

Possible Fairness for Allocating Indivisible Resources

Haris Aziz
UNSW Sydney
Sydney, Australia
haziz@cse.unsw.edu.au

Shiji Xing
Hong Kong Polytechnic University
Hong Kong, China
shiji.xing@connect.polyu.hk

Bo Li
Hong Kong Polytechnic University
Hong Kong, China
comp-bo.li@polyu.edu.hk

Yu Zhou
Hong Kong Polytechnic University
Hong Kong, China
csyzhou@comp.polyu.edu.hk

ABSTRACT

Fair division of indivisible resources has attracted significant attention from multi-agent systems and computational social choice. Two popular solution concepts are envy-freeness up to any item (EFX) and maximin share (MMS) fairness which are defined using agents' cardinal preferences. On one hand, accurate cardinal values are hard to express in real-life applications, and on the other hand, with cardinal values, MMS and EFX may not be easy to satisfy. In this work, we study a new setting where agents have arbitrary ordinal preferences for the items (possibly with indifferences), and an allocation is called possible EFX (p-EFX) or possible MMS (p-MMS) if there exist cardinal preferences that are consistent with the ordinal ones so that the allocation is EFX or MMS.

We first design a polynomial-time algorithm to compute an allocation that is p-EFX and p-MMS under lexicographic preferences. This result also strengthens a result of Hosseini et al. (AAAI 2021) who proved the existence of EFX and MMS allocations under strict lexicographic preferences (i.e., the items do not have ties). Although it has been well justified that lexicographic preferences are natural and common, there are situations where they do not fit appropriately, especially when the items have similar types. Therefore, on top of p-EFX and p-MMS, we want the allocation to be balanced (i.e., the numbers of items allocated to the agents differ by at most one). We then design another algorithm that satisfies p-EFX, p-MMS, and balanced simultaneously.

KEYWORDS

Indivisible Resources; Ordinal Preferences; Possible Fairness

ACM Reference Format:

Haris Aziz, Bo Li, Shiji Xing, and Yu Zhou. 2023. Possible Fairness for Allocating Indivisible Resources. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

Fairly allocating a set of indivisible resources among a number of heterogeneous agents has been within the center of multi-agent systems [2, 9]. Basically, there are two classes of fairness criteria, namely envy-based and share-based. An envy-based criterion evaluates an agent's fairness on an allocation against those of her peers,

and a share-based one evaluates it against the agent's due share. Two representative notions in these two classes are *envy-freeness* (EF) and *proportionality* (PROP). Informally, an envy-free allocation requires no agent to prefer the bundle of resources allocated to any other agent more than her own, and a proportional one requires every agent to have a value no smaller than her average value for all resources. It is widely known that these two requirements are too strong in the sense that such fair allocations seldom exist due to the indivisibility of the resources. Instead, the literature mostly focuses on their relaxation. The first relaxation is *envy-freeness up to any item* (EFX) [17], which means the envy of an agent towards another agent can be eliminated if an arbitrary resource is removed from the envied agent's bundle. The second one is *maximin share* (MMS) fairness [15], which is motivated by an imaginary experiment where an agent is to divide all the resources into n bundles but is the last one to select a bundle (n is the number of agents). The agent's best strategy, in the worst case, is to maximize the minimum value of all bundles, and this value is named her maximin share. Then an MMS allocation is defined so that every agent's value is no smaller than her MMS.

Unfortunately, although a significant amount of effort has been spent, whether an EFX allocation always exists or not is still unknown. It is shown in [21, 28] that an MMS allocation is not guaranteed to exist even when there are only three agents. Thereafter, a growing body of papers works on finding the approximations of EFX and MMS allocations, which will be discussed in Section 1.2. In all these works, it is assumed that every agent has a cardinal valuation function that assigns an exact numerical value for each set of resources. However, in many real-life scenarios, it is easy for the agents to tell which one is better among the items, but it may be hard to quantify how much it is better off.

In this work, we take an epistemic perspective to study the fair allocation of indivisible resources and provides an alternative way to approximate EFX and MMS fairness. Pragmatically, in our model, we assume each agent has an ordinal ranking over the resources. We then define *possible EFX* and *possible MMS* allocations, which requires that for each agent there exists a cardinal valuation that is consistent with the ordinal preference so that the allocation is EFX and MMS. Konczak and Lang [27] first proposed the notion of possible EF and PROP in fair allocation by Aziz et al. [8] and Segal-Halevi et al. [33]. A simple example shows that possible EF and PROP are still hard to satisfy; see Example 1 and formal definitions are introduced in Section 2. Our definitions are weaker and

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

we prove that we can ensure possible EFX and possible MMS simultaneously by a unified algorithm for arbitrary ordinal rankings even when the agents have indifference over the items.

EXAMPLE 1. Consider an instance with two agents and two resources o and o' . Both agents strictly prefer o over o' . In that case, the agent who gets o' is envious of the other agent and does not get half of the total utility he would get if he got both resources. Therefore, allocating each agent one item is possible EFX and possible MMS, but there exists no possible EF or possible PROP allocation.

Besides [8, 33] and our work, the fair allocation has been investigated from an epistemic perspective in the literature, but most of the existing study focuses on the knowledge of allocations. To the best of our knowledge, Aziz et al. [7] were the first to study the epistemic envy-freeness, which only requires each agent to know her own allocated items. An allocation is regarded as epistemic EF to an agent if there exists a re-allocation of the items not allocated to her among the other agents so that the resulting allocation is EF. Very recently, Garg and Sharma [22] extended this notion to epistemic EFX and proved the existence of such an allocation. In another direction, Chan et al. [18] proposed *maximin aware* allocations, which requires that for each agent, no matter how the items not allocated to her are distributed among the others, there exists one agent whose bundle is no better than hers. Our work is fundamentally different from these works in the sense that the knowledge in our setting is about valuations instead of allocations.

1.1 Our Contribution

In this work, we study the epistemic fair allocation problem of indivisible resources, where the agents express their ordinal preferences over the resources. We propose two *possible fairness* definitions, namely, possible EFX (p-EFX) and possible MMS (p-MMS). Informally, an allocation is p-EFX or p-MMS if for every agent there exists a cardinal additive valuation that is consistent with her ordinal preference so that the allocation is EFX or MMS to her. We justify the two concepts by the following two facts. First, it is usually unrealistic for the agents to express their accurate cardinal values or the agents simply do not have them; Second, as we show in this work, although possible fairness is weaker than the cardinal ones, they are guaranteed to be satisfiable and are compatible with extra properties such as balancedness and Pareto optimality (PO).

To prove the existence of possible fair allocations, we first show that an EFX, MMS, and PO allocation exists when the valuations are weak lexicographic. Hosseini et al. [26] proved the existence of such an allocation for strict lexicographic preferences and left the weak case open. Our result answers this question affirmatively.

Main Result 1. A simultaneously EFX, MMS and PO allocation exists and can be computed in polynomial time for arbitrary (weak) lexicographic preferences.

An interesting distinction between strict and weak lexicographic preferences is that EFX implies MMS in the former case but not in the latter one. As we will see, one of the difficulties in proving the above result is to design a unified algorithm to compute an allocation that is both EFX and MMS.

The existence of fair allocation under lexicographic preferences implies the existence of the corresponding possibly fair allocation.

It has been well justified that lexicographic preferences are natural and common in many fields such as psychology, computer science, and economics[32], but there are situations where they do not fit appropriately, especially when the items have similar types. Therefore, to justify the fairness of possibly fair allocations, an extra property we may want to satisfy is *balancedness*. An allocation being balanced means the numbers of items allocated to the agents differ by at most one so that the agents have almost the same number of items. Note that with lexicographic preferences, EFX and MMS are not compatible with balancedness, but we show that we can compute in polynomial time an allocation that is p-EFX, p-MMS, and balanced. Moreover, the valuation profiles that give EFX and MMS are the same, and under this valuation profile, the allocation is also PO.

Main Result 2. A simultaneously p-EFX, p-MMS, and balanced allocation exists and can be computed in polynomial time. Moreover, p-EFX and p-MMS use the same valuation profile under which the allocation is PO.

1.2 Other Related Works

Since envy-free or proportional allocations seldom exist when the resources are indivisible, a significant amount of effort has been spent on their approximations. Two of the most notable relaxations of envy-freeness are envy-freeness up to one item (EF1) [15] and EFX that we have discussed [17]. EF1 is much weaker than EFX where envy between two agents can be eliminated after the removal of some (instead of “any” as in EFX) item. Compared with EFX, an EF1 allocation is ensured to exist even when the valuations are combinatorial and monotone [29]. Although the existence of EFX allocations is still unknown, there are constant approximations [5, 30]. Caragiannis et al. [16], Chaudhury et al. [20] and Chaudhury et al. [19] also showed that an EFX allocation exists if we allow a small number of items not to be allocated.

we can donate a small number of items to a charity.

MMS is one of the most widely studied relaxations of proportionality [15]. Kurokawa et al. [28] gave the first instance for which no allocation is MMS. Recently, a constant impossibility result is proven by Feige et al. [21]. On the positive side, Kurokawa et al. [28] gave the first $2/3$ -approximate MMS fair algorithm, albeit not in polynomial time. Later Amanatidis et al. [4], Barman and Krishnamurthy [11] provided polynomial-time algorithms with the same approximation ratio. The approximation ratio is further improved to $3/4$ in [24] and $3/4 + 1/(12n)$ in [23].

In another direction, using ordinal preferences to approximate MMS allocations has been investigated in [3, 25], but in these works, the agents are still equipped with cardinal valuation functions but the algorithm does not know. More ordinal fairness notions can be found in [12–14, 31].

2 PRELIMINARIES

2.1 Model

For any integer $k \geq 1$, we denote $[k] = \{1, \dots, k\}$. A resource allocation instance is a triple $I = (N, M, \succ)$ where $N = [n]$ is a set of agents, $M = [m]$ is a set of items, and $\succ = (\succ_1, \dots, \succ_n)$ specifies for each agent i a preference \succ_i over M . For any two items a and

$b, a \succeq_i b$ means agent i (weakly) prefers a to b . The agents may be indifferent among items. Agent i is indifferent between a and b , denoted by $a \sim_i b$, if $a \succeq_i b$ and $b \succeq_i a$. Agent i strictly prefers a to b , denoted by $a \succ_i b$, if $a \succeq_i b$ and $a \not\sim_i b$. The preference relation of each agent $i \in N$ can be expressed by k_i equivalence classes $E_i^1, \dots, E_i^{k_i}$ (in decreasing order of preferences), where each set E_i^j is a maximal equivalence class of items for which agent i is indifferent, and k_i is the number of equivalence classes of agent i . For any two items $a \in E_i^j$ and $b \in E_i^k$, $a \succeq_i b$ implies that $j \leq k$, $a \sim_i b$ implies that $j = k$, and $a \succ_i b$ implies that $j < k$. A bundle is a subset of items $X \subseteq M$. We sometimes denote a bundle X by a vector $(a_1, \dots, a_{k_i})_{\succeq_i}$ if X contains $a_j \in \mathbb{N}$ items in E_i^j for $j \in [k_i]$. An allocation $\mathbf{A} = (A_1, \dots, A_n)$ is an n -partition of M if $A_1 \cup \dots \cup A_n = M$ and $A_i \cap A_j = \emptyset$ for any $i \neq j$, where A_i is the bundle allocated to agent i . An allocation \mathbf{A} is called *partial* if $A_1 \cup \dots \cup A_n \subset M$. An allocation $\mathbf{A} = (A_1, \dots, A_n)$ is *balanced* if for any $i, j \in N$, $\|A_i\| - \|A_j\| \leq 1$.

One way to extend preferences over items to preferences over bundles of items is via *downward lexicographic dominance* (DL) (see, e.g., [10]).

DEFINITION 1 (DL-PREFERENCE). *For an agent $i \in N$, given two bundles $A = (a_1, \dots, a_{k_i})_{\succeq_i}$ and $B = (b_1, \dots, b_{k_i})_{\succeq_i}$, we say that agent i DL-prefers A to B , or $A \succeq_i^{DL} B$, if either of the following two conditions holds:*

- *there exists $l \leq k_i$ such that $a_l > b_l$ and $a_j \geq b_j$ for any $j < l$.*
- *$a_j = b_j$ for any $j \in [k_i]$.*

Moreover, we say that agent i strictly DL-prefers A to B , denoted by $A \succ_i^{DL} B$, if $A \succeq_i^{DL} B$ and $B \not\sim_i^{DL} A$; agent i is DL-indifferent between A and B , denoted by $A \sim_i^{DL} B$, if $A \succeq_i^{DL} B$ and $B \succeq_i^{DL} A$.

DL-preferences are also abbreviated as lexicographic preferences in this paper.

Although agents only express ordinal preferences, they may have cardinal valuations $u_i : 2^M \rightarrow \mathbb{R}^+$. A valuation u_i is consistent with \succeq_i if $u_i(g) \geq u_i(g') \Leftrightarrow g \succeq_i g'$. The set of all cardinal valuations consistent with \succeq_i is denoted by $\mathcal{U}(\succeq_i)$. When we consider agents' cardinal valuations, it is assumed they are additive; that is $u_i(X) = \sum_{g \in X} u_i(g)$ for all $i \in N$ and $X \subseteq M$. Note that DL preferences can be expressed by additive cardinal valuations (as will be shown in the proof of Proposition 2).

2.2 Fairness Notions

Next, we introduce the fairness notions we are interested in. Note that these concepts are defined using cardinal valuations. An allocation $\mathbf{A} = (A_1, \dots, A_n)$ is envy-free (EF) if $u_i(A_i) \geq u_i(A_j)$ for all agents i and j . Since EF is too hard to satisfy, we consider its relaxation, envy-free up to any item. An allocation is *envy-free up to any item* (EFX) if for all agents i and j

$$u_i(A_i) \geq u_i(A_j \setminus \{g\}) \text{ for any } g \in A_j.$$

A weaker notion than EF is proportionality. An allocation $\mathbf{A} = (A_1, \dots, A_n)$ is proportional (PROP) if $u_i(A_i) \geq \frac{1}{n} \cdot u_i(M)$ for any agent i . One popular relaxation of PROP is maximin share (MMS) fairness. Let $\Pi_n(M)$ be the set of all n -partitions of M , agent i 's

maximin share $\text{MMS}_i^n(u_i)$ is

$$\text{MMS}_i^n(u_i) = \max_{X=(X_1, \dots, X_n) \in \Pi_n(M)} \min_{j \in [n]} u_i(X_j).$$

Given an n -partition (X_1, \dots, X_n) of M , we say it is an MMS-defining partition for agent i if $u_i(X_j) \geq \text{MMS}_i^n(u_i)$ for all $j \in [n]$. An allocation $\mathbf{A} = (A_1, \dots, A_n)$ satisfies *MMS fairness* if each agent $i \in N$ gets utility at least $\text{MMS}_i^n(u_i)$, i.e.,

$$u_i(A_i) \geq \text{MMS}_i^n(u_i).$$

An allocation is called Pareto optimal (PO) if there is no alternative allocation which strictly improves one agent without hurting any agent. Formally, $\mathbf{A} = (A_1, \dots, A_n)$ is Pareto-dominated by $\mathbf{B} = (B_1, \dots, B_n)$ if $u_i(A_i) \leq u_i(B_i)$ for all i and the inequality is strict for some i . An allocation \mathbf{A} is PO if it is not Pareto-dominated by any allocation.

In this work, we consider the ordinal versions of these fairness notions. The ordinal concepts are defined in the same way as possible EF and possible PROP that have been defined by Aziz et al. [8]. An allocation $\mathbf{A} = (A_1, \dots, A_n)$ is possible EF if for each $i \in N$, there exists $u_i \in \mathcal{U}(\succeq_i)$ such that for all $j \in N$, $u_i(A_i) \geq u_i(A_j)$; it is possible PROP if for each $i \in N$, there exists $u_i \in \mathcal{U}(\succeq_i)$ such that $u_i(A_i) \geq \frac{1}{n} u_i(M)$. As shown by Example 1, a possible EF or PROP allocation seldom exists. In the following, we adapt these notions to possible EFX and possible MMS, which will be shown to be always satisfiable.

DEFINITION 2 (p -EFX). *An allocation $\mathbf{A} = (A_1, \dots, A_n)$ is possible envy-free up to any item (p -EFX) if for each $i \in N$, there exists $u_i \in \mathcal{U}(\succeq_i)$ such that for any $j \in N$,*

$$u_i(A_i) \geq u_i(A_j \setminus \{g\}) \text{ for any } g \in A_j.$$

DEFINITION 3 (p -MMS). *An allocation $\mathbf{A} = (A_1, \dots, A_n)$ is possible maximin share (p -MMS) fair if for each $i \in N$, there exists $u_i \in \mathcal{U}(\succeq_i)$ such that*

$$u_i(A_i) \geq \text{MMS}_i^n(u_i).$$

At the end of this paper, we will introduce a stronger requirement than possible fairness, namely, necessary fairness, which requires the allocation to be fair under all cardinal valuations that are consistent with the given ordinal preferences. This definition is hard to satisfy and we provide a brief discussion in Section 5.

DEFINITION 4 (n -EFX). *An allocation $\mathbf{A} = (A_1, \dots, A_n)$ is necessary envy-free up to any item (n -EFX) if for each $i \in N$, for any $u_i \in \mathcal{U}(\succeq_i)$ and for any $j \in N$,*

$$u_i(A_i) \geq u_i(A_j \setminus \{g\}) \text{ for any } g \in A_j.$$

DEFINITION 5 (n -MMS). *An allocation $\mathbf{A} = (A_1, \dots, A_n)$ is necessary maximin share (n -MMS) fair if for each $i \in N$ and for any $u_i \in \mathcal{U}(\succeq_i)$, it holds that*

$$u_i(A_i) \geq \text{MMS}_i^n(u_i).$$

3 COMPUTING FAIR ALLOCATIONS UNDER WEAK LEXICOGRAPHIC PREFERENCES

In this section, we prove the existence of EFX and MMS allocations under (weak) lexicographic preferences, which also implies the existence of p -EFX and p -MMS allocations.

3.1 DL-EFX, DL-MMS and DL-PO

The requirements of EFX, MMS and PO can be equivalently stated as the following DL-EFX, DL-MMS and DL-PO.

DEFINITION 6 (DL-EFX). Let A_i and A_j be the bundles of agents i and j , respectively. We say that i DL-envies j if $A_j \succ_i^{DL} A_i$. An allocation $\mathbf{A} = \{A_1, \dots, A_n\}$ is DL-envy free up to any item (DL-EFX) if for any $i, j \in N$,

$$A_i \succeq_i^{DL} A_j \setminus \{g\} \text{ for any } g \in A_j.$$

The DL-maximin share (DL-MMS) of an agent is the most DL-preferred bundle the agent can guarantee herself if she is to partition M into n bundles while knowing that she can only get the bundle she least DL-prefers. Formally, agent i 's DL-MMS is defined as

$$\text{DLMMS}_i^n(\succeq_i) = \text{Max}_{X \in \Pi_n(M)} \text{Min}\{X_1, \dots, X_n\},$$

where $\text{Max}\{\cdot\}$ and $\text{Min}\{\cdot\}$ denote the most DL-preferred and the least DL-preferred bundles to agent i . Note that we write DLMMS_i^n for simplicity when \succeq_i is clear from the context. We want to highlight that DLMMS_i^n actually represents a bundle of items instead of a cardinal value.

DEFINITION 7 (DL-MMS). An allocation $\mathbf{A} = (A_1, \dots, A_n)$ satisfies DL-MMS fairness if each agent $i \in N$ gets a bundle that she (weakly) DL-prefers to DLMMS_i^n , i.e.,

$$A_i \succeq_i^{DL} \text{DLMMS}_i^n \text{ for any } i \in N.$$

Given an n -partition (X_1, \dots, X_n) of M , we say it is a DL-MMS-defining partition for agent i if

$$X_j \succeq_i^{DL} \text{DLMMS}_i^n \text{ for all } j \in [n].$$

DEFINITION 8 (DL-PO). Given two allocations $\mathbf{A} = (A_1, \dots, A_n)$ and $\mathbf{B} = (B_1, \dots, B_n)$, we say \mathbf{B} DL-Pareto dominates \mathbf{A} if $B_i \succeq_i^{DL} A_i$ for every $i \in N$ and the inequality is strict for some $i \in N$. Moreover, we say an allocation is DL-Pareto optimal (DL-PO) if it is not DL-Pareto dominated by any other allocation.

Note that possible PO is equivalent to DL-PO. Actually, it can be proven that if there exists a profile of valuations such that the allocation is PO, it is also PO under the lexicographic valuation [6].

It is shown by Hosseini et al. [26] that with strict preferences, a DL-EFX allocation is also DL-MMS. Thus the existence of DL-EFX allocations implies that of DL-MMS allocations. However, when the preferences have ties, this is not true any more.

PROPOSITION 1. A DL-EFX allocation may not be DL-MMS if the preferences contain indifferences.

PROOF. Consider an instance with $N = \{1, 2, 3\}$ and $M = \{a, b, c, d, e, f\}$, the agents' preferences over the items are defined as follows:

- Agent 1: $(a \sim_1 b \sim_1 c) \succ_1 (d \sim_1 e \sim_1 f)$.
- Agent 2: $(a \sim_2 b) \succ_2 (c \sim_2 d \sim_2 e \sim_2 f)$.
- Agent 3: $(a \sim_3 b \sim_3 c \sim_3 d \sim_3 e \sim_3 f)$.

It is easy to verify that the allocation $\mathbf{A} = \{A_1 = \{c\}, A_2 = \{a, b\}, A_3 = \{d, e, f\}\}$ is DL-EFX. However, the MMS bundle for agent 1 contains one item from $\{a, b, c\}$ and one item from $\{d, e, f\}$, which agent 1 strictly DL-prefers to A_1 . Hence, \mathbf{A} is not DL-MMS. \square

It deserves noting that DL-EFX and DL-MMS are not compatible with balancedness. A simple example to see this is when the preferences are strict and identical. Letting $m = 2n$, a balanced allocation requires that every agent gets two items. However, in any DL-EFX allocation, the agent who gets the most favorite item cannot get any other item any more; in any DL-MMS allocation, the agent who does not get one of the top $n - 1$ items must get all the other items.

Finally, it is not hard to prove that DL fairness implies possible fairness, which is proven in the full paper.

PROPOSITION 2. A DL-EFX allocation is p -EFX. Similarly, a DL-MMS allocation is p -MMS.

3.2 Computing DL-MMS

Before presenting our first main algorithm, we first show that DL-MMS can be computed efficiently, which is contrary to MMS defined under cardinal valuations. Intuitively, we want to allocate the items in every equivalence class evenly to n bundles. However, this is usually not possible. Considering items in the best equivalence class, if they cannot be evenly allocated, we should only allocate the items in the other classes to the bundles that contain fewer items in the best class since they are at a disadvantage. Similarly, if the items in the second-best equivalence class cannot be evenly allocated to these disadvantageous bundles, we only allocate the other items to the bundles that are at a disadvantage in this second round. Continuing this, we obtain the DL-MMS bundle.

For any ordinal preference \succeq_i , recall that it can be expressed by k_i equivalence classes $E_i^1, \dots, E_i^{k_i}$. Letting

- $n_1 = n$ and if $k_i \geq 2$,
- $n_j = n_{j-1} - (|E_i^{j-1}| \bmod n_{j-1})$ for every $j \in [2, k_i]$,

we have the following lemma, whose proof is in the full paper.

LEMMA 1. DLMMS_i^n contains $\lfloor \frac{|E_i^j|}{n_j} \rfloor$ items in E_i^j for every $j \in [k_i]$.

3.3 The Algorithm

We then introduce some necessary notions and technical lemmas used in our algorithm. To visualize our algorithm, we define *exchange graph*, which is similar to the envy graph defined in [1, 6]. At each step of our algorithm where the (partial) allocation is $\mathbf{A} = (A_1, \dots, A_n)$, we draw the exchange graph $G(\mathbf{A})$. Each item $g \in M$ is a vertex in $G(\mathbf{A})$ and for any two items $g_1, g_2 \in M$, there is a directed edge from g_1 to g_2 if $g_1 \in A_i$ for some $i \in N$ and $g_2 \succeq_i g_1$. We say that an item g is (1) *unallocated* if $g \notin A_j$ for all $j \in N$, (2) *exchangeable* if $g \in A_j$ for some $j \in N$ and there exists a path from g to some unallocated item in $G(\mathbf{A})$ (we call such a path an *exchange path*), and (3) *finalized* if $g \in A_j$ for some $j \in N$ but there does not exist an exchange path from g to any unallocated item.

For any exchange path $P = v_1 \rightarrow \dots \rightarrow v_k$, we denote by i_j the agent who owns the item v_j for every $j \in [k - 1]$, then the *path-exchange allocation* $\mathbf{A}^P = (A_1^P, \dots, A_n^P)$ is obtained by exchanging items backwards along the path. That is, we let $A_i^P = A_i$ for every $i \in N$, then

$$A_{i_j}^P = A_{i_j}^P \setminus \{v_j\} \cup \{v_{j+1}\} \text{ for every } j \in [k - 1].$$

Note that the items that the agents get by exchange are at least as preferred as the items they owned before.

Once an item becomes finalized, its owner will never exchange it. We have the following lemmas regarding finalized items, whose proofs are in the full paper.

LEMMA 2. *Given a (partial) allocation $\mathbf{A} = (A_1, \dots, A_n)$ and a finalized item $g \in A_i$, $g \succ_i g'$ holds for any unallocated or exchangeable item g' . In other words, any item g' such that $g' \succeq_i g$ is finalized.*

LEMMA 3. *Given a (partial) allocation $\mathbf{A} = (A_1, \dots, A_n)$ and a finalized item g , any item that can be reached by g in $G(\mathbf{A})$ is also finalized.*

Our algorithm is formally described in Algorithm 1. Note that the parameter t limits the number of rounds (Steps 5 to 26) and is only considered in Section 4. For now, we can assume that t is infinity. In each round, each of the remaining agents picks one item in an arbitrary order in the way described in Algorithm 2; each agent i picks her favorite item g among the unallocated or exchangeable items. If g is exchangeable before i picks it, which means that there exists at least one exchange path from g to an unallocated item, the agents who own the items on one of such paths exchange their items according to the path-exchange allocation. At the end of each round, two sets of agents are removed from the algorithm:

- The first set contains all agents who get an item which another agent strictly prefers to the item she gets herself in that round.
- The second set contains all agents who are not in the first set but get an item that can be reached in the exchange graph by the items allocated to the first set of agents in that round.

The first set of agents are removed to ensure that they are not DL-envied by the agents who strictly prefer their items after removing any one of their items. The second set of agents are removed to ensure that they are not DL-envied by the agents in the first set after removing any one of their items. Note that another interpretation of the second set that this set contains all agents who get an item which some agent who is to be removed (weakly) prefers to the item she gets herself in that round.

We first have the following observations.

OBSERVATION 1. *In any round, the agent who lastly picks an item is not removed.*

PROOF. We consider round l and denote by T^l the agents who get an item in round l , by $i \in T^l$ the agent who lastly picks an item in round l , by g_j^l the item that any agent $j \in T^l$ gets in round l .

For the sake of contradiction, we suppose that i is removed at the end of round l . Since i is the last to pick an item, $g_j^l \succeq_j g_i^l$ holds for any agent $j \in T^l$, thus i is not in the first set of removed agents. Then i must be in the second set, which means that there exists an agent $j \in T^l$ who is removed in the first set and whose item g_j^l can reach g_i^l in the exchange graph.

The fact that j is in the first set of removed agents means there exists an agent $k \in T^l$ such that $g_j^l \succ_k g_k^l$. This implies that g_j^l is finalized when i picks an item, since otherwise, agent k should have picked g_j^l and got an item that she strictly prefers to g_k^l in round l . By Lemma 3, all items that can be reached by g_j^l are also finalized

Algorithm 1: Computing a DL-MMS, DL-EFX and DL-PO allocation.

```

1 Input: An instance  $\mathcal{I} = (N, M, \succeq)$  over  $n$  agents and  $m$ 
   items, and a round limit  $t$  whose default value is  $+\infty$ .
2 Output: An allocation  $\mathbf{A}$  and the remaining items in  $M$ .
3 Initialize  $\mathbf{A} = (A_1, \dots, A_n)$  where  $A_j \leftarrow \emptyset$  for each  $j \in N$ ,
   and  $c \leftarrow 1$ .
4 while  $M \neq \emptyset$  and  $c \leq t$  do
5   Initialize  $T \leftarrow N$ .
6   // Allocate items to the agents in  $T$ 
7   while  $T \neq \emptyset$  do
8      $i \leftarrow$  an arbitrary agent in  $T$ .
9      $\mathbf{A}', g, M' \leftarrow$  Algorithm 2( $i, M, \mathbf{A}$ ).
10     $\mathbf{A} \leftarrow \mathbf{A}', M \leftarrow M', T \leftarrow T \setminus \{i\}$ .
11    if  $M = \emptyset$  then
12      | Return  $\mathbf{A}$  and  $M$ .
13    end
14  end
15  Denote by  $g_j$  the item that agent  $j$  gets in this round for
   every  $j \in N$ .
16  Initialize  $S_1 \leftarrow \emptyset, S_2 \leftarrow \emptyset$ .
17  // Select the first set of agents who will be removed
18  while there exists an agent  $k \in N \setminus S_1$  such that  $g_k \succ_j g_j$ 
   for some  $j \in N \setminus \{k\}$  do
19    |  $S_1 \leftarrow S_1 \cup \{k\}$ .
20  end
21  // Select the second set of agents who will be removed
22  for each  $k \in S_1$  do
23    |  $S_2 \leftarrow S_2 \cup \{k' \in N \setminus S_1 :$ 
   |    $g_{k'} \text{ can be reached by } g_k \text{ in } G(\mathbf{A})\}$ .
24  end
25   $N \leftarrow N \setminus S_1 \setminus S_2$ .
26   $c \leftarrow c + 1$ .
27 end
28 Return  $\mathbf{A}$  and  $M$ .

```

when i picks an item, thus i can not have got an item that can be reached by g_j^l , a contradiction. \square

The proof of Observation 1 also gives the second observation.

OBSERVATION 2. *In any round, the items that the removed agents get are finalized when the last agent picks an item.*

By Observation 1, there is always at least one agent picking items until all the items are allocated, thus we have the following lemma.

LEMMA 4. *For any instance with ordinal preferences, Algorithm 1 can allocate all the items.*

The following lemma significantly simplifies the analysis.

LEMMA 5. *Without loss of generality, we can assume that there is no exchange during the execution of Algorithm 1.*

The intuition of Lemma 5 is that for every instance, we can somehow design a new algorithm that explicitly specifies who gets

Algorithm 2: Picking an item.

```

1 Input: An agent  $i$ , a set of items  $M$ , and a (partial) allocation
   $A = (A_1, \dots, A_n)$ .
2 Output: A new allocation  $A' = (A'_1, \dots, A'_n)$ , the item  $g$  that
  agent  $i$  picks, the remaining items in  $M$ .
3 Initialize  $A'_j \leftarrow A_j$  for every  $j \in N$ .
4  $g \leftarrow i$ 's favorite item among the unallocated or
  exchangeable items (tie breaks arbitrarily).
5 if  $g$  is unallocated then
6    $M \leftarrow M \setminus \{g\}$ .
7 else
8   Let  $P$  be an exchange path in  $G(A)$  that starts at  $g$  and
   ends at  $g'$ .
9   Compute the path-exchange allocation  $A^P$ .
10   $A' \leftarrow A^P$ ,  $M \leftarrow M \setminus \{g'\}$ .
11 end
12  $A'_i \leftarrow A'_i \cup \{g\}$ .
13 return  $A'$ ,  $g$  and  $M$ .

```

what at Step 9 so that there is no exchange and it outputs exactly the same allocation with Algorithm 1. We defer the complete proof of Lemma 5 to the full paper. Thus in the following, it suffices to consider that there is no exchange during Algorithm 1; that is, the item that an agent picks in a round is exactly the item she gets in that round.

LEMMA 6. *For any instance with ordinal preferences, Algorithm 1 computes a DL-EFX allocation.*

PROOF. Consider any two agents i and j , whose bundles returned by Algorithm 1 are A_i and A_j , respectively. We suppose that agent i is not removed from the algorithm before round k and agent j is not removed before round k' . We denote $A_i = \{a_i^1, \dots, a_i^k\}$ and $A_j = \{a_j^1, \dots, a_j^{k'}\}$, where a_i^l for any $l \in [k]$ and $a_j^{l'}$ for any $l' \in [k']$ are the items that i and j get in rounds l and l' , respectively. Since agent i always picks her favorite items, we have

$$a_i^1 \succeq_i \dots \succeq_i a_i^k. \quad (1)$$

Moreover, since neither agent i nor j is removed before round $\min\{k, k'\}$, we have

$$a_i^l \succeq_i a_j^l \text{ for every } l \in [\min\{k, k'\} - 1]. \quad (2)$$

We consider the following three cases:

Case 1: $k \geq k' = 1$. In this case, clearly, agent i does not DL-envy agent j after removing the only item in A_j .

Case 2: $k \geq k' \geq 2$. In this case, since agent i picks $a_i^{k'-1}$ before agent j picks $a_j^{k'}$, we have

$$a_i^{k'-1} \succeq_i a_j^{k'}. \quad (3)$$

Combining Equations (1), (2), (3), it follows that $A_i \succeq_i^{\text{DL}} A_j \setminus \{g\}$ for any $g \in A_j$.

Case 3: $k < k'$. In this case, agent i is removed at the end of round k but agent j is not. First, we have

$$a_i^k \succ_i a_j^k, \quad (4)$$

since otherwise, $a_j^k \succeq_i a_i^k$ implies that j should have been removed in the second set at the end of round k , a contradiction. Second, according to Observation 2, a_i^k is finalized at the end of round k , which according to Lemma 2, gives us

$$a_i^k \succ_i a_j^l \text{ for every } l \in [k+1, k']. \quad (5)$$

Combining Equations (1), (2), (4), (5), it follows that $A_i \succ_i^{\text{DL}} A_j \setminus \{g\}$ for any $g \in A_j$. \square

LEMMA 7. *For any instance with ordinal preferences, Algorithm 1 computes a DL-MMS allocation.*

PROOF. Consider any agent i , whose bundle returned by Algorithm 1 is $A = (a_1, \dots, a_{k_i})_{\succeq_i}$. Denote her DL-MMS by $B = (b_1, \dots, b_{k_i})_{\succeq_i}$, from Lemma 1, $b_j = \lfloor \frac{|E_i^j|}{n_j} \rfloor$ for every $j \in [k_i]$. Our target is to show $A \succeq_i^{\text{DL}} B$.

We consider the process of allocating items in each equivalence class of agent i , from the one that i prefers the most (i.e., E_i^1) to the one that she prefers the least (i.e., $E_i^{k_i}$). Each time, we focus on only one equivalence class and after all items in that equivalence class have been allocated, we turn to the next equivalence class. Note that when items in one equivalence class are being allocated, items in later equivalence classes may also be allocated. Formally, for every $j \in [k_i]$, let r_j be the first round by the end of which all items in the first j equivalence classes (i.e., $\bigcup_{l \leq j} E_i^l$) have been allocated. For every $j \in [k_i - 1]$, let x_j be the number of items in later equivalence classes (i.e., $\bigcup_{l \geq j+1} E_i^l$) that are allocated before round r_j . Then, we have the following claim.

CLAIM 1. *For every $j \in [k_i]$, one of the following two holds:*

- (1) $a_i > \lfloor \frac{|E_i^j|}{n_i} \rfloor$ for some $l \in [j]$.
- (2) $a_j = \lfloor \frac{|E_i^j|}{n_j} \rfloor$. Besides, if $j < k_i$, at most $n_{j+1} - x_j$ agents (including agent i) can pick the remaining items in $\bigcup_{l \geq j+1} E_i^l$.

The reasonings for the base case and the induction step are similar. Thus here we prove Claim 1 for the base case and defer the proof of the induction step to the full paper.

PROOF OF CLAIM 1 FOR BASE CASE. We first show by contradiction that agent i is not removed before round r_1 . Suppose that i is removed at the end of round $r' < r_1$, according to Observation 2, the item she picks in round r' (which is an item in E_i^1) is finalized when the last agent picks an item in round r' . Then Lemma 2 gives that all items in E_i^1 are also finalized, which means that $r_1 = r'$, a contradiction. Since agents pick items in rounds, agent i can get at least

$$\lfloor \frac{|E_i^1| + x_1}{n_1} \rfloor \geq \lfloor \frac{|E_i^1|}{n_1} \rfloor$$

items in E_i^1 . If i gets more than $\lfloor \frac{|E_i^1|}{n_1} \rfloor$ items in E_i^1 , the first statement is true. Thus it remains to consider that i gets exactly $\lfloor \frac{|E_i^1|}{n_1} \rfloor$ items in E_i^1 , which first gives

$$\lfloor \frac{|E_i^1| + x_1}{n_1} \rfloor = \lfloor \frac{|E_i^1|}{n_1} \rfloor. \quad (6)$$

We further consider the following cases.

Case 1: $((|E_i^1| + x_1) \bmod n_1) = 0$. This case happens only when the following conditions are satisfied simultaneously,

- $(|E_i^1| \bmod n_1) = 0$;
- no agent is removed before round r_1 ;
- each of the n_1 agents picks $\frac{|E_i^1|}{n_1}$ items in E_i^1 in the first r_1 rounds.

In this case, clearly, $x_1 = 0$ and $a_1 = \frac{|E_i^1|}{n_1}$. Besides, agent i is not removed at the end of round r_1 , since otherwise, according to Observation 2, the last agent in round r_1 can not have picked an item in E_i^1 , a contradiction. Therefore, at most $n_1 = n_2 = n_2 - x_1$ agents (including agent i) can pick the remaining items in $\bigcup_{l \geq 2} E_i^l$ in later rounds, thus the second statement is true.

Case 2: $((|E_i^1| + x_1) \bmod n_1) \neq 0$. In this case, agent i does not get an item in E_i^1 in round r_1 . Observe that at least

$$(|E_i^1| + x_1) \bmod n_1 = (|E_i^1| \bmod n_1) + x_1$$

agents get an item in E_i^1 in round r_1 , where the equality is because Equation 6. These agents are removed at the end of round r_1 since agent i strictly prefers their items to her own. Therefore, at most $n_1 - (|E_i^1| \bmod n_1) - x_1 = n_2 - x_1$ agents (including agent i) can pick the remaining items in $\bigcup_{l \geq 2} E_i^l$ in round r_1 and later rounds, thus the second statement is true. \square

From Claim 1, it directly follows that $A \succeq_i^{\text{DL}} B$, thus completing the proof of the lemma. \square

LEMMA 8. *For any instance with ordinal preferences, Algorithm 1 computes a DL-PO allocation.*

PROOF. Following Theorem 5 of [6], it can be easily seen that a (partial) allocation is not DL-PO if and only if there exists a cycle in the exchange graph which contains at least one edge corresponding to a strict preference. Therefore, it suffices prove that the exchange graph at the end of our algorithm does not contain such a cycle.

Denote by $\mathbf{A} = (A_1, \dots, A_n)$ the allocation returned by Algorithm 1, and suppose for the sake of contradiction that there is a cycle $v_1 \rightarrow \dots \rightarrow v_k \rightarrow v_1$ in $G(\mathbf{A})$ where $v_1 \in A_i$ and $v_2 \succ_i v_1$. Then, the path $v_2 \rightarrow \dots \rightarrow v_1$ implies that v_2 is exchangeable before agent i picks v_1 . Therefore, agent i should have picked v_2 instead of v_1 , a contradiction. \square

THEOREM 1. *Given any instance with arbitrary ordinal preferences, an allocation that is simultaneously DL-EFX, DL-MMS and DL-PO exists and can be computed in polynomial time.*

PROOF. From Lemmas 6, 7 and 8, the allocation returned by Algorithm 1 is simultaneously DL-EFX, DL-MMS and DL-PO. To complete the proof, it suffices to show that the algorithm runs in polynomial time. As has been shown in Lemma 4, at least one agent picks an item in each round, thus at least one unallocated item is removed from M and there are at most m rounds. In each round, at most n agents pick items. When an agent picks an item, it takes polynomial time to find the agent's favorite item and to search the exchange path in the exchange graph. At the end of each round, it takes polynomial time to find the first set of agents to be removed and to search the exchange graph to find the second set. Therefore, Algorithm 1 runs in polynomial time. \square

4 COMPUTING POSSIBLY FAIR AND BALANCED ALLOCATIONS

Next, we design an algorithm to compute a simultaneously p-EFX, p-MMS and balanced allocation. Moreover, under the same valuation profile to achieve EFX and MMS, the allocation is also PO.

4.1 The Selection of Cardinal Valuation

For any agent i whose ordinal preference is \succeq_i , we define her cardinal valuation u_i as follows. For any $g \in M$, let

$$u_i(g) = W_i + v_i(g), \quad (7)$$

where v_i is *lexicographic* and satisfies

$$v_i(o) > \sum_{o': o \succ_i o'} v_i(o')$$

and W_i is a constant that satisfies $W_i > \sum_{o \in M} v_i(o)$. The insights behind this cardinal valuation are that v_i can guarantee the allocation being both EFX and MMS, and W_i can guarantee the allocation being balanced.

For any valuation as defined in (7), the maximin share can be computed in polynomial time. In particular, we have the following lemma, whose proof is in the full paper.

LEMMA 9. *For any agent whose cardinal valuation is as defined in (7), her MMS bundle contains the first $\lfloor \frac{m}{n} \rfloor$ items in the DL-MMS bundle that she prefers the most.*

4.2 The Algorithm

Now we are ready to present our second main algorithm (see Algorithm 3). In the algorithm, the agents are divided into two groups. The first group N_1 contains the first $n - (m \bmod n)$ agents of N , and the second group N_2 contains the others. We first let the agents in N_1 pick items by running Algorithm 1 for $\lfloor \frac{m}{n} \rfloor$ rounds. Then, for the agents in N_1 who get less than $\lfloor \frac{m}{n} \rfloor$ items in Algorithm 1, they pick items in the way described in Algorithm 2 until they have got in total $\lfloor \frac{m}{n} \rfloor$ items. Lastly, each of the agents in N_2 picks $\lfloor \frac{m}{n} \rfloor + 1$ items in the way described in Algorithm 2.

LEMMA 10. *Given that the cardinal valuation of each agent is as defined in (7), the allocation returned by Algorithm 3 is EFX.*

PROOF. Consider any two agents i and j , whose bundles returned by Algorithm 3 are A_i and A_j , respectively. If $|A_i| \geq |A_j|$, for any $g \in A_j$, we have

$$u_i(A_i) \geq |A_j| \cdot W_i > u_i(A_j \setminus \{g\}).$$

If $|A_i| < |A_j|$, it follows that $i \in N_1$, $j \in N_2$, and $|A_i| = |A_j| - 1$. Since i picks her items before j and she always picks her favorite items, $u_i(g) \geq u_i(g')$ holds for any $g \in A_i$ and $g' \in A_j$, which gives $u_i(A_i) \geq u_i(A_j \setminus \{g'\})$ for any $g' \in A_j$. \square

LEMMA 11. *Given that the cardinal valuation of each agent is as defined in (7), the allocation returned by Algorithm 3 is MMS fair.*

PROOF. Consider any agent i , denote by A_i her bundle returned by Algorithm 3 and by B her MMS bundle. From Lemma 9, B contains the first $\lfloor \frac{m}{n} \rfloor$ items in DLMMS_i^n that i prefers the most.

Algorithm 3: Computing a p-MMS, p-EFX and balanced allocation.

```

1 Input: An instance  $\mathcal{I} = (N, M, \succ)$  over  $n$  agents and  $m$ 
  items.
2 Output: An allocation  $\mathbf{A}$ .
3 Initialize  $\mathbf{A} = (A_1, \dots, A_n)$  where  $A_j \leftarrow \emptyset$  for each  $j \in N$ .
4  $N_1 \leftarrow \lfloor n - (m \bmod n) \rfloor$ .
5  $N_2 \leftarrow N \setminus N_1$ .
6  $\{A'_1, \dots, A'_{|N_1|}\}, M' \leftarrow \text{Algorithm 1}(N_1, M, \lfloor \frac{m}{n} \rfloor)$ .
7  $A_j \leftarrow A'_j$  for every  $j \in N_1$ ,  $M \leftarrow M'$ .
8 for  $i \in N_1$  do
9   while  $|A_i| < \lfloor \frac{m}{n} \rfloor$  do
10     $A', \_ , M' \leftarrow \text{Algorithm 2}(i, M, \mathbf{A})$ .
11     $\mathbf{A} \leftarrow \mathbf{A}', M \leftarrow M'$ .
12  end
13 end
14 for  $j \in N_2$  do
15   while  $|A_j| < \lfloor \frac{m}{n} \rfloor + 1$  do
16     $A', \_ , M' \leftarrow \text{Algorithm 2}(j, M, \mathbf{A})$ .
17     $\mathbf{A} \leftarrow \mathbf{A}', M \leftarrow M'$ .
18  end
19 end
20 Return  $\mathbf{A}$ .
```

If $i \in N_2$, A_i contains $(\lfloor \frac{m}{n} \rfloor + 1)$ items and we have

$$u_i(A_i) \geq (\lfloor \frac{m}{n} \rfloor + 1) \cdot W_i > u_i(B).$$

If $i \in N_1$, denote by A'_i agent i 's bundle after running Algorithm 1 for $\lfloor \frac{m}{n} \rfloor$ rounds. From Lemma 7, i (weakly) DL-prefers A'_i to the first $\lfloor \frac{m}{n} \rfloor$ items in $\text{DLMMS}_i^{|N_1|}$ that she prefers the most. Moreover, since $|N_1| \leq n$, $\text{DLMMS}_i^{|N_1|} \succeq_i^{\text{DL}} \text{DLMMS}_i^n$. Thus, $A'_i \succeq_i^{\text{DL}} B$, which gives $u_i(A_i) \geq u_i(A'_i) \geq u_i(B)$. \square

LEMMA 12. *Given that the cardinal valuation of each agent is as defined in (7), the allocation returned by Algorithm 3 is PO.*

PROOF. Denote by $\mathbf{A} = (A_1, \dots, A_n)$ the allocation returned by Algorithm 3. Following the proof of Lemma 8, it is easy to see that \mathbf{A} is also DL-PO since agents in Algorithm 3 also always pick items in the way described in Algorithm 2.

Suppose for the sake of contradiction that there exists an allocation $\mathbf{B} = (B_1, \dots, B_n)$ which Pareto dominates \mathbf{A} , given that each agent's valuation function is as defined in (7). Assume without loss of generality that $u_1(B_1) > u_1(A_1)$ and $u_j(B_j) \geq u_j(A_j)$ for any $j \in N$. First observe that $|A_j| = |B_j|$ for every $j \in N$, since otherwise, there exists an agent $k \in N$ such that $|B_k| < |A_k|$, leading to the following contradiction,

$$u_k(B_k) < |A_k| \cdot W_k \leq u_k(A_k).$$

This observation, combining with our assumption, gives us $B_1 \succ_1^{\text{DL}} A_1$ and $B_j \succeq_j^{\text{DL}} A_j$ for any $j \in N$, which contradicts the fact that \mathbf{A} is DL-PO. Therefore, such an allocation \mathbf{B} does not exist. \square

THEOREM 2. *For any instance with ordinal preferences, Algorithm 3 computes an allocation that is simultaneously p-EFX, p-MMS and balanced in polynomial time.*

PROOF. Clearly, Algorithm 3 can allocate all the items and returns a balanced allocation. From Lemmas 10 and 11, the allocation is also p-EFX and p-MMS. Since it has been shown in Theorem 1 that both Algorithm 1 and Algorithm 2 run in polynomial time, Algorithm 3 also runs in polynomial time. \square

5 NECESSARY FAIRNESS

In this section, we focus on necessary MMS (n-MMS) and necessary EFX (n-EFX). It is not hard to see that n-MMS and n-EFX are hard to satisfy since they have stronger requirements than the cardinal definitions. Consider the following example.

EXAMPLE 2. *Let $M = \{o_1, o_2, o_3, o_4, o_5\}$ and $N = \{1, 2\}$ such that for both $i \in N$:*

$$o_1 \succ_i o_2 \sim_i o_3 \sim_i o_4 \sim_i o_5.$$

If both agents have lexicographic valuations on items, then the only MMS or EFX allocation is that one agent gets o_1 and the other agent gets o_2, o_3, o_4, o_5 . However, if o_1 is only slightly more valuable than the other items, the only MMS or EFX allocation is that one agent gets o_1 together with one of o_2, o_3, o_4, o_5 , and the other agent gets the remaining. Thus no allocation can simultaneously satisfy both cases.

Before the end of this section, we show a simple necessary condition for n-MMS and n-EFX allocations. For each agent i and item o , let B_i^o be the set of the items that are at least as good as o .

PROPOSITION 3. *If an assignment A is n-MMS or n-EFX, then for each $i \in N$, $|A_i \cap B_i^o| \geq \lfloor \frac{|B_i^o|}{n} \rfloor$ for all $o \in M$.*

PROOF. Assume that there exists some o such that $|A_i \cap B_i^o| < \lfloor \frac{|B_i^o|}{n} \rfloor$. Consider $u_i \in \mathcal{U}(\succ_i)$ such that $W \leq u_i(o') \leq W + \epsilon$ for $o' \in B_i^o$ and $u_i(o') \leq \epsilon$ for $o' \in M \setminus B_i^o$ where ϵ is an arbitrarily small value. This means $u_i(A_i) < W \lfloor \frac{|B_i^o|}{n} \rfloor$. However agent i can guarantee at least $W \lfloor \frac{|B_i^o|}{n} \rfloor$ utility by partitioning the elements of B_i^o in a balanced way. Therefore the allocation cannot be n-MMS. The same situation applies to n-EFX; Since $|A_i \cap B_i^o| < \lfloor \frac{|B_i^o|}{n} \rfloor$, there must be an agent who takes more than $\lfloor \frac{|B_i^o|}{n} \rfloor$ items in B_i^o . Using the same valuation function, EFX also fails. \square

Example 2 also shows that $|A_i \cap B_i^o| \geq \lfloor \frac{|B_i^o|}{n} \rfloor$ for all $o \in M$ for all $i \in N$ is not enough for an allocation to be n-MMS or n-EFX.

6 CONCLUSION

In this work, we study possible fairness for allocating indivisible resources. We first prove that an EFX, MMS and PO allocation exists if the preferences are (weak) lexicographic. Based on this result, we further prove that a simultaneously possible EFX, MMS, and balanced allocation exists and can be computed in polynomial time. Moreover, we use the same valuation profile to achieve EFX and MMS, under which the allocation is PO. We believe it is of both theoretical interest and practical importance to characterize what possible fairness can be achieved in resource allocation problems. Several interesting future directions include extending our work to the allocation of chores or a mixture of goods and chores, and the more general setting when the agents have asymmetric weights.

ACKNOWLEDGMENTS

Aziz is supported by the Defence Science and Technology Group through the Centre for Advanced Defence Research in Robotics and Autonomous Systems under the project “Task Allocation for Multi-Vehicle Coordination” (UA227119). Li is supported by NSFC under Grant No. 62102333, HKSAR RGC under Grant No. PolyU 25211321 and 15221420, CCF-HuaweiLK2022006, and HK PolyU under Grant No. P0034420.

REFERENCES

- [1] David J. Abraham, Katarína Cechlárová, David F. Manlove, and Kurt Mehlhorn. 2005. Pareto Optimality in House Allocation Problems. In *ISAAC (Lecture Notes in Computer Science, Vol. 3827)*. Springer, 1163–1175.
- [2] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. 2022. Fair Division of Indivisible Goods: A Survey. *CoRR abs/2202.07551* (2022).
- [3] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2016. On Truthful Mechanisms for Maximin Share Allocations. In *IJCAI. IJCAI/AAAI Press*, 31–37.
- [4] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. 2017. Approximation Algorithms for Computing Maximin Share Allocations. *ACM Trans. Algorithms* 13, 4 (2017), 52:1–52:28.
- [5] Georgios Amanatidis, Evangelos Markakis, and Apostolos Ntokos. 2020. Multiple birds with one stone: Beating $1/2$ for EFX and GMMS via envy cycle elimination. *Theor. Comput. Sci.* 841 (2020), 94–109.
- [6] Haris Aziz, Péter Biró, Jérôme Lang, Julien Lesca, and Jérôme Monnot. 2019. Efficient reallocation under additive and responsive preferences. *Theor. Comput. Sci.* 790 (2019), 1–15.
- [7] Haris Aziz, Sylvain Bouveret, Ioannis Caragiannis, Ira Giagkousi, and Jérôme Lang. 2018. Knowledge, Fairness, and Social Constraints. In *AAAI. AAAI Press*, 4638–4645.
- [8] Haris Aziz, Serge Gaspers, Simon Mackenzie, and Toby Walsh. 2015. Fair assignment of indivisible objects under ordinal preferences. *Artif. Intell.* 227 (2015), 71–92.
- [9] Haris Aziz, Bo Li, Hervé Moulin, and Xiaowei Wu. 2022. Algorithmic Fair Allocation of Indivisible Items: A Survey and New Questions. *CoRR abs/2202.08713* (2022).
- [10] Haris Aziz and Jérôme Monnot. 2020. Computing and testing Pareto optimal committees. *Auton. Agents Multi Agent Syst.* 34, 1 (2020), 24.
- [11] Siddharth Barman and Sanath Kumar Krishnamurthy. 2020. Approximation Algorithms for Maximin Fair Division. *ACM Trans. Economics and Comput.* 8, 1 (2020), 5:1–5:28.
- [12] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang. 2010. Fair Division under Ordinal Preferences: Computing Envy-Free Allocations of Indivisible Goods. In *ECAI (Frontiers in Artificial Intelligence and Applications, Vol. 215)*. IOS Press, 387–392.
- [13] Steven J. Brams, Paul H. Edelman, and Peter C. Fishburn. 2001. Paradoxes of fair division. *The Journal of philosophy* 98, 6 (2001), 300–314.
- [14] Steven J. Brams, Paul H. Edelman, and Peter C. Fishburn. 2003. Fair division of indivisible items. *Theory and Decision* 55, 2 (2003), 147–180.
- [15] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- [16] Ioannis Caragiannis, Nick Gravin, and Xin Huang. 2019. Envy-Freeness Up to Any Item with High Nash Welfare: The Virtue of Donating Items. In *EC. ACM*, 527–545.
- [17] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Economics and Comput.* 7, 3 (2019), 12:1–12:32.
- [18] Hau Chan, Jing Chen, Bo Li, and Xiaowei Wu. 2019. Maximin-Aware Allocations of Indivisible Goods. In *IJCAI. ijcai.org*, 137–143.
- [19] Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta, and Pranabendu Misra. 2021. Improving EFX Guarantees through Rainbow Cycle Number. In *EC. ACM*, 310–311.
- [20] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. 2021. A Little Charity Guarantees Almost Envy-Freeness. *SIAM J. Comput.* 50, 4 (2021), 1336–1358.
- [21] Uriel Feige, Ariel Sapir, and Laila Tauber. 2021. A Tight Negative Example for MMS Fair Allocations. In *WINE (Lecture Notes in Computer Science, Vol. 13112)*. Springer, 355–372.
- [22] Jugal Garg and Eklavya Sharma. 2022. Existence and Computation of Epistemic EFX. *CoRR abs/2206.01710* (2022).
- [23] Jugal Garg and Setareh Taki. 2021. An improved approximation algorithm for maximin shares. *Artif. Intell.* 300 (2021), 103547.
- [24] Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2021. Fair Allocation of Indivisible Goods: Improvement. *Math. Oper. Res.* 46, 3 (2021), 1038–1053.
- [25] Daniel Halpern and Nisarg Shah. 2021. Fair and Efficient Resource Allocation with Partial Information. In *IJCAI. ijcai.org*, 224–230.
- [26] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. 2021. Fair and Efficient Allocations under Lexicographic Preferences. In *AAAI. AAAI Press*, 5472–5480.
- [27] Kathrin Konczak and Jérôme Lang. 2005. Voting procedures with incomplete preferences. In *Proc. IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling*, Vol. 20. Citeseer.
- [28] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2018. Fair Enough: Guaranteeing Approximate Maximin Shares. *J. ACM* 65, 2 (2018), 8:1–8:27.
- [29] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *EC. ACM*, 125–131.
- [30] Benjamin Plaut and Tim Roughgarden. 2020. Almost Envy-Freeness with General Valuations. *SIAM J. Discret. Math.* 34, 2 (2020), 1039–1068.
- [31] Kirk Pruhs and Gerhard J. Woeginger. 2012. Divorcing Made Easy. In *FUN (Lecture Notes in Computer Science, Vol. 7288)*. Springer, 305–314.
- [32] Leonard J. Schulman and Vijay V. Vazirani. 2012. Allocation of divisible goods under lexicographic preferences. *arXiv preprint arXiv:1206.4366* (2012).
- [33] Erel Segal-Halevi, Avinatan Hassidim, and Haris Aziz. 2020. Fair Allocation with Diminishing Differences. *J. Artif. Intell. Res.* 67 (2020), 471–507.