

Bin Packing and Covering: Pushing the Frontier on the Maximin Share Fairness

Bo Li*

Ankang Sun[†]

Zunyu Wang[‡]

Yu Zhou[§]

Abstract

We study a fundamental fair allocation problem, where the agent’s value is determined by the number of bins either used to pack or cover the items allocated to them. Fairness is evaluated using the maximin share (MMS) criterion. This problem is not only motivated by practical applications, but also serves as a natural framework for studying group fairness. As MMS is not always satisfiable, we consider two types of approximations: cardinal and ordinal. For cardinal approximation, we relax the requirements of being packed or covered for a bin, and for ordinal approximation, we relax the number of bins that are packed or covered. For all models of interest, we provide constant approximation algorithms.

1 Introduction

Fair allocation of indivisible items has recently attracted considerable interest in the fields of computational economics [Akrami *et al.*, 2025], artificial intelligence [Seddighin and Seddighin, 2024], and theoretical computer science [Garg *et al.*, 2025]. Most theoretical work focuses on broad valuation classes, including additive, submodular, and subadditive valuations [Akrami and Garg, 2024; Seddighin and Seddighin, 2025]. Motivated by real-world applications, there is a recent research trend on how to integrate the combinatorial structure of items into the design of fair algorithms. Various models have been investigated, such as the order delivery problem [Hosseini *et al.*, 2025], budget-feasible constraints [Barman *et al.*, 2023], and scheduling constraints [Eqbal *et al.*, 2024].

In this paper, we revisit the model introduced in [Li *et al.*, 2023], where items are allocated to agents who have bin packing costs [Martello and Toth, 1990; Coffman Jr *et al.*, 2013]. This model captures a broad range of real-world applications. For example, when cloud providers allocate computing resources (RAM, CPU, storage) to different users or tenants. Users want to minimize the number of virtual machines used to run their tasks. A similar situation arises in logistics, where parcels are distributed among delivery agents. Each agent seeks to minimize the number of trucks needed to complete their deliveries. Formally, the problem can be modeled as allocating m indivisible chores (interpreted as items) among n agents, where the size of each item is agent-specific. Upon receiving a bundle of items, each agent must pack them into bins such that the total size in any bin does not exceed a given capacity. The cost incurred by an agent is defined as the minimum number of bins required to pack all her allocated items feasibly.

Li *et al.* [2023] considered the maximin share (MMS) fairness introduced by Budish [2011]. The criterion is defined through a thought experiment: each agent divides all items into n bundles, but only receives the bundle with the largest cost, where n is the number of agents. The MMS of the agent is the smallest cost they can guarantee themselves in this experiment. For the bin packing model, the MMS is essentially the minimum κ such that all items can be partitioned into n bundles so that the items in every bundle can be packed into κ bins. An allocation is regarded as MMS fair if every agent’s cost is at

*Department of Computing, The Hong Kong Polytechnic University, China. comp-bo.li@polyu.edu.hk

[†]Department of Computing, The Hong Kong Polytechnic University, China ankang.sun@polyu.edu.hk

[‡]Department of Computing, The Hong Kong Polytechnic University, China. zunyu1.wang@connect.polyu.hk

[§]Institute of AI and Future Networks, Beijing Normal University, China. yu.zhou@bnu.edu.cn

Setting	Ordinal Approximation	Cardinal Approximation
Bin Covering	$\frac{3}{4}$ (Theorem 2)	$\frac{2}{3}$ (Theorem 1)
Bin Packing	$\frac{4}{3}$ (Theorem 3)	$\frac{4}{3}$ [Li <i>et al.</i> , 2025]

Table 1: Main results: For the results on ordinal approximations, we admit a small constant additive loss.

most her MMS. However, Li *et al.* [2023] demonstrate that MMS fair allocations may not always exist. As a result, they study an ordinal relaxation of MMS fairness. An allocation is said to be α -ordinal-approximate MMS fair if the number of bins each agent uses to pack the assigned items does not exceed α times her MMS. They show that a simple round-robin algorithm achieves a 2-ordinal approximation, and this ratio is tight. Nevertheless, they observe that allowing a small constant additive loss enables an improved multiplicative approximation ratio of $\frac{3}{2}$. In this work, we further improve the asymptotic ordinal approximation ratio. Furthermore, we propose to study the *cardinal* approximation. Informally, an allocation is α -cardinal-approximate MMS fair if each agent can pack the assigned items using MMS bins, with each bin exceeding its capacity by at most a factor of α .

While the bin packing model has been studied, the dual problem – bin covering model where items are goods [Assmann *et al.*, 1984] – has not been explored. That is, upon receiving a set of items, the agent uses them to cover bins, where a bin is considered covered if the total size of items in it meets or exceeds the bin’s capacity. Here, the agent’s utility is the maximum number of bins they are able to cover. This setting also widely appears in the real world. For example, consider the medical kit assembly problem, where hospitals are allocated medical supplies to assemble emergency or surgical kits. A kit is considered usable only when it includes all essential items meeting or exceeding a required threshold (e.g., a complete set of PPE). Similarly, in IT hardware deployment problem, an IT department receives various hardware components (CPUs, RAM, drives) and wants to assemble as many fully functional computers (bins) as possible, each meeting a minimum specification. For the bin covering model, the MMS of each agent is the maximum κ such that all items can be partitioned into n bundles, with the items in each bundle capable of covering κ bins. Accordingly, an MMS fair allocation requires every agent’s utility to be no smaller than MMS. In this work, for the first time, we study the ordinal and cardinal approximations of MMS fairness for the bin covering model.

The bin packing and covering models not only reflect practical applications with such combinatorial structures, but also provide a natural model to study group fairness by treating agents as groups and bins as individual group members. The cardinal approximation quantifies how much each member may be overburdened or undercompensated, while the ordinal approximation seeks to minimize the number of additional members required – ensuring no member is over-utilized – or to maximize the number of fully compensated members.

1.1 Main Contribution

In this paper, we study the extent to which MMS fairness can be satisfied in the bin covering and packing models. The MMS criterion is relaxed in two dimensions, either cardinal or ordinal. Our algorithmic results are summarized in Table 1.

Bin Covering Model We start with the bin covering model, where the items are goods and bring non-negative utilities to the agents. We first observe that exact MMS cannot be satisfied even when there are two agents. That is, there are instances with two agents for which no matter how items are allocated, there is at least one agent whose items can cover strictly fewer bins than their MMS. This result shows a sharp contrast to the traditional additive setting, for which an MMS allocation always exists for two agents. We then focus on the approximations.

We first design an algorithm that guarantees $\frac{2}{3}$ -cardinal-approximate MMS fair: it always computes an allocation where each agent can use their assigned items to cover their MMS number of bins, with the total size in each bin reaching at least $\frac{2}{3}$ of its capacity. We propose an algorithm that, at each

step, selects an agent to partition the remaining items into subsets, with the number of subsets equal to the number of remaining agents. Then, the algorithm matches the subsets to the agents in a way that each unmatched agent is not satisfied with any of the matched subsets. This idea has been considered in the literature [Aigner-Horev and Segal-Halevi, 2022], but our algorithm differs in the following two aspects: (1) instead of arbitrarily selecting an agent, we choose a specific remaining agent according to a certain rule, and (2) our algorithm imposes constraints on the subsets partitioned by the selected agent. These adaptations are necessary as simply applying the existing algorithm does not guarantee $\frac{2}{3}$ -cardinal-approximate MMS fair solution.

Next, we consider the ordinal approximation, but find that it is impossible to guarantee any non-zero approximation in general. This limitation occurs in cases where each agent’s MMS is 1, yet in every possible allocation, there exists at least one agent whose items cannot cover even a single bin. However, as the MMS value increases, the approximation improves significantly. This observation leads us to consider allowing a small constant additive loss in the approximation. Based on this, we design an algorithm that guarantees each agent receives utility of at least $\frac{3}{4}\kappa_i - \frac{7}{4}$, where κ_i is agent a_i ’s MMS.

Bin Packing Model For the bin packing model, the items can be viewed as chores and bring cost to the agents. For the cardinal approximation, we observe an equivalent relationship with the *homogeneous* job scheduling model studied in [Li *et al.*, 2025].¹ Thus, the results in [Li *et al.*, 2025] implies a $\frac{4}{3}$ -cardinal-approximation. For the ordinal approximation, we prove that there always exists an allocation in which each agent a_i has a cost no greater than $\frac{4}{3}\kappa_i + \frac{4}{3}$, where κ_i is the MMS of a_i . For any $\kappa_i \geq 2$, our result improves upon that in Li *et al.* [2023] which showed the existence of an allocation where each a_i has cost at most $\frac{3}{2}\kappa_i + 1$. This improvement relies on a different classification of the items and a more refined analysis on the structure of the MMS partition of agents.

1.2 Other Related Work

The problem of fair division has been well studied since 1948 when Steinhaus [1948] first formalized the fairness notion of proportionality (PROP), where each agent receives at least $\frac{1}{n}$ of his total value for all items. Another popular fairness notion is envy-freeness (EF) [Foley, 1967], where no agent wants to exchange her bundle with any other agent’s. For indivisible items, PROP and EF are hard to satisfy, which gives rise to their relaxations, such as envy-freeness up to one item (EF1) [Lipton *et al.*, 2004; Budish, 2011], envy-freeness up to any item (EFX) [Gourvès *et al.*, 2014; Caragiannis *et al.*, 2019], and maximin share (MMS) [Budish, 2011]. EF1 allocations always exist for both goods and chores with monotone valuations [Lipton *et al.*, 2004; Bhaskar *et al.*, 2021]. MMS allocations may not exist even for additive valuations [Kurokawa *et al.*, 2016, 2018] and the existence of EFX allocations is still unknown. Accordingly, the approximations of MMS and EFX are extensively studied by follow-up works, e.g., [Akrami and Garg, 2024; Ghodsi *et al.*, 2018; Barman and Krishnamurthy, 2020] for goods and [Zhou and Wu, 2024; Huang and Segal-Halevi, 2023; Huang and Lu, 2021; Garg *et al.*, 2024] for chores. Besides additive valuations, the approximation of MMS with submodular, XoS and subadditive valuations has been studied in [Barman and Krishnamurthy, 2020; Ghodsi *et al.*, 2018; Uziah and Feige, 2023; Akrami *et al.*, 2023; Seddighin and Seddighin, 2024] for goods and in [Li *et al.*, 2023] for chores.

¹Given any bin packing instance, we can transfer it into a homogeneous job scheduling instance by setting the number of machines owned by each group to the MMS value of the corresponding agent while keeping the item sizes unchanged. One can see that the MMS value (i.e., the makespan) of each group in the job scheduling instance is equal to the capacity of the bins owned by the agent. Accordingly, an α -GMMS allocation in the job scheduling instance leads to an α -CMMS allocation in the bin packing instance. For the other direction, we can transfer any homogeneous job scheduling instance to a bin packing one by setting the MMS value of each group to the capacity of the bins owned by the corresponding agent.

2 Preliminaries

For any integer $k \in \mathbb{Z}^+$, denote $[k] = \{1, \dots, k\}$. We consider the problem of allocating m indivisible items $M = \{e_1, \dots, e_m\}$ among n agents $N = \{a_1, \dots, a_n\}$. The agents have subjective opinions on the sizes of items, and the size of item e_j to agent a_i is denoted as $s_{i,j}$. Without loss of generality, assume that the bin capacity of every agent is 1 and $0 < s_{i,j} \leq 1$ for all i, j . For simplicity, let $s_i(S) = \sum_{e_j \in S} s_{i,j}$ and $s_i(e_j) = s_{i,j}$. Let $\mathbf{s} = (s_1, \dots, s_n)$. A fair allocation instance is denoted by $\mathcal{I} = (N, M, \mathbf{s})$. An allocation, denoted by $\mathbf{A} = (A_1, \dots, A_n)$, is an n -partition of M where A_i is the items allocated to agent a_i such that $A_1 \cup \dots \cup A_n = M$ and $A_i \cap A_j = \emptyset$ for all $i \neq j$. When items are goods, agents have bin covering valuations $v_i : 2^M \rightarrow \mathbb{Z}_{\geq 0}$, where $v_i(S)$ is the maximum number of bins that can be covered by items in $S \subseteq M$, i.e., the maximum $\tau \geq 0$ such that there is a τ -partition $(S_1, \dots, S_\tau)$ such that $s_i(S_j) \geq 1$ for all j . When items are chores, agents have bin packing cost functions $c_i : 2^M \rightarrow \mathbb{Z}_{\geq 0}$, where $c_i(S)$ is the minimum number of bins that can pack all items in S , i.e., the minimum $\tau \geq 0$ such that there is a τ -partition $(S_1, \dots, S_\tau)$ such that $s_i(S_j) \leq 1$ for all j . Denote by $\mathbf{v} = (v_1, \dots, v_n)$ and $\mathbf{c} = (c_1, \dots, c_n)$.

We consider the maximin share (MMS) fairness. For any positive integer k and set of items S , denote by $\Pi(S, k)$ the set of all k -partitions of S . For the bin covering model, agent a_i 's MMS is defined as

$$\kappa_i = \max_{(X_1, \dots, X_n) \in \Pi(M, n)} \min_{j \in [n]} v_i(X_j).$$

That is, κ_i is the maximum κ such that a_i can partition all items into n bundles and every bundle can cover no smaller than κ bins. A partition that reaches κ_i is called an MMS partition of a_i . An allocation $(A_1, \dots, A_n)$ is MMS fair if $v_i(A_i) \geq \kappa_i$ for all agents. It is worth noting that an MMS fair allocation may not exist even for instances with only two agents, which shows a sharp contrast to the traditional additive setting. Feige *et al.* [2021] designed a hard instance for four agents with additive valuations, where the first two agents have the same valuation and so do the other two agents. Every agent can partition all the items into four bundles and each bundle is valued exactly 1 (after normalization). However, no matter how the items are allocated, there is at least one agent whose value is strictly smaller than 1. This example can be adapted to the bin covering setting with two agents, where the values correspond to item sizes and each bin has a capacity of 1. In this case, the MMS values of both agents are 2, however, there will always be an agent who cannot cover two bins.

Therefore, we consider two ways to relax the requirements of MMS fairness.

Definition 1. Given a bin covering instance and $0 \leq \alpha \leq 1$, an allocation $(A_1, \dots, A_n)$ is α -cardinal-approximate MMS fair (α -CMMS) if every agent a_i can partition A_i into κ_i bundles $(A_{i,1}, \dots, A_{i,\kappa_i})$ such that $s_i(A_{i,j}) \geq \alpha$ for all $j \in [\kappa_i]$.

Definition 2. Given a bin covering instance and $0 \leq \alpha \leq 1$, an allocation $(A_1, \dots, A_n)$ is α -ordinal-approximate MMS fair (α -OMMS) if $v_i(A_i) \geq \alpha \kappa_i - c$ for some constant $c \geq 0$.

Note that the constant additive loss $-c$ is important in the above definition. If $c = 0$, we are not able to have $\alpha > 0$, even when $\kappa_i = 1$ and $n = 3$. We can adapt the hard instance with three agents designed by Feige *et al.* [2021] to the bin covering setting with three agents like above. Here, there will always be an agent whose items cannot pack a bin, resulting in α being 0 in Definition 2. Thus, for ordinal approximation, we focus solely on the asymptotic approximation ratio.

For the bin packing model, agent a_i 's MMS is defined as

$$\kappa_i = \min_{(X_1, \dots, X_n) \in \Pi(M, n)} \max_{j \in [n]} c_i(X_j).$$

That is, κ_i is the minimum κ such that a_i can partition all items into n bundles and the items in every bundle can be packed into κ bins. Similarly, a partition that reaches κ_i is called an MMS partition of a_i . An allocation $(A_1, \dots, A_n)$ is MMS fair if $c_i(A_i) \leq \kappa_i$ for all agents, and cardinal and ordinal approximations are considered to relax the requirements.

Definition 3. Given a bin packing instance and $\alpha \geq 1$, an allocation $(A_1, \dots, A_n)$ is α -cardinal-approximate MMS fair (α -CMMS) if every agent a_i can partition A_i into κ_i bundles $(A_{i,1}, \dots, A_{i,\kappa_i})$ such that $s_i(A_{i,j}) \leq \alpha$ for all $j \in [\kappa_i]$.

Definition 4. Given a bin packing instance and $\alpha \geq 1$, an allocation $(A_1, \dots, A_n)$ is α -ordinal-approximate MMS fair (α -OMMS) if $c_i(A_i) \leq \alpha\kappa_i + c$ for some constant $c \geq 0$.

Reading the ordinal approximation, if $c = 0$, it is proved in [Li *et al.*, 2023] that the tight approximation ratio is 2. But if we allow a small constant additive loss, the multiplicative ratio can be improved to $\frac{3}{2}$. In our paper, we continue studying the asymptotic bound and improve this ratio to $\frac{4}{3}$.

Finally, we call an instance identical ordering (IDO) if all agents agree on an ordering of the items regarding their sizes, i.e., $s_{i,1} \geq s_{i,2} \geq \dots \geq s_{i,m}$ for all a_i . It is widely known that IDO instances are the hardest to approximate MMS for the instances with additive valuations [Barman and Krishnamurthy, 2020; Huang and Lu, 2021]. This is also true for bin covering and bin packing models, and the proof is almost the same as that in these references. Thus we state the following and omit the proof.

Lemma 1. *If there is an algorithm that ensures α -CMMS (or α -OMMS) allocations for all IDO instances, there is another algorithm that ensures α -CMMS (or α -OMMS) allocations for arbitrary instances. Further, if the former algorithm runs in polynomial time, the latter does as well.*

In all the subsequent sections, we only consider IDO instances.

3 Cardinal Approximation for the Bin Covering Model

In this section, we present the algorithm to compute $\frac{2}{3}$ -CMMS allocations for the bin covering model. That is, given each agent a_i 's MMS value κ_i – the maximum number of bins each bundle can cover in the MMS partition – we compute an allocation in which the items received by each agent a_i can be partitioned into κ_i bundles, each with a total size of at least $\frac{2}{3}$.

Theorem 1. *For any bin covering instance, a $\frac{2}{3}$ -CMMS allocation exists.*

To simplify the analysis, we assume in the MMS partition of each agent a_i , every bin is exactly covered. In other words, M can be partitioned into $n\kappa_i$ bundles $X_1, \dots, X_{\kappa_i \cdot n}$ such that $s_i(X_j) = 1$ for all $j \in [\kappa_i \cdot n]$. This assumption is made without loss of generality: we can always rescale the item sizes so that for any $e_l \in M$, if $e_l \in X_j$, $s'_{i,l} = \frac{s_{i,l}}{s_i(X_j)}$. Then, if we can achieve $\frac{2}{3}$ -CMMS in the scaled instance, the same guarantee holds for the original instance.

3.1 The Algorithm

For each agent a_i , item e is large if $\frac{2}{3} \leq s_i(e) \leq 1$, and is medium if $\frac{1}{3} < s_i(e) < \frac{2}{3}$. If an item is neither large nor medium, it is small. As the capacity of the bin is 1, for each bundle in the MMS partition for a_i , aside from small items, it may contain at most (1) a large item, (2) a medium item, and (3) two medium items. We now adjust the MMS partition as follows: if there exists medium items e, e' such that $s_i(e) > s_i(e')$, and in the partition, e is packed in a bin with another medium item while e' is the only medium item in the bin, then swap the positions of e and e' . Repeat the process until no two medium items satisfying the swap condition remains, resulting in the *refined-partition*.

For each bundle in the refined-partition, again, aside from small items, it may contain at most (1) a large item, (2) a medium item, and (3) two medium items. Let F denote the set of large and medium items that, in the refined-partition, are placed solely in a bundle without any other medium items. Let H denote the set of medium items that, in the refined-partition, are placed in a bundle with another medium item. Thus, $F \cup H$ equals the set of large and medium items. For each a_i , as there are $\kappa_i \cdot n$ bundles in the refined-partition and items in H are placed in pairs, it holds that $|F| + \frac{|H|}{2} \leq \kappa_i \cdot n$. The sets F and

H can differ among agents. We abuse notation and use F and H for every agent, and will clarify the underlying agent from the context when needed.

Next, we introduce the high-level idea of the algorithm. At each step, we select an agent a_i from the remaining agents, denoted as $\mathcal{N}$. Agent a_i partitions the remaining items into $|\mathcal{N}|$ parts such that she is happy with every part. In our context, this translates to that every part S can be partitioned into κ_i bundles, each with a total item size of at least $\frac{2}{3}$ (when targeting $\frac{2}{3}$ -CMMS). Then create a bipartite graph with remaining agents on one side and the constructed $|\mathcal{N}|$ parts on the other side. There is an edge between an agent and a part if the agent is happy when receiving that part. *An envy-free matching with regards to the agents* is a matching in which every unmatched agent is not incident to any matched part. Aigner-Horev and Segal-Halevi [2022] proved that an envy-free matching with regards the agents always exists once the selected group can create $|\mathcal{N}|$ satisfied parts. Moreover, they provided a polynomial time algorithm to compute the maximin cardinality envy-free matching. Match agents to parts based on the envy-free matching, which guarantees that for every pair of unmatched a_j and a matched part S , agent a_j is not happy with receiving S .

Similar algorithmic ideas have been applied in computing approximate MMS allocations in other contexts [Hummel, 2025; Akrami and Rathi, 2025]. Our techniques differ from those existing algorithms in the following two aspects: at each step, (1) instead of choosing an arbitrary agent, we select the agent with the maximum total number of large and medium items, and (2) instead of allowing the selected agent to create $|\mathcal{N}|$ satisfied parts in an arbitrary way, we require that these $|\mathcal{N}|$ parts must partition large and medium items in a balanced manner. In specific, the selected agent creates the parts based on a (predetermined) $|\mathcal{N}|$ -arrangement, defined as follows: according to the non-descending order of item size, place remaining items into $|\mathcal{N}|$ columns such that j -th column contains j -th, $(2|\mathcal{N}| - j + 1)$ -th, $(2|\mathcal{N}| + j)$ -th, ... largest size item. When the selected agent creates $|\mathcal{N}|$ parts, each bundle must contain every large and medium item placed in some column in the $|\mathcal{N}|$ -arrangement. The proposed algorithm is formally introduced as Algorithm 1.

Algorithm 1 Computing cardinal approximate MMS allocations

Input: An instance (N, M, s) and κ_i for every $a_i \in N$.

Output: A $\frac{2}{3}$ -CMMS allocation $\mathbf{A} = (A_1, \dots, A_n)$.

- 1: Create the $|\mathcal{N}|$ -arrangement and initialize $\mathcal{N} \leftarrow N$.
 - 2: **while** $\mathcal{N} \neq \emptyset$ **do**
 - 3: Select $a_i \in \mathcal{N}$ with the maximum total number of large and medium items, breaking ties arbitrarily.
 - 4: Create $|\mathcal{N}|$ pairwise disjoint parts $S_1, \dots, S_{|\mathcal{N}|}$ as follows. For each $S_j, \forall j \leq |\mathcal{N}|$, first distribute all large and medium items in the j -th column of the $|\mathcal{N}|$ -arrangement into at most κ_i bundles such that the total size of items in each bundle is at most 1 for a_i . Then add small items to these κ_i bundles until the total size of items in each bundle is at least $\frac{2}{3}$ for a_i .
 - 5: Create a bipartite graph with vertices $\mathcal{N} \cup (S_1, \dots, S_{|\mathcal{N}|})$. Add edge (a_ℓ, S_j) if and only if S_j can be partitioned into κ_ℓ bundles, each with total item size at least $\frac{2}{3}$ for a_ℓ . Compute a maximum-cardinality envy-free matching and assign each matched part to its matched agent.
 - 6: Update $\mathcal{N}$ and M by removing matched agents and items. Then update the $|\mathcal{N}|$ -arrangement.
 - 7: **end while**
-

3.2 The Proof of Theorem 1

We now prove that Algorithm 1 can return an allocation. It suffices to show that in each round, the selected agent can create the $|\mathcal{N}|$ parts described in Line 4. First, consider the first round of the algorithm.

Lemma 2. *At the first round of the while-loop, for each agent a_i , she can create n pairwise disjoint parts $P_1, \dots, P_n$ following Line 4 such that for any $j \in [n]$, P_j can be partitioned into κ_i bundles, each with a total item size of at least $\frac{2}{3}$.*

Proof. For any $j \in [n]$, let $F_j \subseteq F$ (resp., $H_j \subseteq H$) denote the items in F (resp., H) in the j -th column of the n -arrangement. We focus on the j -th column and first show that a_i can place all items in $F_j \cup H_j$ into at most κ_i bundles (or bins), each having at most two items with total size at most 1 for a_i . In specific, we always place each item in F_j into a single bin, and thus, it suffices to show that items in H_j can be placed into at most $\kappa_i - |F_j|$ bundles, each having at most two items with total size at most 1.

Suppose $\lfloor \frac{|F|}{2n} \rfloor = x$ with $x \in \mathbb{N}^+ \cup \{0\}$ and $H = \{h_1, \dots, h_{|H|}\}$ with $s_i(h_1) \geq s_i(h_2) \geq \dots \geq s_i(h_{|H|})$. For simplicity, instead of $s_i(\{e\})$, we write $s_i(e)$. In the MMS refined-partition of a_i , items in H form $\frac{|H|}{2}$ pairs, satisfying that the two items in each pair have total size at most 1. Then it must hold that $s_i(h_1 \cup h_{|H|}) \leq 1$, as otherwise, the total size of h_1 and any item in H exceeds 1, contradicting the refined-partition. By the same reasoning, $s_i(h_2 \cup h_{|H|-1}) \leq 1$ holds, and indeed, for any r, r' with $r + r' \geq |H| + 1$, $s_i(h_r \cup h_{r'}) \leq 1$.

Observation 1. For any r, r' with $r + r' \geq |H| + 1$, $s_i(h_r \cup h_{r'}) \leq 1$.

Next, we below split the proof based on the values of $|F|$ and j .

Case 1: $|F| = 2nx + y$ with $y \in \{1, \dots, n\}$, and $j \in \{1, \dots, y\}$. In this case, we have $|F_j| = 2x + 1$. Then the remaining is to show that one can always place items in H_j into at most $\kappa_i - 2x - 1$ bundles, each having at most two items with total size at most 1. Note that $|F| = 2nx + y \leq \kappa_i n$ holds, and hence, $\kappa_i \geq 2x + 1$, as κ_i, x, y are integers and $y \geq 1$. That $|H_j| = 0$ is a trivial case, and then we focus on $|H_j| \geq 1$. By $|F| = 2nx + y$ and $|F| + \frac{|H|}{2} \leq \kappa_i n$, it holds that $|H| \leq 2\kappa_i n - 4nx - 2y$, which implies $|H_j| \leq 2\kappa_i - 4x - 2$.

First, consider the case when $|H_j| = 2\kappa_i - 4x - 2$. Since $|H_j| \geq 1$ and κ_i, x are integers, we have $\kappa_i \geq 2x + 2$. We enumerate items in H_j as follows:

$$H_j = \{h_r \mid r = 2np - y - j + 1 \text{ and } r = 2np - y + j \text{ for } p = 1, \dots, \kappa_i - 2x - 1\}.$$

Due to Observation 1 and $|H| \leq 2\kappa_i n - 4nx - 2y$, for each $s \in \{1, \dots, \frac{|H_j|}{2}\}$, the total size of the s -th largest size item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Therefore, one can place items in H_j into $\frac{|H_j|}{2} = \kappa_i - 2x - 1$ bundles, each having two items with total size at most 1.

When $|H_j| = 2\kappa_i - 4x - 3$, it must hold that $|H| \leq 2n(\kappa_i - 2x - 1) + j - y - 1$. In the above enumeration, item h_r with $r = 2n(\kappa_i - 2x - 1) - y + j$ no longer exists. If $|H_j| \leq 2$, we have $\kappa_i \geq 2x + 1 + |H_j|$, as κ_i is an integer. Accordingly, $|H_j| \leq \kappa_i - |F_j|$, and we are done. Focus on $|H_j| \geq 3$. By Observation 1, for each $s \in \{1, \dots, \lfloor \frac{|H_j|}{2} \rfloor\}$, the total size of the $(s + 1)$ -th largest size item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Then, excluding the largest value item in H_j , the remaining items can be placed into $\kappa_i - 2x - 2$ bundles, each having two items with total size at most 1. Therefore, a_i can place all items in H_j into $\kappa_i - 2x - 1$ bundles, each having at most two items with total size at most 1.

When $|H_j| \leq 2\kappa_i - 4x - 4$, it must hold that $|H| \leq 2n(\kappa_i - 2x - 2) + 2n - j - y$. As $\kappa_i \geq 2x + 2 + \frac{|H_j|}{2}$ and κ_i is an integer, we have $\kappa_i \geq 2x + 1 + |H_j| = |F_j| + |H_j|$ when $|H_j| \leq 3$. We can focus on $|H_j| \geq 4$. By Observation 1, for each $s \in \{1, \dots, \lfloor \frac{|H_j|}{2} \rfloor - 1\}$, the total size of the $(s + 2)$ -th largest value item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Thus, aside from the two largest size items in H_j , the remaining items can be placed into $\lceil \frac{|H_j| - 2}{2} \rceil \leq \kappa_i - 2x - 3$ bundles, each having at most two items with total size at most one. Therefore, agent a_i can place all items in H_j into at most $\kappa_i - 2x - 1$ bundles, each having at most two items with total size at most 1.

Case 2: $|F| = 2nx + y$ with $y \in \{1, \dots, n - 1\}$, and $j \in \{y + 1, \dots, n\}$. The case $y = j = n$ is covered in Case 1. In Case 2, we have $|F_j| = 2x$. Then the remaining is to show that a_i can always place items in H_j into at most $\kappa_i - 2x$ bundles, each having at most two items with total size at most 1. Note that $|F| = 2nx + y \leq \kappa_i n$ holds, and hence, $\kappa_i \geq 2x + 1$, as κ_i, x, y are integers and $y \geq 1$. That $|H_j| = 0$ is a trivial case, and then we focus on $|H_j| \geq 1$. By $|F| = 2nx + y$ and $|F| + \frac{|H|}{2} \leq \kappa_i n$, it holds that $|H| \leq 2\kappa_i n - 4nx - 2y$, which implies $|H_j| \leq 2\kappa_i - 4x$ as $j \geq y + 1$.

First, consider the case when $|H_j| = 2\kappa_i - 4x$. Enumerate items in H_j as follows:

$$H_j = \{h_r \mid r = 2n(p-1) + j - y \text{ and } r = 2np + 1 - j - y \text{ for } p = 1, \dots, \kappa_i - 2x\}.$$

Due to Observation 1 and $|H| \leq 2\kappa_i n - 4nx - 2y$, for each $s \in \{1, \dots, \lfloor \frac{|H_j|}{2} \rfloor\}$, the total size of the s -th largest size item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Therefore, one can place items in H_j into $\lfloor \frac{|H_j|}{2} \rfloor = \kappa_i - 2x$ bundles, each having two items with total size at most 1.

When $|H_j| = 2\kappa_i - 4x - 1$, it must hold that $|H| \leq 2n(\kappa_i - 2x) - j - y$. In the above enumeration, item h_r with $r = 2n(\kappa_i - 2x) + 1 - j - y$ no longer exists. If $|H_j| \leq 2$, we have $\kappa_i \geq 2x + |H_j|$, as κ_i is an integer. Accordingly, $|H_j| \leq \kappa_i - |F_j|$, and we are done. Focus on $|H_j| \geq 3$. By Observation 1, for each $s \in \{1, \dots, \lfloor \frac{|H_j|}{2} \rfloor\}$, the total size of the $(s+1)$ -th largest size item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Then, aside from the largest size item in H_j , the remaining items can be placed into $\kappa_i - 2x - 1$ bundles, each having two items with total size at most 1. Therefore, agent a_i can place all items in H_j into $\kappa_i - 2x$ bundles, each having at most two items with total size at most 1.

When $|H_j| \leq 2\kappa_i - 4x - 2$, it must hold that $|H| \leq 2n(\kappa_i - 2x - 1) + j - y - 1$. If $|H_j| \leq 3$, we have $\kappa_i \geq 2x + |H_j| = |F_j| + |H_j|$, as κ_i is an integer. Thus, we can focus on $|H_j| \geq 4$. Due to Observation 1, for each $s \in \{1, \dots, \lfloor \frac{|H_j|}{2} \rfloor - 1\}$, the total size of the $(s+2)$ -th largest size item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Thus, aside from the two largest size items in H_j , the remaining items can be placed into $\lceil \frac{|H_j| - 2}{2} \rceil \leq \kappa_i - 2x - 2$ bundles, each having at most two items with total size at most 1. Therefore, agent a_i can always place all items in H_j into at most $\kappa_i - 2x$ bundles, each having at most two items with total size at most 1.

Case 3: $|F| = 2nx + n + y$ with $y \in \{1, \dots, n\}$, and $j = \{n - y + 1, \dots, n\}$. In this case, we have $|F_j| = 2x + 2$. Then the remaining is to show that one can always place items in H_j into at most $\kappa_i - 2x - 2$ bundles, each having at most two items with total size at most 1. By $\kappa_i n \geq 2nx + n + y$ and $y \geq 1$, we have $\kappa_i \geq 2x + 2$ as κ_i is an integer. That $|H_j| = 0$ is a trivial case, and we focus on $|H_j| \geq 1$. By $|F| + \frac{|H|}{2} \leq \kappa_i n$, it holds that $|H| \leq 2\kappa_i n - 4nx - 2n - 2y$, which implies $|H_j| \leq 2\kappa_i - 4x - 4$.

First, consider the case when $|H_j| = 2\kappa_i - 4x - 4$. As $|H_j| \geq 1$, we have $2\kappa_i \geq 4x + 5$, implying $\kappa_i \geq 2x + 3$ as κ_i is an integer. Enumerate items in H_j as follows:

$$H_j = \{h_r \mid r = 2n(p-1) + n + j - y \text{ and } r = 2np + n - j - y + 1 \text{ for } p = 1, \dots, \kappa_i - 2x - 2\}.$$

Due to Observation 1 and $|H| \leq 2\kappa_i n - 4nx - 2n - 2y$, for each $s \in \{1, \dots, \lfloor \frac{|H_j|}{2} \rfloor\}$, the total size of the s -th largest size item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Therefore, one can place items in H_j into $\lfloor \frac{|H_j|}{2} \rfloor = \kappa_i - 2x - 2$ bundles, each having two items with total size at most 1.

When $|H_j| = 2\kappa_i - 4x - 5$, it must hold that $|H| \leq 2n(\kappa_i - 2x - 2) + n - j - y$. In the above enumeration, item h_r with $r = 2n(\kappa_i - 2x - 2) + n - j - y + 1$ no longer exists. If $|H_j| \leq 2$, we have $\kappa_i \geq 2x + 2 + |H_j| = |F_j| + |H_j|$, as κ_i is an integer. Thus, we focus on $|H_j| \geq 3$. By Observation 1, for each $s \in \{1, \dots, \lfloor \frac{|H_j|}{2} \rfloor\}$, the total size of the $(s+1)$ -th largest size item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Then, aside from the largest size item in H_j , the remaining items can be placed into $\kappa_i - 2x - 3$ bundles, each having two items with total size at most 1. Therefore, one can place all items in H_j into $\kappa_i - 2x - 2$ bundles, each having at most two items with total size at most 1.

When $|H_j| \leq 2\kappa_i - 4x - 6$, it must hold that $|H| \leq 2n(\kappa_i - 2x - 3) + n + j - y - 1$. If $|H_j| \leq 3$, we have $\kappa_i \geq 2x + 2 + |H_j| = |F_j| + |H_j|$, as κ_i is an integer. Thus, we focus on $|H_j| \geq 4$. Due to Observation 1, for each $s \in \{1, \dots, \lfloor \frac{|H_j|}{2} \rfloor - 1\}$, the total size of the $(s+2)$ -th largest size item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Thus, aside from the two largest size items in H_j , the remaining items can be placed into $\lceil \frac{|H_j| - 2}{2} \rceil \leq \kappa_i - 2x - 4$ bundles, each having at most two items with total size at most 1. Therefore, agent a_i can always place all items in H_j into at most $\kappa_i - 2x - 2$ bundles, each having at most two items with total size at most 1.

Case 4: $|F| = 2nx + n + y$ with $y \in \{1, \dots, n-1\}$, and $j = \{1, \dots, n-y\}$. In this case, we have $|F_j| = 2x + 1$. Then the remaining is to show that one can always place items in H_j into at most $\kappa_i - 2x - 1$ bundles, each having at most two items with total size at most 1. That $|H_j| = 0$ is a trivial case, and we focus on $|H_j| \geq 1$. By $|F| + \frac{|H|}{2} \leq \kappa_i n$, it holds that $|H| \leq 2\kappa_i n - 4nx - 2n - 2y$, which implies $|H_j| \leq 2\kappa_i - 4x - 2$.

First, consider the case when $|H_j| = 2\kappa_i - 4x - 2$. Enumerate items in H_j as follows:

$$H_j = \{h_r \mid r = 2np + n - j - y + 1 \text{ and } r = 2np + n + j - y \text{ for } p = 0, \dots, \kappa_i - 2x - 2\}.$$

Due to Observation 1 and $|H| \leq 2\kappa_i n - 4nx - 2n - 2y$, for each $s \in \{1, \dots, \frac{|H_j|}{2}\}$, the total size of the s -th largest size item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Therefore, one can place items in H_j into $\frac{|H_j|}{2} = \kappa_i - 2x - 1$ bundles, each having two items with total size at most 1.

When $|H_j| = 2\kappa_i - 4x - 3$, it must hold that $|H| \leq 2n(\kappa_i - 2x - 2) + n + j - y - 1$. If $|H_j| \leq 2$, we have $\kappa_i \geq 2x + 1 + |H_j| = |F_j| + |H_j|$, as κ_i is an integer. Thus, we focus on $|H_j| \geq 3$. By Observation 1, for each $s \in \{1, \dots, \lfloor \frac{|H_j|}{2} \rfloor\}$, the total size of the $(s+1)$ -th largest size item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Then, aside from the largest value item in H_j , the remaining items can be placed into $\kappa_i - 2x - 2$ bundles, each having two items with total size at most 1. Therefore, one can place all items in H_j into $\kappa_i - 2x - 1$ bundles, each having at most two items with total size at most 1.

When $|H_j| \leq 2\kappa_i - 4x - 4$, it must hold that $|H| \leq 2n(\kappa_i - 2x - 2) + n - j - y$. If $|H_j| \leq 3$, we have $\kappa_i \geq 2x + 1 + |H_j| = |F_j| + |H_j|$, as κ_i is an integer. Thus, we focus on $|H_j| \geq 4$. Due to Observation 1, for any $s \in \{1, \dots, \lfloor \frac{|H_j|}{2} \rfloor - 1\}$, the total size of the $(s+2)$ -th largest size item and the $(|H_j| + 1 - s)$ -th largest size item in H_j is at most 1. Thus, aside from the two largest size items in H_j , the remaining items can be placed into $\lceil \frac{|H_j| - 2}{2} \rceil \leq \kappa_i - 2x - 3$ bundles, each having at most two items with total size at most 1. Therefore, agent a_i can always place all items in H_j into at most $\kappa_i - 2x - 1$ bundles, each having at most two items with total size at most 1. $\square$

Now suppose that agents $a_1, \dots, a_k$ are matched and each matched agent a_j receives subsets of items B_j . In the next round, agent a_i is selected. We prove that agent a_i can still create $n - k$ parts as required in Line 4.

Lemma 3. *For each $k \in [n-1]$ and any k subsets $B_1, \dots, B_k$ allocated before a_i is selected, agent a_i can create $n - k$ pairwise disjoint parts $P_1, \dots, P_{n-k}$ following Line 4 such that for any $j \in [n - k]$, P_j can be partitioned into κ_i bundles, each with a total item size of at least $\frac{2}{3}$.*

Proof. Without loss of generality, assume that each matched B_j contains the large and medium items in j -th column of the n -arrangement (created in the first round). At each round, the selected agent is the one with the maximum total number of large and medium items among all remaining agents. Then for each $j \in [k]$, B_j includes a_i 's large and medium items in j -th column of the n -arrangement. Moreover for each $j \in [k]$, let $S_j \subseteq B_j$ be the small items for a_j , and then each item in $S_1 \cup \dots \cup S_k$ is also small for a_i . Thus for any $\ell \geq k + 1$, the large and medium items (from a_i 's perspective) in ℓ -th column of the n -arrangement are unallocated when a_i becomes the selected agent.

By the property of envy-free matching with regards the agents, for any $B_j, \forall j \in [k]$, it must be the case that B_j cannot be partitioned into κ_i bundles, each with a total item size of at least $\frac{2}{3}$. For each $j \in [k]$, by Lemma 2, the large and medium items for a_i in the j -th column of the n -arrangement, denoted as $F_j \cup H_j$, can be placed into at most $\kappa_i n$ bundles, each with total item size at most 1. Then we claim that

$$v_i(F_j) + v_i(H_j) + v_i(S_j) < \kappa_i.$$

To see the claim, assume for a contradiction that $v_i(F_j) + v_i(H_j) + v_i(S_j) \geq \kappa_i$. We next create a κ_i -partition of B_j such that each bundle has total item size at least $\frac{2}{3}$, contradicting the envy-free matching. First, place $F_j \cup H_j$ into at most κ_i bundles, each with a total item size of at most 1. Lemma 2 guarantees the existence of such a partition. Then, if a bundle has total item size less than $\frac{2}{3}$, we add small items such

that the total size of the bundle is at least $\frac{2}{3}$. Since each small item has size at most $\frac{1}{3}$, after adding small items, the total size of items in the bundle does not exceed 1. Therefore, $v_i(F_j) + v_i(H_j) + v_i(S_j) \geq \kappa_i$ implies that there are enough small items to make each bundle having total item size at least $\frac{2}{3}$, leading to the desired contradiction.

For a_i and each $\ell \geq k + 1$, place the large and medium items in ℓ -th column of the n -arrangement into at most κ_i bundles, each with total item size at most 1 (by Lemma 2). If the number of these bundles is less than $\kappa_i(n - k)$, create empty bundles until the total number reaches $\kappa_i(n - k)$. Next, if a bundle has total item size less than $\frac{2}{3}$, we add small items such that the total size of the bundle is at least $\frac{2}{3}$.

We below show that unallocated small items are enough for producing $\kappa_i(n - k)$ bundles, each with total item size at least $\frac{2}{3}$. For the sake of contradiction, assume that after adding all small items, there exists a bundle with total size less than $\frac{2}{3}$. Since each small item has size at most $\frac{1}{3}$ and a_i stops adding small items to a bundle when the total item size becomes at least $\frac{2}{3}$, then for each created bundle created, the total item size is at most 1. Therefore, the total size of items in $M \setminus (B_1 \cup \dots \cup B_k)$ is less than $\kappa_i(n - k) - \frac{1}{3}$. Combining the upper bound of total size of items in B_j 's, we have that the total size of items for a_i is less than $\kappa_i n - \frac{1}{3}$, a contradiction. $\square$

By Lemmas 2 and 3, at each round of the algorithm, the selected agent can create disjoint parts in Line 4. Thus, at each round, there exists at least one matched agent, and therefore, Algorithm 1 returns $\frac{2}{3}$ -CMMS allocations. We remark that given the oracle of valuations and MMS partitions, the algorithm runs in time polynomial in n and m by making n partitions queries and n^3 valuation queries (for building bipartite graphs). In fact, we can save the n^3 valuation queries by modifying the construction of bipartite graphs. Whether there exists an edge between S_j and a_ℓ can be determined as follows: Let a_ℓ fill κ_ℓ bins with S_j in the same way as that in the proof of Lemma 2, and if each bin has size at least $2/3$, then create an edge between a_ℓ and S_j . The above analysis still holds.

To complement the result, we present an example where the existing algorithm, namely Lone Divider, in [Aigner-Horev and Segal-Halevi, 2022] fails to output a $\frac{2}{3}$ -CMMS allocation. The differences between Lone Divider and Algorithm 1 are: (1) in Line 3, Lone Divider selects an arbitrary agent, and (2) in Line 4, Lone Divider does not impose requirements on the $|\mathcal{N}|$ parts created by the selected agent.

Consider an instance with 20 agents. Focusing on agent a_n and assuming $\kappa_n = 3$, agent a_n 's MMS partition has 60 bundles, each composed of three items with size $\frac{2}{3} - \epsilon$, $\frac{1}{3} - \epsilon$, and 2ϵ , where $\epsilon > 0$ is arbitrarily small. It is possible that a_n is the last agent and the first 19 allocated bundles are: 12 bundles each with 5 items of size $\frac{2}{3} - \epsilon$, and 7 bundles each with 8 items of size $\frac{1}{3} - \epsilon$. From a_n 's perspective, none of these 19 bundles can satisfy her, as they cannot form 3 bundles each with a total size of at least $\frac{2}{3}$. Thus, this situation can happen in Lone Divider. There remains 4 items each with size $\frac{1}{3} - \epsilon$ and 60 items each with size 2ϵ . It is easy to see that with these remaining items, agent a_n cannot form 3 bundles each with a total size of at least $\frac{2}{3}$.

4 Ordinal Approximation for the Bin Covering Model

In this section, we elaborate on the algorithm that computes $\frac{3}{4}$ -OMMS allocations for the bin covering model. That is, for any bin covering instance, we can compute an allocation such that each agent a_i can cover at least $\frac{3}{4}\kappa_i - \frac{7}{4}$ bins using her allocated items.

Theorem 2. *For any bin covering instance, a $\frac{3}{4}$ -OMMS allocation exists.*

For ease of presentation, in this and the next sections, for each agent a_i , let $L_i = \{e_j \in M \mid s_i(e_j) > \frac{1}{2}\}$ be the set of *large* items or *L-item(s)* for short and $M_i = \{e_j \in M \mid \frac{1}{2} \geq s_i(e_j) > \frac{1}{3}\}$ be the set of *medium* items or *M-item(s)*. Moreover, we say an item e is *small* or *S-item* for a_i if it is neither large nor medium (i.e., $s_i(e) \leq \frac{1}{3}$).

4.1 The Algorithm

We employ the classic Round-Robin algorithm to compute the desired allocation $\mathbf{A} = (A_1, \dots, A_n)$. That is, the agents sequentially pick items from the unallocated items, each time choosing the one with the largest size. Since the instances we focus on are IDO, the items allocated to each agent a_i are actually $A_i = \{e_{i+jn} \mid j \in \mathbb{N}, i+jn \leq m\}$. The technical contribution here is to show how each agent a_i distributes her allocated items such that at least $\frac{3}{4}\kappa_i - \frac{7}{4}$ bins can be covered, i.e., $v_i(A_i) \geq \frac{3}{4}\kappa_i - \frac{7}{4}$ for every $a_i \in N$. By the characteristic of the Round-Robin algorithm, we have $v_i(A_i) \geq v_i(A_n)$ for every a_i . Thus, it suffices to prove that $v_i(A_n) \geq \frac{3}{4}\kappa_i - \frac{7}{4}$ holds for every a_i . In the following, we shall focus on a fixed agent a_i and the same arguments apply to the other agents.

We first construct a set of synthesized items $P_i = \{o_1, \dots, o_k\}$ (see Algorithm 2), which consists of all a_i 's large items in A_n and synthesizes as many pairs of medium items in A_n as possible. The last remaining medium item (which exists when $|M_i \cap A_n|$ is odd) is either directly put into P_i or combined with the smallest synthesized item, depending on the sum of its size and that of the smallest item in P_i . In the following, instead of directly considering the large and medium items in A_n , we will first distribute the synthesized items in P_i and then replace them with their corresponding items in A_n .

Algorithm 2 Constructing Synthesized Items for Agent a_i

Input: Bundle A_n , agent a_i 's large items L_i and medium items M_i .

Output: A set of synthesized items $P_i = \{o_1, \dots, o_k\}$.

- 1: Initialize $P_i \leftarrow A_n \cap L_i$.
 - 2: **while** $|A_n \cap M_i| \geq 2$ **do**
 - 3: Arbitrarily pick two items e_a, e_b in $A_n \cap M_i$.
 - 4: Synthesize e_a and e_b into one item o , i.e., $s_i(o) = s_i(e_a) + s_i(e_b)$.
 - 5: $P_i \leftarrow P_i \cup \{o\}$, $A_n \leftarrow A_n \setminus \{e_a, e_b\}$.
 - 6: **end while**
 - 7: **if** $|A_n \cap M_i| = 1$ **then**
 - 8: Let e^* be the only item in $A_n \cap M_i$ and o^* be the item with the smallest size in P_i .
 - 9: **if** $s_i(e^*) + s_i(o^*) \geq 1$ **then**
 - 10: $P_i \leftarrow P_i \cup \{e^*\}$.
 - 11: **else**
 - 12: Synthesize e^* and o^* into one item o' , i.e., $s_i(o') = s_i(e^*) + s_i(o^*)$.
 - 13: $P_i \leftarrow P_i \setminus \{o^*\} \cup \{o'\}$.
 - 14: **end if**
 - 15: **end if**
 - 16: Sort the items in P_i in decreasing order of the sizes and rename them as $o_1, \dots, o_k$.
 - 17: **return** P_i .
-

We then present how to cover bins using the items in A_n . We consider two cases with respect to the number of synthesized items in P_i .

Case 1: $k \leq \frac{3}{4}\kappa_i - \frac{1}{2}$

For this case, we distribute the items in two steps:

- First, we distribute each of the synthesized items in P_i to one bin and cover these bins using the small items in A_n .
- Second, we greedily create new bins and cover them using the small items in A_n .

Case 2: $k > \frac{3}{4}\kappa_i - \frac{1}{2}$

For this case, we let $z = \lceil \frac{3}{4}\kappa_i - \frac{7}{4} \rceil$ and distribute the items in four steps:

- First, we cover $k - z$ bins using the smallest $2 \cdot (k - z)$ items in P_i .

- Second, as long as there remain two items in P_i whose sizes are at most $\frac{2}{3}$, we distribute the smallest two of them to cover a new bin.
- Third, we distribute each of the remaining items in P_i to one bin and cover these bins using the small items in A_n . Note that here we specify the order of the bins to be covered, that is, in increasing order of the sizes of the P_i items that they already have, starting from the one with the second smallest P_i item. We lastly cover the bin with the smallest P_i item.
- Fourth, we greedily create new bins and cover them using the small items in A_n .

4.2 The Proof of Theorem 2

We are now ready to prove Theorem 2.

Proof of Theorem 2. We will show that the above algorithm can cover at least $\frac{3}{4}\kappa_i - \frac{7}{4}$ bins for a_i in both of the above two cases.

Case 1: $k \leq \frac{3}{4}\kappa_i - \frac{1}{2}$.

For this case, we first have the following claim.

Claim 1. *There are enough small items in A_n to cover the bins in the first step.*

Proof of Claim 1. For the sake of contradiction, suppose that we cannot cover all the bins in the first step. Since the size of each small item is at most $\frac{1}{3}$ and the item that cover a bin in the first step is small, the total size that a covered bin in the first step is at most $\frac{4}{3}$. Thus we have an upper bound of the total size of A_n to a_i , i.e., $s_i(A_n) < \frac{4}{3} \cdot (k-1) + 1 \leq \kappa_i - 1$. However, by the design of the Round-Robin algorithm, we have the following lower bound of the total size of A_n to a_i ,

$$s_i(A_n) \geq \frac{1}{n} \cdot (s_i(M) - (n-1)) \geq \kappa_i - 1 + \frac{1}{n},$$

where the last inequality is because $s_i(M) \geq n \cdot \kappa_i$ by the definition of MMS, a contradiction. $\square$

By Claim 1, each bin in the first step is covered and has a total size of at most $\frac{4}{3}$. Since all items distributed in the second step are small, each bin in the second step also has a total size of at most $\frac{4}{3}$. Thus, letting ALG be the number of covered bins in the above two steps, we have

$$\frac{4}{3}ALG + 1 > s_i(A_n). \quad (1)$$

Besides, by the design of the Round-Robin algorithm, we have the following lower bound of $s_i(A_n)$,

$$s_i(A_n) \geq \frac{1}{n} \cdot (s_i(M) - (n-1)) \geq \kappa_i - 1 + \frac{1}{n}, \quad (2)$$

where the last inequality is because the items M can cover at least $n \cdot \kappa_i$ bins in a_i 's MMS partition. Combining Inequalities (1) and (2), we have

$$v_i(A_n) \geq ALG > \frac{3}{4}\kappa_i - \frac{3}{2} > \frac{3}{4}\kappa_i - \frac{7}{4},$$

which completes the proof of Theorem 2 for Case 1.

Case 2: $k > \frac{3}{4}\kappa_i - \frac{1}{2}$

For this case, we define more notations here. Let l be the number of undistributed items in P_i after the first two steps. Note that these l items are the l largest ones in P_i (i.e., $o_1, \dots, o_l$). For each $j \in [l]$, let $\Delta_j = 1 - s_i(o_j)$. Note that $\Delta_1 \leq \dots \leq \Delta_l$, $\Delta_{l-1} < \frac{1}{3}$, and Δ_l may be at least $\frac{1}{3}$ but at most $\frac{2}{3}$. Then we have an upper bound of the total size of the items in A_n that are large or medium for agent a_i .

Claim 2. $s_i(A_n \cap (L_i \cup M_i)) \leq \sum_{j=1}^{l-1} (1 - \Delta_j) + (k - l + 1) \cdot (1 - \Delta_l)$.

Proof of Claim 2. Recall the process of constructing the set P_i and we know that the total size of the items in P_i equals to the total size of the large and medium items in A_n . We can derive the upper bound of the total size of P_i as the right side of the inequality by scaling the sizes of $o_{l+1}, \dots, o_k$ to that of o_l (i.e., $1 - \Delta_l$) and summing up these scaled sizes and the sizes of $o_1, \dots, o_l$. $\square$

We further consider two subcases.

Subcase 2-1: No bin is created in the second step

For this case, there are $2z - k$ undistributed items in P_i after the first two steps, i.e., $l = 2z - k$. We are now going to show that there are enough small items in A_n to cover the bins in the third step. If it is, together with the $k - z$ covered bins in the first step, we have at least $z = \lceil \frac{3}{4}\kappa_i - \frac{7}{4} \rceil$ covered bins in total, which will prove Theorem 2 for Subcase 2-1.

We first provide a lower bound of the total size of the small items in A_n to agent a_i .

Claim 3. $s_i(A_n \cap S_i) \geq \sum_{j=1}^{l-1} \Delta_j + (k - l + 1) \cdot \Delta_l + \kappa_i - k - \frac{4}{3} + \frac{4}{3n}$.

Proof of Claim 3. We first derive an upper bound of the total size of the items in M that are large or medium for a_i . Recall that $A_n = \{e_{jn} \mid j \in \mathbb{N}, jn \leq m\}$. For each large or medium item e_{jn} in A_n , let Q_j be the set of the items between e_{jn} and $e_{(j+1)n}$ (or e_m if $(j+1)n > m$), i.e., $Q_j = \{e_{jn+t} \mid 1 \leq t \leq n-1, jn+t \leq m\}$. Notice that the large and medium items in A_n and the items in their corresponding sets Q_j s cover all the items in M that are large or medium for a_i except the first $n-1$ largest items. By the characteristic of the Round-Robin algorithm, we know that each item in Q_j has a size of at most e_{jn} . Therefore, we have that

$$\begin{aligned} s_i(L_i \cup M_i) &\leq n - 1 + n \cdot s_i(A_n \cap (L_i \cup M_i)) \\ &\leq n - 1 + n \cdot \sum_{j=1}^{l-1} (1 - \Delta_j) + n \cdot (k - l + 1) \cdot (1 - \Delta_l), \end{aligned}$$

where $n - 1$ is the upper bound of the total size of the first $n - 1$ largest items and the second inequality is due to Claim 2. Then we have a lower bound of total size of the items in M that are small for a_i ,

$$\begin{aligned} s_i(S_i) &= s_i(M) - s_i(L_i \cup M_i) \geq n\kappa_i - n + 1 - n \cdot \sum_{j=1}^{l-1} (1 - \Delta_j) - n \cdot (k - l + 1) \cdot (1 - \Delta_l) \\ &= n \cdot \sum_{j=1}^{l-1} \Delta_j + n \cdot (k - l + 1) \cdot \Delta_l + (\kappa_i - k)n + 1 - n. \end{aligned}$$

By the Round-Robin algorithm, we have that

$$s_i(A_n \cap S_i) \geq \frac{1}{n} \cdot (s_i(M \cap S_i) - \frac{1}{3}(n-1)) \geq \sum_{j=1}^{l-1} \Delta_j + (k - l + 1) \cdot \Delta_l + \kappa_i - k - \frac{4}{3} + \frac{4}{3n},$$

which completes the proof of the claim. $\square$

Next, consider when the largest small item in A_n has a size smaller than Δ_{l-1} . In order to cover a bin with size $1 - \Delta$, the total size of small items we need does not exceed $\Delta + \Delta_{l-1}$. Therefore, the total size of small items we need to cover the bins in the third step does not exceed

$$\sum_{j=1}^l (\Delta_j + \Delta_{l-1}) = \sum_{j=1}^{l-1} \Delta_j + l \cdot \Delta_{l-1} + \Delta_l.$$

By Claim 3, it suffices to show that

$$\sum_{j=1}^{l-1} \Delta_j + (k - l + 1) \cdot \Delta_l + \kappa_i - k - \frac{4}{3} \geq \sum_{j=1}^{l-1} \Delta_j + l \cdot \Delta_{l-1} + \Delta_l,$$

which is equivalent to $l \cdot \Delta_{l-1} - (k-l) \cdot \Delta_l \leq \kappa_i - k - \frac{4}{3}$. Since $\Delta_l \geq \Delta_{l-1}$ and $k \geq l$, it suffices to show

$$(2l-k) \cdot \Delta_{l-1} \leq \kappa_i - k - \frac{4}{3} \Leftrightarrow \Delta_{l-1} \leq \frac{\kappa_i - k - \frac{4}{3}}{4z - 3k}.$$

Since $\Delta_{l-1} < \frac{1}{3}$, it suffices to show

$$\Delta_{l-1} \leq \frac{\kappa_i - k - \frac{4}{3}}{4z - 3k} \geq \frac{1}{3} \Leftrightarrow z = \lceil \frac{3}{4}\kappa_i - \frac{7}{4} \rceil \leq \frac{3}{4}\kappa_i - 1,$$

which is true since κ_i is an integer.

Now, consider when the largest small item in A_n has a size at least Δ_{l-1} . More generally, let t be the largest integer in $[l-1]$ such that the j -th largest small item in A_n has a size at least Δ_{l-j} for every $j \in [t]$. This means that for every $j \in [t]$, the j -th largest small item alone is enough to cover the bin with item o_{l-j} . Since the sizes of the other small items are smaller than Δ_{l-t-1} , the total size of small items we need to cover other bins that have already had a size of $1 - \Delta$ does not exceed $\Delta + \Delta_{l-t-1}$. Therefore, in order to cover the bins in the third step, the total size of small items we need does not exceed

$$\frac{t}{3} + \sum_{j=1}^{l-t-1} (\Delta_j + \Delta_{l-t-1}) + \Delta_l + \Delta_{l-t-1} = \frac{t}{3} + \sum_{j=1}^{l-t-1} \Delta_j + (l-t) \cdot \Delta_{l-t-1} + \Delta_l,$$

where $\frac{t}{3}$ is the upper bound of the total size of the first t largest small items in A_n . By Claim 3, it suffices to show that

$$\sum_{j=1}^{l-1} \Delta_j + (k-l+1) \cdot \Delta_l + \kappa_i - k - \frac{4}{3} \geq \frac{t}{3} + \sum_{j=1}^{l-t-1} \Delta_j + (l-t) \cdot \Delta_{l-t-1} + \Delta_l,$$

which is equivalent to

$$(l-t) \cdot \Delta_{l-t-1} - \sum_{j=l-t}^{l-1} \Delta_j - (k-l) \cdot \Delta_l \leq \kappa_i - k - \frac{4}{3} - \frac{t}{3}.$$

Since $\Delta_{l-t-1} \leq \dots \leq \Delta_l$, it suffices to show that

$$(l-t) \cdot \Delta_{l-t-1} - t \cdot \Delta_{l-t-1} - (k-l) \cdot \Delta_{l-t-1} \leq \kappa_i - k - \frac{4}{3} - \frac{t}{3},$$

which is equivalent to

$$\Delta_{l-t-1} \leq \frac{\kappa_i - k - \frac{4}{3} - \frac{t}{3}}{4z - 3k - 2t},$$

which is true since $\Delta_{l-t-1} < \frac{1}{3}$ and $\frac{\kappa_i - k - \frac{4}{3} - \frac{t}{3}}{4z - 3k - 2t} > \frac{1}{3}$.

Subcase 2-2: Some bins are created in the second step. For this case, first observe that there are enough small items in A_n such that all the bins in the third step can be satisfied. This is because compared to Subcase 2-1, there are now less undistributed items in P_i after the first two steps and the lower bound of the total size of a_i 's small items in A_n does not change (i.e., Claim 3 still applies here). Also observe that the size of each item in P_i that is distributed to a bin in the first two steps is at most $\frac{2}{3}$. This means that each covered bin in the first two steps has a size of at most $\frac{4}{3}$. Note that the total size of the items distributed to each covered in the last two steps is obviously at most $\frac{4}{3}$ since the last item given to each bin is a small one. Following the same arguments in Case 1, we conclude that there are at least $\frac{3}{4}\kappa_i - \frac{7}{4}$ covered bins. $\square$

We remark that although the analysis uses MMS partitions and valuations, the algorithm does not require any oracle and runs in polynomial time with respect to n and m .

5 Ordinal Approximation for the Bin Packing Model

In this section, we consider the ordinal approximation for the bin packing model and present the existence of $\frac{4}{3}$ -OMMS allocations.

Theorem 3. *For any bin packing instance, a $\frac{4}{3}$ -OMMS allocation exists.*

5.1 The Algorithm

We now informally introduce the algorithm. First, hypothetically arrange all items into n columns (also called n -arrangement) based on the descending order of size, such that the j -th column contains items $e_j, e_{j+n}, e_{j+2n}, \dots$ for every $j \in [n]$. The algorithm proceeds in n rounds. In each round j , we focus on the items in j -th column and let a selected agent leave with a bag of items, which is determined through two phases: bag initialization and bag filling. In the bag initialization phase (Line 4), based on the ascending order of item indices, items from the underlying column are added to the bag as long as they are large for at least one remaining agent. In the bag filling phase (Line 8), if there exist agents who believe that the total size of items in the bag is no greater than $\frac{1}{n}$ fraction of the total size of all items and still have unallocated small items, let one of them be the selected agent and add the smallest unallocated item into the bag. At the end of each round, the selected agent leaves with all items in the bag. The formal description is provided in Algorithm 3. Our algorithm borrows the idea of the algorithm in Li *et al.* [2023], but we classify large and medium items differently, which necessitates different approaches in the analysis.

Algorithm 3 Computing $\frac{4}{3}$ -OMMS allocations

Input: A bin packing instance $\mathcal{I} = (N, M, \mathbf{s})$.

Output: Allocation $\mathbf{A} = (A_1, \dots, A_n)$.

- 1: Initialize $R \leftarrow M$.
 - 2: **for** $j = 1, \dots, n$ **do**
 - 3: Initialize $B_j \leftarrow \emptyset, t \leftarrow j, a_k \leftarrow \emptyset$.
 - 4: **while** $e_t \in R$ and there exists agent a_i such that $e_t \in L_i \cup M_i$ **do**
 - 5: $B_j \leftarrow B_j \cup \{e_t\}, R \leftarrow R \setminus \{e_t\}, t \leftarrow t + n$.
 - 6: $a_k \leftarrow a_i$.
 - 7: **end while**
 - 8: **while** there exists agent a_i such that $s_i(B_j) \leq \frac{1}{n}s_i(M)$ and there still exist unallocated S -items for agent a_i **do**
 - 9: $a_k \leftarrow a_i$.
 - 10: Let $e \in R$ be the item with the largest index and $B_j \leftarrow B_j \cup \{e\}, R \leftarrow R \setminus \{e\}$.
 - 11: **end while**
 - 12: $A_k \leftarrow B_j, N \leftarrow N \setminus \{a_k\}$.
 - 13: **end for**
-

5.2 The Proof of Theorem 3

We will first prove that all items can be allocated. Before that, let us present several propositions regarding the Algorithm 3.

Proposition 1. *For an agent a_i , suppose that e^* is the last item added to A_i . If e^* is small for i , then $s_i(A_i \setminus \{e^*\}) \leq \frac{1}{n} \cdot s_i(M)$.*

Proof. Since e^* is small for a_i , then e^* must be added to A_i in the bag filling phase. By Line 8, agent a_i has total size no greater than $\frac{1}{n}s_i(M)$ right before adding e^* , and hence, $s_i(A_i \setminus \{e^*\}) \leq \frac{1}{n} \cdot s_i(M)$. $\square$

Proposition 2. For agents a_i and a_j , suppose that a_j receives her bundle A_j in some round earlier than a_i . If after the round in which a_j receives A_j , there remains unallocated S -items for a_i , then $s_i(A_j) > \frac{1}{n}s_i(M)$.

These propositions directly follow from the condition in Line 8 of Algorithm 3.

Lemma 4. All items are allocated at the termination of Algorithm 3.

Proof. Suppose that a_n is the last agent to receive a bundle. If at the beginning of round n (Line 2), all unallocated items are large or medium for a_n , then all of them will be added to the bag in Line 4 and then will be allocated to a_n after the bag filling phase. Focus on the case when, at the beginning of round n , there exist unallocated S -items for a_n . By Proposition 2, it holds that for each $j \in [n-1]$, $s_n(A_j) > \frac{1}{n}s_n(M)$, which implies that the total size of remaining items for a_n at the beginning of round n is less than $\frac{1}{n}s_n(M)$. Therefore, by Line 8, all items will be added to the bag and then allocated to a_n . $\square$

Next we present the performance guarantee of the returned allocation. Focus on agent a_i and define $H_i := L_i \cup M_i$. Suppose that $c_i(H_i) = |H_i| - k$. Then we claim that the $2k$ items with the smallest size in H_i can be packed into k bins as follows: for each $\ell \in [k]$, pack the item with the ℓ -th smallest size and the item with the $(2k+1-\ell)$ -th smallest size into a bin. To see the claim, we begin with a solution achieving $c_i(H_i)$ and exchange items as follows: if there exist e, e' such that $s_i(e) < s_i(e')$ with e packed solely in a bin and e' packed with another item, then exchange their positions in the solution. Then it is easy to see that the $2k$ items with the smallest size in H_i can be packed into κ_i bins.

Let $H_i^* \subseteq H_i$ be the set of $2k$ items with the smallest size in H_i , and let $H_i' := H_i \setminus H_i^*$. We below prove that items in $A_i \cap H_i$ can be packed into $\kappa_i + 2$ bins.

Lemma 5. For agent a_i , it holds that $c_i(A_i \cap H_i) \leq \kappa_i + 2$.

Proof. We create a packing solution for $A_i \cap H_i$. For each $e \in A_i \cap H_i'$, it will be solely packed into a bin. For items in $A_i \cap H_i^*$, due to the construction of the n -arrangement, one can verify that the second largest size item can be packed into a bin with the t -th largest size item; the third largest size item can be packed into a bin with the $(t-1)$ -th largest size item; and so on, where $t = |A_i \cap H_i^*|$. Thus, excluding the largest size item in $A_i \cap H_i^*$, we can pack the rest items in $\lceil \frac{|A_i \cap H_i^*| - 1}{2} \rceil$ bins, and therefore, $c_i(A_i \cap H_i^*) \leq \lceil \frac{|A_i \cap H_i^*| - 1}{2} \rceil + 1$.

The construction of n -arrangement implies that $|A_i \cap H_i'| \leq \lfloor \frac{|H_i'|}{n} \rfloor + 1$ and $|A_i \cap H_i^*| \leq \lfloor \frac{|H_i^*|}{n} \rfloor + 1$. Then we have

$$\begin{aligned} c_i(A_i \cap H_i) &\leq |A_i \cap H_i'| + \lceil \frac{|A_i \cap H_i^*| - 1}{2} \rceil + 1 \\ &\leq \lfloor \frac{|H_i'|}{n} \rfloor + \lceil \frac{\lfloor \frac{|H_i^*|}{n} \rfloor}{2} \rceil + 2 \\ &\leq \frac{1}{n} \cdot \left(|H_i'| + \frac{|H_i^*|}{2} \right) + \frac{5}{2}. \end{aligned}$$

As $|H_i'| + \frac{|H_i^*|}{2} = c_i(H_i) \leq \kappa_i n$ and $c_i(A_i \cap H_i)$ is an integer, the above inequality implies $c_i(A_i \cap H_i) \leq \kappa_i + 2$. $\square$

Now we are ready to prove Theorem 3.

Proof of Theorem 3. If A_i does not contain S -items for a_i , then by Lemma 5, agent a_i is 1-OMMS. Then we focus on the case when A_i contains S -items for a_i and let e^* be the last S -item added to A_i in the bag filling phase.

We now pack items in A_i into $\kappa_i + 2$ bins, such that the total size of items in a bin is at most $\frac{4}{3}$. Consider the following way of packing

- Step 1: Pack large and medium items in A_i into $\kappa_i + 2$ bins, each with a total size of at most 1.
- Step 2: If there exists a bin with total size no greater than 1, add S -items to it until the total size just exceeds 1. Repeat the process until all S -items are allocated.

Lemma 5 ensures the packing solution for Step 1. For Step 2, by Proposition 1, we have $s_i(A_i \setminus \{e^*\}) \leq \frac{1}{n}s_i(M) = \kappa_i$, implying $s_i(A_i) \leq \kappa_i + \frac{1}{3}$ as e^* is small. Thus, after allocating all S -items, the total size of each bin is at most $\frac{4}{3}$; otherwise, $s_i(A_i) > \frac{4}{3}(\kappa_i + 2)$, a contradiction.

Begin with the constructed $\kappa_i + 2$ bins. If there exists a bin with total size at most $\frac{2}{3}$, reassign an S -item from the bin with total size larger than 1 to it. Repeat the process, and at the end, one of the following happens: (i) each bin has total size at most 1, or (ii) each bin has total size at least $\frac{2}{3}$. If (i) happens, then the theorem statement holds. If (ii) happens, we first prove that at the end, there are at most $\kappa - 4$ bins, each with total size larger than 1. Assume for a contradiction. As each bin has total size at least $\frac{2}{3}$, the total size of items in A_i is

$$s_i(A_i) > (\kappa_i - 3) \times 1 + 5 \times \frac{2}{3} = \kappa_i + \frac{1}{3} \geq s_i(A_i),$$

deriving a contradiction. For each bin with total size larger than 1, we remove the largest size S -item from the bin, ensuring that the total size is at most 1 after the removal (by Step 2). Then we pack these removed S -items into extra bins, each containing 3 small items. Since each S -item has size at most $\frac{1}{3}$, each of these bins has total size no greater than 1. Given that there are at most $\kappa_i - 4$ removed S -items, we have

$$c_i(A_i) \leq \kappa_i + 2 + \lceil \frac{\kappa_i - 4}{3} \rceil \leq \frac{4}{3}\kappa_i + \frac{4}{3},$$

completing the proof. $\square$

The algorithm does not require any oracle and runs in polynomial time with respect to n and m . We remark that the $\frac{4}{3}$ -CMMS allocations in [Li *et al.*, 2025] do not translate to $\frac{4}{3}$ -OMMS allocations. In a $\frac{4}{3}$ -CMMS allocation, each agent i 's bundle can be packed into κ_i bins with capacity $\frac{4}{3}$. If in a $\frac{4}{3}$ -CMMS allocation, each agent i 's bundle can be packed such that: whenever the total size of a bin exceeds 1, there exists a small item in the bin whose removal reduces the total item size to at most 1, then removing these small items and packing them into $\frac{1}{3}\kappa_i$ bins with capacity 1 yields a $\frac{4}{3}$ -OMMS allocation. The $\frac{4}{3}$ -CMMS allocation computed in [Li *et al.*, 2025] do not possess the property just described, and therefore, cannot be converted into $\frac{4}{3}$ -OMMS allocations.

6 Conclusion

We investigate the cardinal approximation of MMS fair allocation for groups. From a technical standpoint, our paper contributes novel analyses for both goods and chores. Our result for chores improves the best-known approximation in Li *et al.* [2023]. On a conceptual level, our work provides a more general fair allocation model with non-atomic agents which generalizes the traditional model with atomic agents. Although the individuals' valuations can be integrated into a unified utility function of a group, the model uncovers several additional insights. First, the utility function (i.e., the types of social welfare) of each group can be adjusted for different scenarios. Second, the group model provides a new dimension to relax MMS fairness, which does not modify the value of MMS and is natural in the situation when we want to maximize the number of satisfied agents. Third, we assume the agents in a group have the same valuation function. This assumption is natural since agents in the same group typically encounter similar conditions and resources that shape their preferences. This assumption also builds the relationship between our model and the atomic model with bin-packing and bin-covering valuation functions. Looking ahead, a promising direction for future research involves extending our model to accommodate heterogeneous valuations.

References

- Elad Aigner-Horev and Erel Segal-Halevi. Envy-free matchings in bipartite graphs and their applications to fair division. *Inf. Sci.*, 587:164–187, 2022.
- Hannaneh Akrami and Jugal Garg. Breaking the $3/4$ barrier for approximate maximin share. In *SODA*, pages 74–91. SIAM, 2024.
- Hannaneh Akrami and Nidhi Rathi. Achieving maximin share and EFX/EF1 guarantees simultaneously. In *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, pages 13529–13537. AAAI Press, 2025.
- Hannaneh Akrami, Kurt Mehlhorn, Masoud Seddighin, and Golnoosh Shahkarami. Randomized and deterministic maximin-share approximations for fractionally subadditive valuations. In *NeurIPS*, 2023.
- Hannaneh Akrami, Noga Alon, Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, and Ruta Mehta. EFX: A simpler approach and an (almost) optimal guarantee via rainbow cycle number. *Oper. Res.*, 73(2):738–751, 2025.
- Susan F. Assmann, David S. Johnson, Daniel J. Kleitman, and Joseph Y.-T. Leung. On a dual version of the one-dimensional bin packing problem. *J. Algorithms*, 5(4):502–525, 1984.
- Siddharth Barman and Sanath Kumar Krishnamurthy. Approximation algorithms for maximin fair division. *ACM Trans. Economics and Comput.*, 8(1):5:1–5:28, 2020.
- Siddharth Barman, Arindam Khan, Sudarshan Shyam, and K. V. N. Sreenivas. Guaranteeing envy-freeness under generalized assignment constraints. In *EC*, pages 242–269. ACM, 2023.
- Umang Bhaskar, A. R. Sricharan, and Rohit Vaish. On approximate envy-freeness for indivisible chores and mixed resources. In *APPROX-RANDOM*, volume 207 of *LIPICs*, pages 1:1–1:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. *ACM Trans. Economics and Comput.*, 7(3):12:1–12:32, 2019.
- Edward G Coffman Jr, János Csirik, Gábor Galambos, Silvano Martello, and Daniele Vigo. Bin packing approximation algorithms: survey and classification. In *Handbook of combinatorial optimization*, pages 455–531. Springer, 2013.
- Sarfaraz Equbal, Rohit Gurjar, Yatharth Kumar, Swaprava Nath, and Rohit Vaish. Fair interval scheduling of indivisible chores. *CoRR*, abs/2402.04353, 2024.
- Uriel Feige, Ariel Sapir, and Laliv Tauber. A tight negative example for MMS fair allocations. In *WINE*, volume 13112 of *Lecture Notes in Computer Science*, pages 355–372. Springer, 2021.
- Duncan Foley. Resource allocation and the public sector. *Yale Economic Essays*, pages 45–98, 1967.
- Jugal Garg, Aniket Murhekar, and John Qin. Fair division of indivisible chores via earning restricted equilibria. *CoRR*, abs/2407.03318, 2024.
- Jugal Garg, Aniket Murhekar, and John Qin. Constant-factor EFX exists for chores. In *STOC*, pages 1580–1589. ACM, 2025.

- Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods: Improvements and generalizations. In *EC*, pages 539–556. ACM, 2018.
- Laurent Gourvès, Jérôme Monnot, and Lydia Tlilane. Near fairness in matroids. In *ECAI*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 393–398. IOS Press, 2014.
- Hadi Hosseini, Shivika Narang, and Tomasz Was. Fair distribution of delivery orders. *Artificial Intelligence*, page 104389, 2025.
- Xin Huang and Pinyan Lu. An algorithmic framework for approximating maximin share allocation of chores. In *EC*, pages 630–631. ACM, 2021.
- Xin Huang and Erel Segal-Halevi. A reduction from chores allocation to job scheduling. In *EC*, page 908. ACM, 2023.
- Halvard Hummel. Maximin shares in hereditary set systems. *ACM Trans. Econ. Comput.*, 13(3):1–33, 2025.
- David Kurokawa, Ariel D. Procaccia, and Junxing Wang. When can the maximin share guarantee be guaranteed? In *AAAI*, pages 523–529. AAAI Press, 2016.
- David Kurokawa, Ariel D. Procaccia, and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. *J. ACM*, 65(2):8:1–8:27, 2018.
- Bo Li, Fangxiao Wang, and Yu Zhou. Fair allocation of indivisible chores: Beyond additive costs. In *NeurIPS*, 2023.
- Bo Li, Fangxiao Wang, and Shiji Xing. Settling the maximin share fairness for scheduling among groups of machines. In *ICML*, Proceedings of Machine Learning Research. PMLR, 2025.
- Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *EC*, pages 125–131. ACM, 2004.
- Silvano Martello and Paolo Toth. Lower bounds and reduction procedures for the bin packing problem. *Discret. Appl. Math.*, 28(1):59–70, 1990.
- Masoud Seddighin and Saeed Seddighin. Improved maximin guarantees for subadditive and fractionally subadditive fair allocation problem. *Artif. Intell.*, 327:104049, 2024.
- Masoud Seddighin and Saeed Seddighin. Beating the logarithmic barrier for the subadditive maximin share problem. *arXiv preprint arXiv:2506.05613*, 2025.
- Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
- Gilad Ben Uziah and Uriel Feige. On fair allocation of indivisible goods to submodular agents. *CoRR*, abs/2303.12444, 2023.
- Shengwei Zhou and Xiaowei Wu. Approximately EFX allocations for indivisible chores. *Artif. Intell.*, 326:104037, 2024.