Centralized Group Equitability and Individual Envy-Freeness in the Allocation of Indivisible Items

Ying Wang[1] *, **Jiaqian Li**[2] *, **Tianze Wei**[3]†, **Hau Chan**[4], **Minming Li**[3]

[1]Columbia University
[2] Boston University
[3] City University of Hong Kong
[4]University of Nebraska Lincoln
yw4360@columbia.edu, jiaqian@bu.edu, t.z.wei-8@my.cityu.edu.hk, hchan3@unl.edu, minming.li@cityu.edu.hk

Abstract

We study the fair allocation of indivisible items to groups of agents from the perspectives of both the agents and a centralized allocator. In our setting, the centralized allocator aims to ensure that the allocation is fair both among the groups and between individual agents. This setting applies to many real-world scenarios, such as when a school administrator allocates resources (e.g., office spaces and supplies) to staff members within departments or when a city council allocates limited housing units to families in need across different communities. To ensure fairness between agents, we consider the classical notion of *envy-freeness* (EF). To ensure fairness among groups, we introduce the notion of *centralized group equitability* (CGEQ), which captures fairness for groups from the centralized allocator's perspective. Because an EF or CGEQ allocation does not always exist in general, we consider their natural relaxations: *envy-freeness to one item* (EF1) and *centralized group equitability up to one item* (CGEQ1). For different classes of valuation functions of the agents and the centralized allocator, we show that allocations satisfying both EF1 and CGEQ1 always exist, and we design efficient algorithms to compute such allocations. We also consider the *centralized group maximin share* (CGMMS) from the centralized allocator's perspective as a group-level fairness objective with EF1 for agents, and present several results.

Extended version — http://arxiv.org/abs/2511.07984

1 Introduction

Fair division of indivisible items often deals with fairly allocating a set of (discrete or indivisible) items to a set of agents who have preferences over the items. Due to both practical and theoretical interests, fair division of indivisible items has received considerable attention in various research communities, such as economics, mathematics, and computer science, for much of the past century (Moulin 2004; Amanatidis et al. 2023). In practice, fair division of indivisible items has many real-world applications ranging from course allocation (i.e., for allocating schedules of courses to students) to goods division (i.e., dividing artworks or furniture among individuals), in which the Course Match (Budish et al. 2017) mechanism and the Spliddit (Goldman and Procaccia 2015) online platform have been developed, to provide fair allocations subject to agents' preferences and appropriate fairness notions (e.g., envy-freeness and its relaxations) for the respective applications. In theory, fair division of indivisible items has led to the development of numerous notions such as envy-freeness up to one/any item, proportionality up to one/any item, and maximin share fairness for quantifying fairness, algorithms (e.g., round-robin or envy-graph procedure (Lipton et al. 2004)) for providing (approximately) fair allocations, and techniques for (partially) characterizing the existence of a fair allocation (Amanatidis et al. 2023).

Allocator's Preference. A main drawback of existing studies of fair division of indivisible items is the lack of consideration from the allocator's perspective, who is responsible for implementing the allocation and has preferences over how the items should be allocated to the set of agents (Bu et al. 2023). As a result, (Bu et al. 2023) initiated the study of fair division of indivisible items from both the allocator's and agents' perspectives, where each agent has a valuation preference over the items, and the allocator has a separate valuation preference for each agent over the items (specifying their internal values for the agent receiving the items). Moreover, they focus on allocations that satisfy relaxations of envy-freeness between agents under both the allocator's valuation for each agent and agents' valuations.

As a motivating example, (Bu et al. 2023) discussed the situations where the government (as the allocator) needs to distribute education resources (e.g., funding and staff members) to different schools (as agents) in which the schools and the government have separate preferences over the educational resources based on their needs and macroeconomic policy for schools, respectively. In addition, (Bu et al. 2023) provided examples where a company allocates resources to different departments, an advisor allocates tasks or projects to students, and conference organizers allocate papers to reviewers that require the consideration of both the allocator's and agents' preferences.

Our Study: Centralized Allocator's Preference. Building on the work of (Bu et al. 2023), we introduce a *centralized*

*These authors contributed equally.
†Corresponding author.

Centralized allocator's valuation	Agents' valuations	EF1+CGEQ1
Arbitrary	Identical	✓ (Poly time) (Thm 1)
Ordered		✓ (Poly time) (Thm 2)
Binary	Arbitrary	✓ (Poly time) (Thm 3)

Table 1: Summary of our main results.

allocator who is interested in ensuring fairness at a group level, where each agent naturally belongs to a predefined group in the fair allocation. For instance, building on the above-mentioned example, a school administrator, tasked with allocating limited resources (e.g., office spaces and supplies) to staff members in departments within the school (Perez 2022), needs to ensure that the allocation is also fair at the department (group) level. A city council, tasked with allocating limited housing units to various neighborhoods in need across different communities (Gray 1976), needs to ensure the allocation is fair across different communities. Finally, a government distributing resources to different schools needs to ensure that the allocation is fair with respect to the schools. Therefore, in this paper, our goal is to explore the fair division of indivisible items, which provides fairness for the agents and guarantees fairness for the centralized allocator.

Our Contribution

We study the fair division of indivisible items to groups of agents from the perspectives of the agents and the centralized allocator. Each agent belongs to a group (e.g., based on their associations) and has an additive valuation function over the items. The centralized allocator has a single additive valuation function indicating their values for the items measured in standardized units (e.g., investment value, monetary amount, and space).

To ensure a fair allocation among agents, we consider the classical envy-freeness (EF) notion. To ensure fairness among the groups, we define the notion of centralized group equitability (CGEQ) to capture the fairness for the groups from the centralized allocator's perspective, which compares the weighted proportion of values received by each group. Because an EF or CGEQ allocation does not always exist—and in fact checking for a pure CGEQ outcome is computationally strongly NP-hard via a reduction from the 3-partition problem, and deciding whether an allocation satisfies EF+CGEQ, EF+CGEQ1, or EF1+CGEQ is likewise NP-hard—we therefore consider their natural relaxations of envy-freeness up to one item (EF1) and centralized group equitability up to one item (CGEQ1). Following the idea from (Bu et al. 2023), we strive to answer the following questions.

Under which conditions can we guarantee the existence of EF1+CGEQ1 allocations? If so, can we design algorithms to compute them efficiently?

To address the above questions, we examine different classes of valuation functions of the agents/allocator. The presence of the centralized allocator introduces a fundamental shift in both the fairness notions and the algorithmic challenges involved. The techniques we develop, though sometimes inspired by classic methods such as round-robin, are nontrivial extensions that integrate allocator-aware priorities and group-level proportionality. Specifically, our key contributions are as follows (summarized in Table 1):

- When each agent has an identical valuation function, even though agents are indistinguishable in terms of preferences, the allocator's separate valuation introduces nontrivial global constraints. Our DM Algorithm (Algorithm 1) constructs a temporary allocation satisfying EF1 for both agents and the allocator, using a match-based process guided by the allocator's preferences. Then, it reallocates bundles to achieve CGEQ1 while preserving EF1 at the agent level. This dynamic reassignment highlights the allocator's role in determining group-level equity even under agent homogeneity.
- When all agents and the allocator share the same ordinal ranking of items, we propose the SPS Algorithm (Algorithm 2) that simultaneously addresses item distribution across groups and within groups. Despite the aligned ordering, the absolute values may differ significantly under separate allocator's and agents' valuation functions. The allocator's proportional fairness criteria—considering each group's value-to-size ratio—guide item assignment to balance CGEQ1 and agent-level EF1. This synchronization of dual fairness notions introduces dependencies absent in standard ordinal settings.
- When the allocator classifies items into two types. Our CD^2P Algorithm (Algorithm 3) innovatively extends round-robin by introducing a reverse round-robin phase. This two-phase procedure ensures EF1 for agents while achieving CGEQ1, and cannot be reduced to a standard Round-Robin without losing group fairness.
- We also propose the Centralized Group Maximin Share (CGMMS) as an allocator-centric optimization benchmark and seek allocations that satisfy CGMMS for the centralized allocator and EF1 for the agents.

The remainder of the paper is organized as follows. In Section 2, we formally define the notations and fairness notions considered in our paper. In Sections 3, 4, and 5, we study EF1+CGEQ1 allocations in identical valuations, ordered valuations, and binary valuations settings, respectively. In Section 6, we discuss EF1+CGMMS allocations. In Section 7, we conclude the paper and provide future research directions. Due to space constraints, we refer readers to the extended version for the omitted proofs.

Related Work

There is an extensive line of work in the fair division of indivisible items. We refer readers to the survey (Amanatidis et al. 2023) for an overview. Below, we review studies focusing on allocations that consider group fairness and fairness from the agents' and allocator's perspectives.

Fairness from the Agents' and Allocator's Perspectives. As discussed earlier, the most relevant work is (Bu et al. 2023), which initiated the study of fair division of indivisible items from the perspectives of the agents and the allocator. There are significant differences between our paper and their model. More specifically,

- (Bu et al. 2023)'s "two-layer" paradigm requires a single allocation to satisfy two distinct sets of valuation functions for individuals (agents vs. allocator). Our model adds yet another layer on top of this: a group structure in which the allocator cares only about each group's aggregate share.
- When we let each group have exactly one agent, our model degenerates to their model, where the allocator has an identical valuation function. Additionally, due to the group setting in our paper, the techniques used in their paper, including envy-cycle elimination and round-robin, cannot be directly applied to our model.
- A concrete example of technology transfer failure arises with (Bu et al. 2023)'s doubly-EF1 algorithm, which hinges on cycle elimination. If one naively ports that step to the group model, moving an item that happens to be a group's extreme-value good immediately violates CGEQ1. To safeguard the group-level guarantees, we must introduce a tailored combination of bundle duplication and a quota inequality, which forms the core of our Lemma 2.

(Flammini, Greco, and Varricchio 2025) also introduced additional valuation functions in the setting, but they focused on the loss of efficiency with respect to the second valuation functions while pursuing fairness among agents, where the target is different from our paper.

Group Fairness. Existing studies have examined group-fair division of indivisible items from the agent perspective only. Some works focus on the predefined group. For example, (Aleksandrov and Walsh 2018) defined group envy-freeness and group Pareto optimality, and studied the price of group envy-freeness. (Benabbou et al. 2019) considered the fair matching among different groups where each agent can pick at most one item. They studied the fairness criteria named typewise envy-freeness up to one item (TEF1), and showed that when agents have binary valuations, TEF1 allocations can be computed in polynomial time. (Feige and Tahan 2022) studied the notion of group proportional share fairness and group any price share fairness in different groups that may have different structures like laminar. (Manurangsi and Suksompong 2024) studied the ordinal maximin share fairness among groups. There are other works that did not consider the predefined group. For example, (Conitzer et al. 2019) studied the group fairness among agents, where they considered any partition of agents, and showed that local optimal Nash welfare allocations satisfy two different relaxations of group fairness that they defined. Later, (Aziz and Rey 2021) extended it to the setting where items include both goods and chores. The survey of (Amanatidis et al. 2023) offers a comprehensive view of recent progress and open problems in this field.

The most related setting to ours is the work of (Scarlett, Teh, and Zick 2023), where they studied the compatibility of individual envy-freeness and group envy-freeness from the agent perspective only. Moreover, they did not consider the centralized allocator and defined the group utility based on the agent's valuation function instead of the centralized allocator's valuation.

2 Preliminaries

In this section, we present the notations and fairness notions for our setting of fair division of indivisible items with a set of agents and a centralized allocator.

Notations

For $r \in \mathbb{N}$, let $[r] = \{1, 2, \dots, r\}$. Let $\mathcal{O} = \{o_1, o_2, \dots, o_m\}$ be a set of m indivisible items, and $\mathcal{N} = [n]$ be a set of n agents. The set of n agents is partitioned into $k \in \mathbb{N}$ groups denoted by $\mathcal{G} = \{G_1, \dots, G_k\}$. Each agent belongs to exactly one group, i.e., $G_p \cap G_q = \varnothing$ for any $p, q \in [k]$ and $\bigcup_{p\in[k]} G_p = \mathcal{N}$. Additionally, our setting includes a *centralized* allocator.

Each agent $i \in \mathcal{N}$ has an additive valuation function $v_i : 2^{\mathcal{O}} \to \mathbb{R}_{\geq 0}$, i.e., for any $S \subseteq \mathcal{O}$, $v_i(S) = \sum_{o\in S} v_i(\{o\})$. Specifically, we assume that $v_i(\varnothing) = 0$. The centralized allocator has their own preferences and is endowed with an additive valuation function u: $2^{\mathcal{O}} \to \mathbb{R}_{\geq 0}$, i.e., for any $S \subseteq \mathcal{O}$, $u(S) = \sum_{o\in S} u(\{o\})$ indicating their values for the items measured in standardized units (e.g., investment value, monetary amount, and space). Additionally, we assume that $u(\varnothing) = 0$.

For simplicity, we use $v_i(o)$ and $u(o)$ instead of $v_i(\{o\})$ and $u(\{o\})$, respectively. Let $\Pi(n, \mathcal{O})$ denote all n-partitions of $\mathcal{O}$. An allocation $\mathcal{A} = (A_1, \dots, A_n) \in \Pi(n, \mathcal{O})$ is an n-partition of $\mathcal{O}$ among n agents, where A_i is the bundle allocated to agent i. We have $\bigcup_{i\in\mathcal{N}} A_i = \mathcal{O}$ and $A_i \cap A_j = \varnothing$ for any two agents $i \neq j$. A fair allocation instance is denoted by $\mathcal{I} = \langle \mathcal{O}, \mathcal{N}, \mathcal{G}, \boldsymbol{v}, u \rangle$, where $\boldsymbol{v} = (v_1, \dots, v_n)$. Next, we provide the formal definitions of the fairness notions for the agents and the centralized allocator.

Fairness Notions

To ensure fairness among agents, we consider the classical envy-freeness (EF) notion from the agents' perspective.

Definition 1 (Envy-Freeness)**.** *An allocation $\mathcal{A}$ is envy-free (EF) if for any two distinct agents $i, j \in \mathcal{N}$, we have $v_i(A_i) \geq v_i(A_j)$.*

However, EF allocations do not always exist. Therefore, we consider a natural and commonly studied relaxation of EF, named envy-free up to one item.

Definition 2 (Envy-Freeness up to One Item)**.** *An allocation $\mathcal{A}$ is envy-free up to one item (EF1) if, for any two agents $i, j \in \mathcal{N}$, either $A_j = \varnothing$ or $v_i(A_i) \geq v_i(A_j \setminus \{o\})$ holds for some item $o \in A_j$.*

Next, we introduce our fairness notion from the centralized allocator's perspective, which we call centralized group equitability (CGEQ).

Definition 3 (Centralized Group Equitability)**.** *An allocation $\mathcal{A}$ is centralized group equitable (CGEQ) if for any two groups $G_p, G_q \in \mathcal{G}$, $\frac{u(\bigcup_{i \in G_p} A_i)}{|G_p|} = \frac{u(\bigcup_{j \in G_q} A_j)}{|G_q|}$ holds.*

This definition reflects an *equitable view from the centralized allocator's perspective*, where the utility function $u(\cdot)$ represents the authority's valuation over bundles allocated to different groups. CGEQ thus requires that each group receives, on average, the same utility as any other group, when judged by the central authority.

Remark 1. *We intentionally label this condition "equitability" (CGEQ) instead of "envy-freeness" because the focus is on equalizing welfare rather than eliminating envy. In other words, the central planner is not checking if group G_p prefers group G_q's bundle (which would imply a notion of group envy). Rather, they are ensuring that their common yardstick values each group's allocation equally on a per-agent basis. This perspective aligns with Conitzer et al. (2019)'s group fairness structure, but we adapt it to heterogeneous valuations by following the framework of Bu et al. (2023), where the centralized allocator's valuation may differ from individual agents' valuations. By synthesizing these approaches, we obtain a group-level fairness notion judged by the central authority's values.*

The exact fairness conditions (e.g., EF) are often too strong to satisfy with indivisible items. Here, CGEQ allocations do not always exist. Consider an instance with only one indivisible item and two groups of equal size: any allocation gives utility to only one group, making CGEQ impossible. Therefore, we propose the following relaxation notion.

Definition 4 (Centralized Group Equitability up to One Item)**.** *An allocation $\mathcal{A}$ is said to be centralized group equitable up to one item (CGEQ1) if, for any two groups $G_p, G_q \in \mathcal{G}$, either $\bigcup_{j \in G_q} A_j = \varnothing$ or $\frac{u(\bigcup_{i \in G_p} A_i)}{|G_p|} \geq \frac{u(\bigcup_{j \in G_q} A_j \setminus \{o\})}{|G_q|}$ holds for some item $o \in \bigcup_{j \in G_q} A_j$.*

The relaxed version, CGEQ1 (Centralized Group Equitability up to One Item), follows the spirit of the classic EF1 relaxation (Lipton et al. 2004) but in terms of equitability—each group's average utility is almost equal, possibly differing only by the value of a single item. We use the term CGEQ1 (not CGEF1) to emphasize this distinction.

We are interested in allocations that satisfy EF1 from the agents' perspective and CGEQ1 from the centralized allocator's perspective. In particular, we study computing EF1+CGEQ1 allocations in various scenarios, focusing on representative valuation functions of the agents and the centralized allocator that have each been the subject of extensive recent study in the fair-division literature—Identical Valuation: (Barman et al. 2025; Birmpas et al. 2024); Ordered Valuation: (Neoh, Peters, and Teh 2025; Garg and Murhekar 2025); Binary Valuation: (Chandramouleeswaran, Nimbhorkar, and Rathi 2025; Bismuth, Bliznets, and Segal-Halevi 2024). If the size of any group is one, i.e., $|G_p| = 1$ for any $G_p \in \mathcal{G}$, CGEQ1 degenerates into EF1. In this case, our setting reduces to a special case of (Bu et al. 2023), where the authors showed that an EF1+EF1 allocation can be computed in polynomial time. Thus, we consider the case where at least one group has size greater than one in the following sections.

3 Identical Valuations

In this section, we consider two cases: one is that the valuations of all agents and the centralized allocator are the same, i.e., $v_1 = \cdots = v_n = u$, and the other is that only the agents' valuations are the same, i.e., $v_1 = \cdots = v_n$ and the centralized allocator's valuation is arbitrary.

For the first case, we observe that it is equivalent to the setting in (Scarlett, Teh, and Zick 2023). That work studies the combination of group and individual fairness, where the group utility is defined based on one set of valuations (i.e., the agents' valuations), and shows that such an allocation can be computed in polynomial time. Thus, it is clear that an EF1+ CGEQ1 allocation can be efficiently computed. Next, we focus on the second case. A motivating example of this case is that in textbook/course-seat allocations, agents have very similar needs for standardized items, while the centralized allocator pursues distributional goals (access, equity).

For simplicity, let v denote the valuation of each agent. In that case, the algorithm in (Scarlett, Teh, and Zick 2023) does not apply directly because its technique depends on the group utility being computed using the agents' valuations. However, the centralized allocator's valuation can be different from the agents' valuations in our setting. Therefore, we propose a new algorithm named Draft-and-Match (DM), as described in Algorithm 1. The algorithm includes two phases. In Phase 1, the intuition of our algorithm is to partition the items into a temporary allocation $\mathcal{A}'$ by allocating the item that has the highest value to the agent whose bundle has the lowest value in each iteration. This approach ensures that the allocation is envy-free up to one item (EF1) with respect to both the agents' and the centralized allocator's valuation functions. In Phase 2, given the temporary allocation, we then follow specific rules to match and reallocate these bundles to the agents, ensuring the final allocation is CGEQ1 for the centralized allocator. Since we consider identical agent valuations, this reallocation does not violate the EF1 property of each agent in $\mathcal{A}'$, and the returned allocation is EF1+CGEQ1.

Theorem 1. *Given an instance where the valuation of each agent is the same, the Draft-and-Match Algorithm (Algorithm 1) computes an EF1+CGEQ1 allocation in polynomial time.*

Before proving Theorem 1, we first present the following lemmas. In Lemma 1, we show that the temporary allocation $\mathcal{A}'$ is EF1 with respect to both the agents' and the centralized allocator's valuation functions. Then in Lemma 2, we show that after the reallocation of bundles in $\mathcal{A}'$, the returned allocation $\mathcal{A}$ is CGEQ1.

Lemma 1. *The allocation $\mathcal{A}'$ (computed in Phase 1) in Algorithm 1 is EF1 with respect to both the agents' and the centralized allocator's valuation functions. That is, for any $i, j \in [n]$, we have $v(A'_i) \geq v(A'_j \setminus \{o\})$ and $u(A'_i) \geq u(A'_j \setminus \{o'\})$ for some $o, o' \in A'_j$.*

Algorithm 1: Draft-and-Match (DM)

Input: An instance $\mathcal{I} = \langle \mathcal{O}, \mathcal{N}, \mathcal{G}, \boldsymbol{v}, u \rangle$ with identical agent valuation functions
Output: An EF1+CGEQ1 allocation $\mathcal{A}$

Let $\mathcal{A}' = (\varnothing, \dots, \varnothing)$ and $\mathcal{A} = (\varnothing, \dots, \varnothing)$;
Phase 1: Partition $\mathcal{O}$ into an allocation $\mathcal{A}'$
Add $n - (m \mod n)$ dummy items, where each agent and the centralized allocator have zero valuation, to $\mathcal{O}$;
Let $\mathcal{O}_s$ be the array of sorted goods by u in non-increasing order;
while $\mathcal{O}_s \neq \varnothing$ **do**
    Let $\mathcal{N}' = \mathcal{N}$;
    Let $\mathcal{O}_n$ be the first n items in $\mathcal{O}_s$ and $\mathcal{O}_s \leftarrow \mathcal{O}_s \setminus \mathcal{O}_n$;
    while $\mathcal{O}_n \neq \varnothing$ **do**
        $A'_{i_{min}} \leftarrow A'_{i_{min}} \cup \{o_{max}\}$, where $i_{min} \in \arg\min_{i \in \mathcal{N}'} v(A'_i)$ and $o_{max} \in \arg\max_{o \in \mathcal{O}_n} v(o)$ (breaking ties arbitrarily);
        $\mathcal{N}' = \mathcal{N}' \setminus \{i_{min}\}$ and $\mathcal{O}_n \leftarrow \mathcal{O}_n \setminus \{o_{max}\}$;
Phase 2: Match the bundles in $\mathcal{A}'$ to agents
Assume that the groups are in non-decreasing order of size, i.e., $|G_1| \leq \ldots \leq |G_k|$;
$t_p \leftarrow 0$, $\forall p \in [k]$, t_p is the total number of times that G_p has been picked so far;
while $\mathcal{A}' \neq (\varnothing, \dots, \varnothing)$ **do**
    if $\exists t_p = 0$ **then**
        $p^* \leftarrow \min\{p \mid p \in [k] \textit{ and } t_p = 0\}$;
    else
        $p^* \leftarrow \arg\min_{p \in [k]} \frac{t_p}{|G_p|}$ (breaking ties by selecting the p that reaches $\min_{p \in [k]} \frac{t_p}{|G_p|}$ the latest);
    Arbitrarily choose an agent $i^* \in G_{p^*}$ with $A_{i^*} = \varnothing$;
    $A_{i^*} \leftarrow A'_{i_{max}}$, where $A'_{i_{max}} \in \arg\max_{A'_i \in \mathcal{A}'} u(A'_i)$;
    $t_{p^*} \leftarrow t_{p^*} + 1$;
    $A'_{i_{max}} \leftarrow \varnothing$;
return $\mathcal{A}$

Since the agents' valuation functions are identical, we can rearrange the bundles without violating fairness at the individual level. Leveraging this flexibility, we aim to rearrange the bundles to satisfy the group-level fairness notion CGEQ1 while preserving EF1 for the agents.

Lemma 2. *The output allocation in Algorithm 1 is CGEQ1.*

Proof Sketch. Let $B_p^1, B_p^2, \dots, B_p^{|G_p|}$ and $B_q^1, B_q^2, \dots, B_q^{|G_q|}$ denote the bundles allocated to agents in G_p and G_q, respectively, during the reallocation process (see Figure 1 for an example). We remove the most valuable item from the centralized allocator's perspective from B_q^1, the first bundle received by G_q. After removing this item, B_q^1 becomes the least valuable among all bundles (from both G_p and G_q).

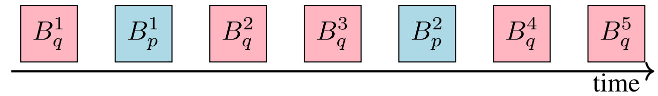

Figure 1: Example where $|G_p| = 2$ and $|G_q| = 5$, with G_q receiving the first bundle during the reallocation process. We denote B_z^i as the bundle received by G_z in the i-th allocation. The red squares represent bundles received by G_q, and the blue squares represent bundles received by G_p.

Next, we duplicate each bundle received by G_p $|G_q|$ times and each bundle received by G_q $|G_p|$ times, ordering all duplicated bundles in non-decreasing order of value. In this sequence, $|G_p|$ copies of B_q^1 (with one item removed) now become the least valuable. This duplication transforms the problem of comparing the average values of the bundles into comparing the sum of the values of all duplicated bundles.

We then perform a pairwise comparison between the duplicated bundles from both groups. In this comparison, each duplicated bundle from G_p is at least as valuable as the corresponding duplicated bundle from G_q. Therefore, the allocation satisfies CGEQ1. □

Proof of Theorem 1. By Lemmas 1 and 2, and the fact that each agent has the same valuation, we conclude that the final allocation is EF1+CGEQ1. Next, we analyze the time complexity. Without loss of generality, we assume that $m \geq n$. In the first part of our algorithm, it takes $O(m \log m)$ to sort the items. In each iteration, selecting an agent takes $O(n)$ time, and choosing their favorite item takes $O(m)$ time. Note that there are $O(\lceil \frac{m}{n} \rceil)$ iterations. In the second part of our algorithm, there are $O(n)$ iterations in the while loop. For each iteration, selecting the target group and agent takes $O(nk)$ time, and choosing the bundle with the highest value from the centralized allocator's perspective takes $O(n)$ time. Therefore, the total running time of our algorithm is $O(m^2 + m \log m + n^2 k)$. □

4 Ordered Valuations

In this section, we consider the instance $\mathcal{I}$ with ordered valuations, where each agent $i \in \mathcal{N}$ and the centralized allocator share the same ranking or preference for all items. Specifically, $v_i(o_1) \geq \dots \geq v_i(o_m)$ and $u(o_1) \geq \dots \geq u(o_m)$. In university settings, students typically agree on the relative order of course preferences (e.g., core before electives), but differ in the intensity of those preferences due to individual constraints such as major, prerequisites, or scheduling. EF1 remains a realistic fairness notion under such variations, while centralized rules determine access priorities. The challenge in this setting arises when certain items are valued differently by the agents and the centralized allocator. In such cases, ordered valuations may help us circumvent this issue.

Inspired by the Weighted Picking Sequence Protocol of (Chakraborty et al. 2021), we propose the Synchronous

Algorithm 2: Synchronous Picking Sequence (SPS)

Input: An instance $\mathcal{I} = \langle \mathcal{O}, \mathcal{N}, \mathcal{G}, \boldsymbol{v}, u \rangle$ with ordered valuation functions
Output: An EF1+CGEQ1 allocation $\mathcal{A}$

Let $\mathcal{A} = (\varnothing, \dots, \varnothing)$;
Add $n - (m \mod n)$ dummy items with zero value for the agents and the centralized allocator to $\mathcal{O}$;
Assume that the groups are ordered in non-decreasing order of size, i.e., $|G_1| \leq \dots \leq |G_k|$;
Set $t_p \leftarrow 0, \forall p \in [k]$ and $\ell_i \leftarrow 0, \forall i \in \mathcal{N}$;
while $\mathcal{O} \neq \varnothing$ **do**
  Let $\mathcal{O}_n$ be the first n items in $\mathcal{O}$ and $\mathcal{O} \leftarrow \mathcal{O} \setminus \mathcal{O}_n$;
  while $\mathcal{O}_n \neq \varnothing$ **do**
    Let $o_{max} \in \arg\max_{o \in \mathcal{O}_n} u(o)$;
    Phase 1: Decide which group gets this item
    if $\exists t_p = 0$ **then**
      $p^* \leftarrow \min\{p \mid p \in [k] \text{ and } t_p = 0\}$;
    else
      $p^* \leftarrow \arg\min_{p \in [k]} \frac{t_p}{|G_p|}$ (breaking ties by selecting p that reaches $\min_{p \in [k]} \frac{t_p}{|G_p|}$ the latest);
    Phase 2: Decide which agent picks this item
    if $\exists i^* \in G_{p^*}$ *such that* $A_{i^*} = \varnothing$ **then**
      $A_{i^*} \leftarrow \{o_{max}\}$ and $\ell_{i^*} \leftarrow t_{p^*}$;
    else
      Find the agent i^* whose label ℓ_{i^*} equals $t_{p^*} \mod |G_p|$, and $A_{i^*} \leftarrow A_{i^*} \cup \{o_{max}\}$;
    $t_{p^*} \leftarrow t_{p^*} + 1$ and $\mathcal{O}_n \leftarrow \mathcal{O}_n \setminus \{o_{max}\}$;
return $\mathcal{A}$

Picking Sequence Algorithm, detailed in Algorithm 2. Our algorithm operates by allocating a set of items in several batches, where each batch has n items. If fewer than n items remain, we add dummy items that have a zero value for the agents and the centralized allocator. For each item within a batch, there are two phases of allocation. In the first phase, the algorithm assigns the item to a group. In the second phase, the item is allocated to a specific agent within that group. Every agent receives exactly one item per batch. By structuring the allocation in this way, the algorithm mirrors a specialized round-robin algorithm. Figure 2 illustrates the allocation process of the first batch (first twelve items) for an instance where there are three groups with 3, 4, and 5 agents, respectively, using Algorithm 2.

Theorem 2. *Given an instance with ordered valuations, the Synchronous Picking Sequence algorithm (Algorithm 2) computes an EF1+CGEQ1 allocation in polynomial time.*

5 Binary Centralized Allocator Valuations

In this section, we consider the instance $\mathcal{I}$ where the centralized allocator has a binary valuation function, i.e., for each

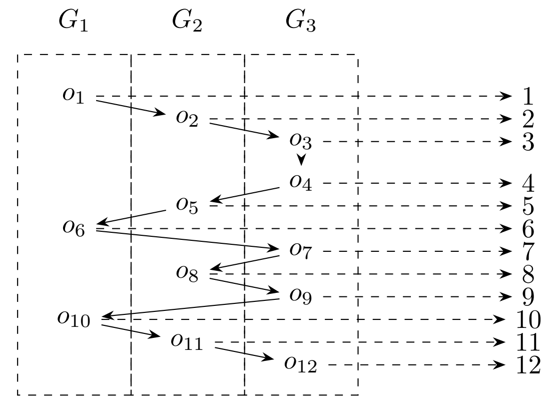

Figure 2: An illustration of the allocation process of the first twelve items. We use ℓ_i to denote the picking order of agent i in their group. Assume that there are three groups G_1, G_2, and G_3, where the groups have 3, 4, and 5 agents, respectively. The arrows represent the allocation sequence of these twelve items. For example, in G_1, agent 1 with $\ell_1 = 1$ picks o_1, agent 6 with $\ell_6 = 2$ picks o_6, and agent 10 with $\ell_{10} = 3$ picks o_{10}. Then, in the following iterations, if G_1 receives an item, the algorithm will follow the order to select the target agent.

item o, either $u(o) = 0$ or $u(o) = 1$ holds. This setting is motivated by by vaccine or emergency-supply alloction, where agencies first distribute critical items (valued as 1 by the allocator) before non-critical ones (valued as 0). In this sense, our $\{0, 1\}$ centralized-allocator valuation model and the rule "allocate $\mathcal{O}^1$ before $\mathcal{O}^0$" mirror public-health prioritization under scarcity (Faden et al. 2020). In this setting, we show that an EF1+CGEQ1 allocation always exists, and it can be computed by the Clustering-Based Dual-Flow Picking Algorithm (CD^{2}P, Algorithm 3) in polynomial time.

Before the execution of our main algorithm, we first cluster the set of items $\mathcal{O}$ into two groups: $\mathcal{O}^1$, where $u(o) = 1$, and $\mathcal{O}^0$, where $u(o) = 0$. The idea behind Algorithm 3 is as follows: First, we determine an order over the agents so that when we allocate items in $\mathcal{O}^1$ using a round-robin approach, it satisfies the CGEQ1 property. Once the allocation of $\mathcal{O}^1$ is completed, we proceed to allocate items in $\mathcal{O}^0$. Note that the allocation of items in $\mathcal{O}^0$ does not affect the CGEQ1 property. For this allocation, we use the reverse of the previously determined sequence and perform another round-robin distribution based on the agents' valuation functions.

Theorem 3. *Given an instance where the centralized allocator has the binary valuation function, the Clustering-Based Dual-Flow Picking Algorithm (Algorithm 3) computes an EF1+CGEQ1 allocation in polynomial time.*

6 Centralized Group Maximin Share

In the above sections, the centralized allocator enforces group-level fairness among agents through an additional fairness requirement (CGEQ1). Now, we shift our focus to optimizing group-level fairness objectives directly, aiming to achieve fairness from a centralized perspective while still

Algorithm 3: Clustering-Based Dual-Flow Picking (CD^{2}P)

Input: An instance $\mathcal{I} = \langle \mathcal{O}, \mathcal{N}, \mathcal{G}, \boldsymbol{v}, u \rangle$ with binary centralized allocator valuation function
Output: An EF1+CGEQ1 allocation $\mathcal{A}$

Let $\mathcal{A} = (\varnothing, \dots, \varnothing)$;
Denote the collection of items with $u(o) = 1$ as $\mathcal{O}^1$ and the collection of items with $u(o) = 0$ as $\mathcal{O}^0$;
Set $t_p \leftarrow 0, \forall p \in [k]$, $\ell_i \leftarrow 0, \forall i \in \mathcal{N}$, and $t \leftarrow 0$;
Phase 1: Allocate items in $\mathcal{O}^1$
while $\mathcal{O}^1 \neq \varnothing$ **do**
    if $\exists i \in \mathcal{N}$ *such that* $A_i = \varnothing$ **then**
        if $\exists t_p = 0$ **then**
            $p^* \leftarrow \min\{p \mid p \in [k] \text{ and } t_p = 0\}$;
        else
            $p^* \leftarrow \arg\min_{p \in [k]} \frac{t_p}{|G_p|}$ (breaking ties by selecting p that reaches $\min_{p \in [k]} \frac{t_p}{|G_p|}$ the latest);
        Find agent $i^* \in G_{p^*}$ with $A_{i^*} = \varnothing$;
        $A_{i^*} \leftarrow \{o_{max}\}$ where $o_{max} \in \arg\max_{o \in \mathcal{O}^1} v_{i^*}(o)$;
        $\ell_{i^*} \leftarrow t$;
        $t_{p^*} \leftarrow t_{p^*} + 1$;
    else
        Find agent i^* whose label ℓ_{i^*} equals t
        $A_{i^*} \leftarrow A_{i^*} \cup \{o_{max}\}$ where $o_{max} \in \arg\max_{o \in \mathcal{O}^1} v_{i^*}(o)$;
    $t \leftarrow (t+1) \mod n$ and $\mathcal{O}^1 \leftarrow \mathcal{O}^1 \setminus \{o_{max}\}$;
$t \leftarrow n - 1$;
Phase 2: Allocate items in $\mathcal{O}^0$
while $\mathcal{O}^0 \neq \varnothing$ **do**
    Find agent i^* with ℓ_{i^*} equals t;
    $A_{i^*} \leftarrow A_{i^*} \cup \{o_{max}\}$ where $o_{max} \in \arg\max_{o \in \mathcal{O}^0} v_{i^*}(o)$;
    $t \leftarrow (t-1) \mod n$ and $\mathcal{O}^0 \leftarrow \mathcal{O}^0 \setminus \{o_{max}\}$;
return $\mathcal{A}$

maintaining EF1 for the agents. This can be understood as the centralized allocator striving to find the "best" fair allocation.

In this case, the utilitarian social welfare ($\sum_{i=1}^{n} u(A_i)$) is not suitable as the optimization objective since it not only remains invariant regardless of the allocation computed but also fails to reflect group-level fairness. Instead, we focus on the share-based fairness objective from the centralized allocator's perspective, which is called centralized group maximin share (CGMMS). This definition is motivated by the well-studied notion—maximin share fairness (MMS) (Budish 2011). Our main goal is to find an allocation satisfying CGMMS from the centralized allocator's perspective and EF1 from the agents' perspective simultaneously.

Definition 5 (Centralized Group Maximin Share)**.** *Let $\mathcal{O}$ be the set of items and $\Pi_n(\mathcal{O})$ be the set of n-partitions of $\mathcal{O}$ (which may be subject to some constraints). The centralized group maximin share* CGMMS *is defined as:*

$$\mathsf{CGMMS} = \max_{\mathcal{A} \in \Pi_n(\mathcal{O})} \min_{G_p \in \mathcal{G}} \frac{u(\cup_{i \in G_p} A_i)}{|G_p|}.$$

An allocation $\mathcal{A}$ is centralized group maximin share fair (CGMMS) if it holds that $\min_{G_p \in \mathcal{G}} \frac{u(\cup_{i \in G_p} A_i)}{|G_p|} = \mathsf{CGMMS}$.

Regarding the existence of an EF1+CGMMS allocation, we have the following example showing that it may not exist.

Example 1. *Consider an instance where there are four items $\{e_1, e_2, e_3, e_4\}$ and two groups $\{G_1, G_2\}$, where G_1 has agent 1 and G_2 has agents 2 and 3. Let $\epsilon > 0$ be a sufficiently small number. For the centralized allocator, we have $u(e_1) = 1$ and $u(e_i) = \epsilon$ for any $i \neq 1$. For agents, we assume that they have the same valuation function v, where $v(e_1) = \epsilon$ and $v(e_i) = 1$ for any $i \neq 1$. It is easy to verify that* $\mathsf{CGMMS} = 3\epsilon$*, and that the unique CGMMS allocation is the following: e_1 is allocated to G_2 and the remaining items are allocated to G_1. In that case, the EF1 property cannot hold for them since agent 1 has three items whose value is 3, whilst there exists one agent in G_2 such that the value of her bundle is zero.*

We next show that computing such an allocation is NP-hard.

Proposition 1. *Computing a CGMMS allocation subject to EF1 for agents is strongly NP-hard.*

Although there is a strong negative result for the computational complexity, we find that when the centralized allocator has a binary valuation function, an EF1+CGMMS allocation can be computed efficiently.

Theorem 4. *When the centralized allocator has a binary valuation function, computing a CGMMS allocation subject to EF1 for agents can be achieved in polynomial time.*

7 Conclusion

In this paper, we study the fair division of indivisible items from the perspectives of agents and a centralized allocator. We propose using EF1 and CGEQ1 to measure the fairness from the agents' and the centralized allocator's perspectives, respectively, and aim to compute allocations that satisfy EF1 and CGEQ1 simultaneously. We show that EF1+CGEQ1 allocations always exist for different classes of agents' and the centralized allocator's valuation functions, and that such allocations can be computed in polynomial time. Regarding optimizing group-level fairness objectives, we show that, in general, finding a CGMMS allocation is hard, but an EF1+CGMMS allocation can be computed in polynomial time when the centralized allocator has a binary valuation function.

For future work, one natural direction is to determine whether an allocation satisfying the above two fairness notions exists in more general settings. *We have searched for a counterexample to existence with the aid of computer programs, but it seems to be hard to find such an instance.* Another interesting direction is to explore whether we can efficiently obtain an approximate CGMMS+EF1 allocation.

Acknowledgments

Prof. Hau Chan is supported by the National Institute of General Medical Sciences of the National Institutes of Health [P20GM130461], the Rural Drug Addiction Research Center at the University of Nebraska-Lincoln, and the National Science Foundation under grants IIS:RI #2302999 and IIS:RI #2414554. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies. Prof. Minming Li is supported by a grant from Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. CityU 11216725)

References

Aleksandrov, M.; and Walsh, T. 2018. Group envy freeness and group pareto efficiency in fair division with indivisible items. In *Proceedings of 41st German Conference on AI*, 57–72.

Amanatidis, G.; Aziz, H.; Birmpas, G.; Filos-Ratsikas, A.; Li, B.; Moulin, H.; Voudouris, A. A.; and Wu, X. 2023. Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.*, 322: 103965.

Aziz, H.; and Rey, S. 2021. Almost group envy-free allocation of indivisible goods and chores. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, 39–45.

Barman, S.; Ebadian, S.; Latifian, M.; and Shah, N. 2025. Fair Division with Market Values. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 13589–13596.

Benabbou, N.; Chakraborty, M.; Elkind, E.; and Zick, Y. 2019. Fairness towards groups of agents in the allocation of indivisible items. In *The 28th International Joint Conference on Artificial Intelligence*, 95–101.

Birmpas, G.; Ezra, T.; Leonardi, S.; and Russo, M. 2024. Fair division with interdependent values. In *International Symposium on Algorithmic Game Theory*, 72–88. Springer.

Bismuth, S.; Bliznets, I.; and Segal-Halevi, E. 2024. Fair division with bounded sharing: Binary and non-degenerate valuations. In *International Symposium on Algorithmic Game Theory*, 89–107. Springer.

Bu, X.; Li, Z.; Liu, S.; Song, J.; and Tao, B. 2023. Fair division with allocator's preference. In *The 19th Conference on Web and Internet Economics*, 77–94.

Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6): 1061–1103.

Budish, E.; Cachon, G. P.; Kessler, J. B.; and Othman, A. 2017. Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research*, 65(2): 314–336.

Chakraborty, M.; Igarashi, A.; Suksompong, W.; and Zick, Y. 2021. Weighted envy-freeness in indivisible item allocation. *ACM Transactions on Economics and Computation*, 9(3): 1–39.

Chandramouleeswaran, H.; Nimbhorkar, P.; and Rathi, N. 2025. Fair Division in a Variable Setting. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*, 472–480.

Conitzer, V.; Freeman, R.; Shah, N.; and Vaughan, J. W. 2019. Group fairness for the allocation of indivisible goods. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, 1853–1860.

Faden, R.; Saad, O.; Kochhar, S.; Kaslow, D.; Pallas, S.; Olayinka, F.; Afolabi, M.; Smith, P.; and Wilder-Smith, A. 2020. WHO SAGE values framework for the allocation and prioritization of COVID-19 vaccination.

Feige, U.; and Tahan, Y. 2022. On allocations that give intersecting groups their fair share. arXiv:2204.06820.

Flammini, M.; Greco, G.; and Varricchio, G. 2025. Fair Division with Social Impact. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 13856–13863.

Garg, J.; and Murhekar, A. 2025. Existence of 2-EFX Allocations of Chores. *arXiv preprint arXiv:2507.19461*.

Goldman, J.; and Procaccia, A. D. 2015. Spliddit: Unleashing fair division algorithms. *ACM SIGecom Exchanges*, 13(2): 41–46.

Gray, F. 1976. Selection and Allocation in Council Housing. *Transactions of the Institute of British Geographers*, 1(1): 34–46.

Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, 125–131.

Manurangsi, P.; and Suksompong, W. 2024. Ordinal Maximin Guarantees for Group Fair Division. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 2922–2930.

Moulin, H. 2004. *Fair division and collective welfare*.

Neoh, T. Y.; Peters, J.; and Teh, N. 2025. Online Fair Division with Additional Information. *arXiv preprint arXiv:2505.24503*.

Perez, A. 2022. The Role of a School Principal's Practice of Resource Allocation and its Relationship with a School's Culture of Collaboration. *Journal of Educational Leadership in Action*, 8(1): 1.

Scarlett, J.; Teh, N.; and Zick, Y. 2023. For One and All: Individual and Group Fairness in the Allocation of Indivisible Goods. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, 2466–2468.