition of the goods into n bundles, at least one of the bundles must be a subset of these $m - (n - 1)$ goods. Hence, the value of them is an upper-bound on the agent's MMS. Next consider some agent $i \in [n - 1]$ that went i^{th} in the order. Suppose their top i favorite goods are g_1, \dots, g_i . Note that they get at least their value for g_i . In any partition of the goods into n bundles, at least $n - i + 1$ of them cannot contain their g_1, \dots, g_{i-1} . Further, as these $n - i + 1$ bundles collectively have at most $m - i + 1$ goods, at least one bundle must have at most $\lfloor \frac{m-i+1}{n-i+1} \rfloor$ goods, all of which are at most as valuable as g_i . Therefore, agent i receives at least a $1/\lfloor \frac{m-i+1}{n-i+1} \rfloor \geq 1/\lfloor \frac{m-(n-1)+1}{n-(n-1)+1} \rfloor = 1/\lfloor \frac{m-n+2}{2} \rfloor$ approximation to their MMS. \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.

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 .

Proof. We will begin with the lower bound for distortion of deterministic rules on instances with $k = n - 1$. Fix some $n \geq 2$ and let $k = n - 1$. We will construct a family of instances with n agents, top- k rankings, increasing number of goods whose distortion grows linearly with the number of goods. This implies the distortion is unbounded. Suppose the agents agree on the ranking of the top k goods $g_1 \succ \dots \succ g_k$ and leave the remaining $m - k$ goods unranked. Consider some allocation A chosen by an α -MMS rule for $\alpha > 0$. By Lemma 4, under A , $n - 1$ agents must receive exactly one of g_1, \dots, g_k . Without loss of generality, suppose agent 1 receives g_1 . Suppose the valuations are such that agent 1 has value $1/m$ for all the goods, while all other agents have value 1 for g_1 and 0 for all others. This means the social welfare is $1/m$ while the highest possible social welfare is 1, so the distortion is at least m . Further, this ratio holds even compared to the best α -MMS allocation, as swapping A_1 and A_2 must remaining α -MMS as all preference rankings are identical, while the social welfare is at least 1.

For randomized rules, however, note that $1/\lfloor \frac{m-n+2}{2} \rfloor$ -MMS can be achieved by a picking sequence as shown in Theorem 5. By Proposition 1, the uniform version of this picking sequence is also $1/\lfloor \frac{m-n+2}{2} \rfloor$ -MMS and has distortion n .

Next, we will consider the case where $k \geq n$. First, deterministic α -MMS rules for $\alpha > 0$ must have distortion $\Omega(n^2)$ is an easy consequence of Theorem 2.

We will now show it's possible to achieve $O(n^2)$ distortion even with an $\frac{k-n+1}{m-n+1} \cdot \frac{1}{2H_n}$ -MMS deterministic rule. The construction is very similar to $k \geq n$ construction of Theorem 5, except that we add one more constraint: we want agent 1 to also have additional appearances in the picking sequence, once at or before position $2nj + 1$ for each j . Further, if we are using top- k rankings, we will give all the goods after the k steps of the picking sequence to the agent 1. Because agent 1 gets the first overall pick as well, this ensures that the utility to agent 1 is at least $1/2n$ fraction of her value for all the goods, i.e., at least $1/2n$. Since n is a trivial upper bound on the social welfare of any allocation, this implies the distortion upper bound of $O(n^2)$.

Now, as in the proof of Theorem 5, we must show that for all d , there are at most d pairs with second coordinate at most d , also counting the additional pairs generated due to the extra constraint above. Since we have not added any pairs with second coordinate at most n , the part of the proof for $d \leq n$ still holds. For $d \geq n + 1$, the number of new pairs added due to the above constraint is $\lfloor \frac{d-1}{2n} \rfloor$. Hence, we need the following strengthening of Lemma 3.

Lemma 5. *For all $n \in \mathbb{N}$ and for all $d \geq n + 1$, $\lfloor \frac{d-1}{2n} \rfloor + \sum_{i=1}^n \left\lfloor \frac{d-i}{2H_n(n-i+1)} \right\rfloor \leq d - n$.*

Proof. First note that when $n = 1$, this is equivalent to showing $\lfloor \frac{d-1}{2} \rfloor + \lfloor \frac{d-1}{2} \rfloor \leq d - 1$ which trivially holds for all $d \in \mathbb{N}$. Therefore, from now on, we restrict to the case of $n \geq 2$.

Similarly to Lemma 3, we consider two cases: $d \geq 4n$ and $d < 4n$. First, suppose $d \geq 4n$. Then,

$$\begin{aligned} \left\lfloor \frac{d-1}{2n} \right\rfloor + \sum_{i=1}^n \left\lfloor \frac{d-i}{2H_n(n-i+1)} \right\rfloor &\leq \frac{d-1}{2n} + \sum_{i=1}^n \frac{d-i}{2H_n(n-i+1)} \\ &\leq \frac{d}{4} + \sum_{i=1}^n \frac{d}{2H_n(n-i+1)} \\ &= \frac{d}{4} + \frac{d}{2H_n} \cdot \sum_{i=1}^n \frac{1}{n-i+1} \\ &= \frac{d}{4} + \frac{d}{2H_n} \cdot H_n \\ &= \frac{3d}{4} \leq d - n, \end{aligned}$$

where the second transition holds because $n \geq 2$, and the final transition holds because $d \geq 4n$.

Next, suppose $d < 4n$. First, we tackle the special case of $n = 2$. Note that when $d \leq 2n$, we have $\lfloor \frac{d-1}{2n} \rfloor = 0$, in which case the desired inequality is already implied by Lemma 3. For $2n < d < 4n$, i.e., $d \in \{5, 6, 7\}$, we have that $\lfloor \frac{d-1}{2n} \rfloor = 1$. Hence, the desired inequality holds as follows:

$$\begin{aligned} d = 5: \quad &1 + \sum_{i=1}^2 \left\lfloor \frac{5-i}{2H_2(2-i+1)} \right\rfloor = 2 \leq 5 - 1 \\ d = 6: \quad &1 + \sum_{i=1}^2 \left\lfloor \frac{6-i}{2H_2(2-i+1)} \right\rfloor = 2 \leq 6 - 1 \\ d = 7: \quad &1 + \sum_{i=1}^2 \left\lfloor \frac{7-i}{2H_2(2-i+1)} \right\rfloor = 3 \leq 7 - 1 \end{aligned}$$

From now on, we restrict to the case of $n \geq 3$. Because we are in the case of $d < 4n$, this implies $\lfloor \frac{d-1}{2n} \rfloor \leq 1$. Hence, it is sufficient to show that

$$\sum_{i=1}^n \left\lfloor \frac{d-i}{2H_n(n-i+1)} \right\rfloor \leq d - n - 1.$$

As in Lemma 3, letting $r = d - n$, the above statement is equivalent to

$$\sum_{i=1}^n \left\lfloor \frac{r+i-1}{2H_n \cdot i} \right\rfloor \leq r - 1.$$

As in Lemma 3, define $f(i) = \left\lfloor \frac{r+i-1}{2H_n \cdot i} \right\rfloor$ and $g(j) = \sum_{i=1}^n \mathbb{1}[f(i) \geq j]$. In Lemma 3, we had $d < 2n$, giving us $f(i) \leq r \leq n$ for all $i \in [n]$. In this case, we still have that $f(i) \leq r$ for all $i \in [n]$, but since we only have $d < 4n$, we have $r \leq 3n$. We note that $g(j) \leq \frac{r-1}{2H_n j - 1}$ can still be derived as in the proof of Lemma 3. Thus, we have that

$$\sum_{i=1}^n f(i) = \sum_{j=1}^{3n} g(j) \leq \sum_{j=1}^{3n} \frac{r-1}{2H_n j - 1} = (r-1) \cdot \sum_{j=1}^{3n} \frac{1}{2H_n j - 1}.$$

To complete the proof, we simply need to show that $\sum_{j=1}^{3n} \frac{1}{2H_n j - 1} \leq 1$ for all $n \geq 3$. When $n = 3$, we have

$$\sum_{j=1}^{3n} \frac{1}{2H_n j - 1} = \sum_{j=1}^9 \frac{1}{2H_3 j - 1} = \frac{19,657,653,727}{21,402,806,880} \leq 1.$$

Suppose $n \geq 4$. First, note that

$$\sum_{j=1}^{3n} \frac{1}{2H_n j - 1} \leq \sum_{j=1}^{3n} \frac{1}{(2H_n - 1)j} = \frac{1}{2H_n - 1} \sum_{j=1}^{3n} \frac{1}{j} = \frac{H_{3n}}{2H_n - 1},$$

where the first transition holds because $j \geq 1$. To show that this final quantity is at most 1, we need to show $H_{3n} \leq 2H_n - 1$ for all $n \geq 4$. We prove this by induction on n .

For the base case of $n = 4$, we have that $H_{12} = \frac{86,021}{27,720} \leq \frac{19}{6} = 2H_4 - 1$. Suppose $H_{3n} \leq 2H_n - 1$ holds for some $n \geq 4$. We need to show that $H_{3(n+1)} \leq 2H_{n+1} - 1$. We have that

$$\begin{aligned} H_{3(n+1)} &= H_{3n} + \frac{1}{3n+1} + \frac{1}{3n+2} + \frac{1}{3n+3} \\ &\leq H_{3n} + \frac{3}{3n} \\ &= 2H_n - 1 + \frac{1}{n} \\ &\leq 2H_n - 1 + \frac{2}{n+1} \\ &\leq 2H_{n+1} - 1 \end{aligned} \quad \square$$

Given Lemma 5, the result for deterministic rules follows.

Finally, for randomized rules, note that the original rule of theorem 6 was a picking sequence, thus the uniform version has distortion n by Proposition 1. \square

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, 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.

Acknowledgements

Nisarg Shah was partially funded by an NSERC Discovery Grant.

References

- 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.
- 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.
- E. Anshelevich, O. Bhardwaj, E. Elkind, J. Postl, and P. Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.
- H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.
- 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.
- 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.
- 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.
- 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.
- 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.
- X. Bei, X. Lu, P. Manurangsi, and W. Suksompong. The price of fairness for indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 81–87, 2019.
- 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.
- 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.
- 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.
- 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.
- S. J. Brams, P. H. Edelman, and P. C. Fishburn. Fair division of indivisible items. *Theory and Decision*, 55(2):147–180, 2003.
- E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 305–322, 2016.
- 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.
- 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.
- 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.

- R. Freeman, N. Shah, and R. Vaish. Best of both worlds: Ex-ante and ex-post fairness in resource allocation. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 21–22, 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.
- M. Ghodsi, 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.
- 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.
- D. Herreiner and C. Puppe. A simple procedure for finding equitable allocations of indivisible goods. *Social Choice and Welfare*, 19(2):415–430, 2002.
- D. Kempe. Communication, distortion, and randomness in metric voting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2087–2094, 2020.
- D. Kurokawa, A. D. Procaccia, and J. Wang. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM*, 65(2):1–27, 2018.
- 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.
- 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.
- 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.
- 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.
- 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.
- H. Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.

Appendix

A Impossible Fairness Properties

Finally, we show that some fairness properties studied in the literature are impossible to guarantee given only the ordinal information, even with complete rankings. Consider an instance $I = (N, M, k, \sigma)$ with valuations v .

Definition 4. An allocation A is said to be *envy-free up to any good (EFX)* if for every pair of agents i, j and every good $g \in A_j$, we have $v_i(A_i) \geq v_i(A_j \setminus \{g\})$.

Definition 5. An allocation A is said to be *equitable up to one good (EQ1)* if for every pair of agents i, j such that $A_j \neq \emptyset$, there exists some good $g \in A_j$ such that $v_i(A_i) \geq v_j(A_j \setminus \{g\})$. A is said to be *equitable up to any good (EQX)* if for every pair of agents i, j such that $A_j \neq \emptyset$ and for every good $g \in A_j$ such that $v_j(g) > 0$, we have that $v_i(A_i) \geq v_j(A_j \setminus \{g\})$.

Proposition 2. *There does not exist an ordinal allocation rule satisfying EFX even for two agents with access to complete rankings.*

Proof. Suppose there are two agents and four goods (denoted g_1, \dots, g_4 for notational convenience). Consider the preference profile wherein both agents have the same preference ranking over the goods, given by $g_1 \succ g_2 \succ g_3 \succ g_4$.

We show that regardless of the allocation chosen by a deterministic ordinal allocation rule, EFX cannot be satisfied with respect to every valuation profile consistent with this preference profile.

Suppose the rule picks allocation $A = (A_1, A_2)$. There are two cases: $|A_1| \neq |A_2|$ and $|A_1| = |A_2|$.

First, suppose $|A_1| \neq |A_2|$, and without loss of generality, suppose $|A_1| > |A_2|$. Since there are four goods, this implies that $|A_1| \geq |A_2| + 2$. In this case, if agent 2 has equal value of $1/4$ for all four goods, then agent 2 would envy agent 1 even after removing any single good from A_1 , violating EFX (and in fact, EF1 too).

Next, suppose that $|A_1| = |A_2|$, which implies that each is equal to 2. Without loss of generality, suppose $g_1 \in A_1$. Then, consider the valuation function of agent 2 which places value 1 on good g_1 and $\epsilon < 1/3$ on every other good. Then, agent 2 would envy agent 1 even after removing the single good in $A_1 \setminus \{g_1\}$, violating EFX. \square

Proposition 3. *There does not exist an ordinal allocation rule satisfying EQ1 or EQX even for two agents with access to complete rankings.*

Proof. Consider the construction used in the proof of Proposition 2, in which two agents have the same preference ranking over four goods given by $g_1 \succ g_2 \succ g_3 \succ g_4$.

We show that regardless of the allocation chosen by a deterministic ordinal allocation rule, EQ1 cannot be satisfied with respect to all valuation profiles consistent with this preference profile. This implies the desired result for EQX as well.

Suppose the rule picks allocation $A = (A_1, A_2)$. There are two cases: $|A_1| \neq |A_2|$ and $|A_1| = |A_2|$.

First, suppose $|A_1| \neq |A_2|$, and without loss of generality suppose $|A_1| > |A_2|$. Since there are four goods, this implies that $|A_1| \geq |A_2| + 2$. Now, consider the valuation profile in which both agents have an equal value of $1/4$ for all the goods. Then, even after removing any single good from A_1 , the utility to agent 1 will be more than the utility to agent 2, violating EQ1.

Next, suppose $|A_1| = |A_2|$, which implies that each is equal to 2. Without loss of generality, suppose $g_1 \in A_1$. Consider the valuation profile in which agent 2 has value 1 for g_1 and 0 for every other good, while agent 1 has value $1/4$ for all the goods. Thus, the utility to agent 2 is 0, while the utility to agent 1 remains higher than that, even after removing any single good from A_1 , violating EQ1. \square