

Average Envy-freeness for Indivisible Items

QISHEN HAN, Rensselaer Polytechnic Institute, USA

BIAOSHUAI TAO, Shanghai Jiao Tong University, China

LIRONG XIA, Rensselaer Polytechnic Institute, USA

In fair division applications, agents may have unequal entitlements reflecting their different contributions. Moreover, the contributions of agents may depend on the allocation itself. Previous fairness notions designed for agents with equal or pre-determined entitlements fail to characterize fairness in these collaborative allocation scenarios.

We propose a novel fairness notion of *average envy-freeness* (AEF), where the envy of agents is defined on the average value of items in the bundles. Average envy-freeness provides a reasonable comparison between agents based on the items they receive and reflects their entitlements. We study the complexity of finding AEF and its relaxation, *average envy-freeness up to one item* (AEF-1). While deciding if an AEF allocation exists is NP-complete, an AEF-1 allocation is guaranteed to exist and can be computed in polynomial time. We also study allocation with *quotas*, i.e. restrictions on the sizes of the bundles. We prove that finding an AEF-1 allocation satisfying quotas is NP-hard. Nevertheless, in the instances with a fixed number of agents, we propose polynomial-time algorithms to find an AEF-1 allocation with quotas for binary valuation and an approximated AEF-1 allocation with quotas for general valuation.

CCS Concepts: • Computing methodologies → Multi-agent systems; • Theory of computation → Algorithmic mechanism design.

Additional Key Words and Phrases: fair division, envy-freeness, resource allocation

ACM Reference Format:

Qishen Han, Biaoshuai Tao, and Lirong Xia. 2023. Average Envy-freeness for Indivisible Items. In *EAAMO '23: Average Envy-freeness for Indivisible Items*, October 30 – November 1, 2023, Boston, MA. ACM, New York, NY, USA, 15 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Fair division aims to allocate items to a group of agents with different preferences and achieve fairness among the agents. It is a classical yet heating topic that has wide application in real-world scenarios including peer review [25], cloud computing [28], and healthcare resource distribution [26]. In fair division applications, it is often the case that agents are *asymmetric*, i.e. having unequal entitlements. For example, in a food bank allocation, a family with more population requires more food to feed themselves. Asymmetric agents characterize a wide range of scenarios where agents have different contributions in a collaboration or agents represent groups with different populations.

Extending fairness notions to agents with different entitlements, Chakraborty et al. [14] proposes *weighted envy-freeness* (WEF), where the entitlements of agents are characterized by *predetermined weights*, and envy is defined on the weighted margin. Its relaxation, *weighted envy-freeness up to one item* (WEF-1), is guaranteed to exist and has been

The full version of this paper, including an appendix with proofs, can be found at <https://arxiv.org/abs/2301.12653>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

Manuscript submitted to ACM

applied to paper-reviewer matching mechanisms [25]. However, agents' entitlement and contribution depend on the allocation in many scenarios. This happens in collaborations on productive activities whose outputs depend on the resource allocated, and in allocating human resources, as shown in the following example.

EXAMPLE 1. *Research resource allocation.* *Multiple research labs decide to start an interdisciplinary research collaboration. They need to allocate the research resources (e.g. funding, computing resources, assistants) and assign tasks to each lab. A lab's contribution (task) depends on the resource it gets, so both plans should be determined simultaneously. Different labs have expertise in different areas and have different preferences for resources. A computer science group would prefer more computing resources while a biological group would prefer more research assistants. Furthermore, the labs want the resources allocated fairly based on their expected contributions. How should they fairly allocate the resources?*

Human resource allocation. *Several team leaders in a company are allocating new employees. Employees can make different levels of contribution to different teams based on the compatibility of their skills to the team's task. Team leaders earn credits based on their total contribution and the employees they occupy. A large team occupies more human resources and is expected to make a larger contribution. Particularly, the value of a team leader is the average contribution of his/her team members. The department wishes to work out an allocation plan so that no team leader would perceive another team with a higher value than his/her own team. How should the department fairly allocate the employees?*¹

There are two challenges in characterizing fairness in such scenarios. Firstly, agents (team leaders and labs) have unequal entitlements, as an agent with more resources ought to have a larger contribution. Secondly, the allocation and the entitlements are decided simultaneously, and the entitlements of the agents depend on the allocation. More specifically, in both scenarios, the value of a bundle of resources is the average value of all its items (research/human resources). Under such a challenge, neither EF nor WEF is able to characterize fairness. EF designed for equal entitlement cannot directly extend to different entitlements. WEF requires a *predetermined* weight for each agent. When the entitlements change with the allocation, a fixed weight is unable to represent them.

Based on the challenges, the following question remains unanswered: **what is a proper fairness notion for fair division with variable entitlements?**

1.1 Average Envy-freeness

We propose the notion of *average envy-freeness* (AEF), where envy between agents is defined on the average value of the items in the bundles.

DEFINITION 1 (AVERAGE ENVY-FREENESS (AEF)). *An allocation A is said to be average envy-free if for any pair of agents $i, h \in N$, $\frac{v_i(A_i)}{|A_i|} \geq \frac{v_i(A_h)}{|A_h|}$.*

In average envy-freeness, the entitlement of each agent is the number of items they received. For example, in the research resource allocation scenario, the contribution of a lab will depend on the research resource they receive. The labs with more resources are expected to make a larger contribution and deserve larger entitlements.

EXAMPLE 2. *Suppose in the human resource allocation problem, there are five employees to be assigned to two teams. The level of contribution of each agent to each team is shown in Table 1, as well as a possible allocation. For Team 1, the leader's utility will be $(9 + 6 + 6)/3 = 7$, and his/her perceived value for Team 2 is $(4 + 2)/2 = 3$. Therefore, the leader of Team 1 will not envy the leader of Team 2. Similarly, the leader of Team 2 will not envy the leader of Team 1. Therefore, the allocation is average envy-free.*

¹We thank an anonymous reviewer for providing this example.

Employees	1	2	3	4	5
Team 1	(9)	(6)	(6)	4	2
Team 2	4	2	4	(8)	(4)

Table 1. The levels of contribution of employees to each team. The circles represent the allocation.

The major difference between WEF and AEF is the modeling of the entitlements. In WEF, the entitlement of each agent is their weight, which is pre-determined and independent of the allocation. In AEF, on the other hand, the entitlement or contribution of an agent is reflected by the size of their bundle. In a scenario where the allocation and the entitlements are decided simultaneously, predetermined weights cannot characterize entitlements that depend on the allocation. On the other hand, when there is no large difference between entitlements brought by different items/individuals, the size of each bundle serves as a natural estimation of the entitlements of each agent.

We believe that average envy-freeness is a proper fairness notion for variable entitlement scenarios, especially for the human resources allocation in Example 1, because it provides a reasonable comparison between the value of teams based on their heterogeneous entitlements.

1.2 Average envy-freeness with a quota

We study the existence and the complexity of average envy-freeness and average envy-freeness up to one item (AEF-1). Deciding whether an AEF allocation exists is NP-complete (Theorem 1), while an AEF-1 allocation always exists and can be computed in polynomial time (Proposition 1). Unfortunately, AEF and AEF-1 allocations sometimes fail to characterize real-world scenarios. AEF and AEF-1 allow allocations where all but one agent receives exactly one item and the rest agent gets all other items, while there are many scenarios with restrictions on the sizes of bundles.

EXAMPLE 3. We follow the employee allocation scenario in Example 2. Note that an allocation where Team 2 gets Employee 4 and Team 1 gets all other employees is also an AEF allocation. The value of Team 1 for leader 1 is $(9 + 6 + 6 + 2)/4 = 5.75$, while the value of Team 2 is 4. Therefore, leader 1 does not envy leader 2. Similarly, leader 2 does not envy leader 1 as well.

However, this allocation may not work in real-world human resource allocations, because the sizes of the teams are usually restricted by the tasks. For example, the task for Team 2 requires at least two employees (leader excluded) so that nobody is overburdened. How can the allocations take these requirements into consideration while achieving fairness on average?

To model the restrictions in these scenarios, we study fair division scenarios with *quotas*, i.e. restrictions on the size of the bundles. A quota reflects the requirements and capability of each agent in the allocation and occurs in many real-world applications. For example, each team has a minimum and maximum requirement of employees. The involvement of quotas allows AEF to be applied to a wider range of scenarios.

1.3 Our contribution

The summary of our result is shown in Table 2. We first consider the existence of AEF and AEF-1 without quotas. Deciding whether an AEF allocation exists is NP-complete, while an AEF-1 allocation always exists and can be computed in polynomial time. The problem becomes much more difficult when quotas are considered, as deciding whether there exists an AEF-1 allocation satisfying the quota is NP-complete. Therefore, we consider the instances with a fixed number

of agents. When the value of each item is either zero or one, we propose a polynomial-time dynamic programming algorithm to decide the existence of an AEF-1 allocation satisfying the quota and find the allocation if it exists. On the other hand, deciding the existence of an AEF-1 allocation satisfying the quota in general valuation is still NP-complete. For the general valuation case, we give an approximation algorithm that finds an $(1 - \frac{4}{mn})$ -AEF-1 allocation, in which agents value their bundles at least $(1 - \frac{4}{mn})$ times of other agent's bundles on the average value after removing one item.

Criterion	n	Valuation	Result	Location
AEF	2	Identical	NP-complete	
	General	General		Theorem 1
AEF-1	General	General	Always Exists	Proposition 1
	General	Binary		
		General	NP-complete	Theorem 2
AEF-1	General	Binary	In P	Theorem 3
with quota	Constant, ≥ 3	General	NP-complete	Theorem 4
			$(1 - \frac{4}{mn})$ -approximation	Theorem 5

Table 2. Result Summary of AEF and AEF-1. n and m are the number of agents and the number of items respectively.

2 RELATED WORKS

There is a large literature on the fair division problem with asymmetric agents, i.e. individuals or groups with heterogeneous entitlements. Chakraborty et al. [14] proposes the notion of *weighted envy-freeness up to one item* (WEF-1), where predetermined weights of each agent represent their entitlements, and shows that a WEF-1 allocation always exists. Payan and Zick [25] and Chakraborty et al. [15] focus on maximizing social welfare of WEF-1 allocations by select the correct picking sequence in a Round-Robin-like mechanism. Chakraborty et al. [16] proposes a parameterized family of weighted envy-freeness, where agents with large or small weights are favored by setting different parameters. Wu et al. [29] studies WEF-1 allocation for indivisible chores. Other fairness notions on asymmetric agents include weighted MMS [5, 6, 18], weighted proportionality [5, 24], and maximizing weighted Nash social welfare [27]. Nevertheless, previous works focuses on fixed entitlements, and does not address scenarios when the entitlements depend on the allocation.

Another perception of average envy-freeness is exactly the envy-freeness under the setting where the valuation functions of agents (rather than the envies) are average valuation, i.e. the value of a subset is the average value of the items. This is related to fair division with non-additive and non-monotone valuation. In a non-additive valuation, the value of a set may not be equal to the sum of the items in the set. In a non-monotone valuation, the value of a set may be less than its subsets. Aziz et al. [1] study envy-freeness with a mixture of goods, chores, and mixed items. Bhaskar et al. [9] shows the existence of an EF-1 allocation under the so-called doubly monotonic valuation. Bérczi et al. [13] proposes an algorithm to find EF-1 allocation between to agents with arbitrary valuation. Other works include envy-freeness with stochastic non-additive valuation [8] and MMS with non-additive valuation [21, 22]. In general, fair division with non-additive and non-monotone valuation is a challenging topic, and so far we haven't observed existing work addressing the average valuation scenario.

Another related line of work is fair division with *cardinality constraints*. Biswas and Barman [10] considers a scenario where items are categorized into multiple groups, and each group has a capacity to contribute to each agent. They show that an EF-1 allocation and a constant-factor approximation of MMS allocation always exist in such a scenario. They also extend this type of constraint to be represented by a matroid. Biswas and Barman [11] designs a polynomial time algorithm to find an EF-1 allocation satisfying the matroid constraint. Dror et al. [17] studies scenarios where agents have heterogeneous matroid constraints, and provides an algorithm to find EF-1 allocations in certain circumstances. Gan et al. [19] and Babaioff et al. [7] study fair division problem where agents have budget constraints on the items. Aziz et al. [3] propose a mechanism that turns a welfare-efficient allocation into a fair allocation while preserving the efficiency constraint.

3 PRELIMINARIES

Problem Instance. An instance of a fair division problem $\mathcal{I} = \langle N, M, V \rangle$ is defined by a set of n agents $N = [n]$, a set of m items M , and a valuation profile $V = \{v_1, v_2, \dots, v_n\}$. We use i to denote a generic agent in N and g to denote a generic item in M .

Valuation and Average Value. Valuation functions represent the preferences of agents among the items. For each agent i , v_i is a mapping from a subset of M to a non-negative value. $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$. We follow the convention to assume *additive valuation*, i.e for any $i \in N$ and $M' \subseteq M$, $v_i(M') = \sum_{g \in M'} v_i(g)$. We say that an instance has *binary valuation* if for any $i \in N$ and $g \in M$, $v_i(g) \in \{0, 1\}$, and an instance has *identical valuation* if $v_1 = v_2 = \dots = v_n$. We also define the *average value* of a subset of item M' as $u_i(M') = \frac{v_i(M')}{|M'|}$, i.e. the average value of the items in the subset. Specifically, $u_i(\emptyset) = 0$.

Allocation. An allocation $A = (A_1, A_2, \dots, A_n)$ is a n -partition of the set of items M , where $A_i \subseteq M$ is the *bundle* allocated to agent i . We sometimes abuse the notation and use A to denote a partial allocation, where there are items unallocated to any agent.

Quota. A quota Q is a constraint on allocations. For each agent i , Q imposes an upper bound and a lower bound to the size of the bundle A_i . An allocation A satisfies a quota Q if all agents satisfy the constraint. A quota is said to be *exact* if the upper bound equals to the lower bound for every agent. An exact quota regulates the exact number of items in each bundle.

DEFINITION 1 (AVERAGE ENVY-FREENESS (AEF)). *An allocation A is said to be average envy-free if for any pair of agents $i, h \in N$, $u_i(A_i) \geq u_i(A_h)$.*

DEFINITION 2 (AVERAGE ENVY-FREENESS UP TO ONE ITEM (AEF-1)). *An allocation A is said to be average envy-free up to one item if for any pair of agents $i, h \in N$, there exists an item $g \in A_i \cup A_h$ such that $u_i(A_i \setminus \{g\}) \geq u_i(A_h \setminus \{g\})$.*

By the definition of AEF-1, agents can remove an item from either bundle under comparison. This is because agents can increase the average value of their own bundle by removing the least preferred item (if it's not the only item). It follows from the definition that an AEF allocation is always AEF-1. An example of an AEF allocation is the allocation in Example 1.

We introduce the computational problems related to AEF and AEF-1.

DEFINITION 3 (AEF-EXISTENCE). *Given an instance \mathcal{I} , does there exist an allocation A such that A is an AEF allocation?*

DEFINITION 4 (AEF-1-EXISTENCE). *Given an instance \mathcal{I} , does there exist an allocation A such that A is an AEF-1 allocation?*

DEFINITION 5 (AEF-EXISTENCE WITH QUOTAS). *Given an instance \mathcal{I} and a quota Q , does there exist an allocation A such that A is an AEF allocation and satisfies Q ?*

DEFINITION 6 (AEF-1-EXISTENCE WITH QUOTAS). *Given an instance \mathcal{I} and a quota Q , does there exist an allocation A such that A is an AEF-1 allocation and satisfies Q ?*

4 AVERAGE ENVY-FREENESS WITHOUT QUOTAS

This section focuses on finding AEF and AEF-1 allocations without quota constraints. We show that deciding the existence of an AEF allocation is NP-complete, while an AEF-1 allocation always exists and can be found in polynomial time.

THEOREM 1. *AEF-EXISTENCE is NP-complete even for two agents with identical valuations.*

PROOF SKETCH. We construct a reduction from PARTITION, which is known to be NP-complete [20]. An instance of PARTITION consists of a multiset $X = \{x_1, x_2, \dots, x_k\}$ where $x_i \in \mathbb{N}$. The goal is to determine whether X can be partitioned into two subsets Y and $X \setminus Y$ with equal sum T .

Given an instance of PARTITION X , we construct an AEF-EXISTENCE instance with two agents and $m = 2k$ items. Two agents share the same value function v , and u is the average valuation function. For each x_i , there exists two items g_i^s and g_i^l such that $v(g_i^s) = (T^2 k^2)^i$, and $v(g_i^l) = (T^2 k^2)^i + x_i$.

(\Rightarrow) Given Y and $X \setminus Y$ being an equal-sum partition of X , we construct an allocation A . For each i , A_1 gets g_i^l if $x_i \in Y$ or g_i^s if $x_i \in X \setminus Y$. A_2 gets the rest of the items. It is not hard to verify that both agents get k items, and $v(A_1) = v(A_2)$. Therefore, $u(A_1) = u(A_2)$, and A is an AEF allocation.

(\Leftarrow) If there exists an AEF allocation A , we show that A induces an equal-sum partition of X in four steps.

First, each agent gets exactly one of the largest items g_k^s and g_k^l . Otherwise, the exponential term $(T^2 k^2)^i$ guarantees that the agent without the two items envies the other agent.

Second, each agent gets exactly k items. g_k^s and g_k^l guarantee that the additive value of bundles is at the same level, and the agent with more items envies the agent with fewer items.

Third, for each $i = 1, 2, \dots, k$, each agent gets exactly one of g_i^s and g_i^l , following the similar reasoning of the first step.

Finally, the allocation induces an equal-sum partition of X . Let $Y = \{x_i | g_i^l \in A_1\}$, and Y and $X \setminus Y$ is an equal-sum partition of X . Given that each agent gets exactly one of g_i^s and g_i^l , the difference between two bundles just come from x_i from each i . Therefore, the sum of x_i in A_1 must equal to the sum of x_i in A_2 , which implies Y and $X \setminus Y$ be an equal-sum partition of X . The full proof is available in the full version of this paper. \square

The fact that AEF-EXISTENCE is already NP-complete without quotas directly implies that AEF-EXISTENCE with quotas is NP-complete.

COROLLARY 1. *AEF-EXISTENCE with quotas is NP-complete even for two agents with identical valuations.*

Despite that AEF allocation is hard to find, we show that AEF-1 allocation always exists just like EF-1.

PROPOSITION 1. *For any instance \mathcal{I} , an AEF-1 allocation always exists and can be found in polynomial time.*

PROOF. Consider the following allocation scheme:

- If $m \leq n$, agents $1, 2, \dots, m$ get their favorite item among the unallocated items in turns. The rest agents get nothing.
- If $m > n$, $1, 2, \dots, n-1$ get their favorite item among the unallocated items in turns, and agent n gets the rest of the items.

We show that allocation induced by this scheme is AEF-1. Consider two agents i and h . We show that i does not envy h up to one item.

$m \leq n$. If $i \leq m$, then i does not envy any $h > i$ because h gets either no item or an item inferior to i 's item under i 's valuation. If $h < i$, then i does not envy h by removing h 's only item. If $i > m$, then i does not envy any other h after removing h 's item (if exists).

$m > n$. If $h \neq n$, i does not envy h after removing h 's only item. If $h = n$, note that i picks their favorite item among all the rest of the items, including all h 's items. Therefore, for any $g \in A_h$, $v_i(A_i) \geq v_i(g)$. Therefore, $u_i(A_i) \geq u_i(A_h)$, and i does not envy h . \square

5 AEF-1 WITH QUOTAS

This section focuses on the complexity of AEF-1-EXISTENCE with quotas. AEF-1 with quotas could cover a wider range of real-world scenarios. Unfortunately, our first result shows that AEF-1-EXISTENCE with quotas is NP-complete even for binary valuations.

THEOREM 2. AEF-1-EXISTENCE with quotas is NP-complete even for binary valuations.

PROOF. We show a reduction from EF-EXISTENCE for binary valuations, which is known to be NP-complete [2, 23].

An EF-EXISTENCE with binary value instance consists of a set of agents, $N = [n]$, a set of items M ($|M| = m$), and a binary additive valuation profile V . The goal is to determine whether there is an envy-free allocation A such that for any pair of agents $i, h \in N$, $v_i(A_i) \geq v_i(A_h)$.

We construct a AEF-1-EXISTENCE with quotas instance as follows: $N' = N$, $M' = M \cup D$, where D is a set of $(n-1)m$ items with value 0 to all agents. Additive valuation profile V' is defined as follows: for each agent i , and item g , if $g \in M$, $v'_i(g) = v_i(g)$; otherwise, $v'_i(g) = 0$. $v'_i(\emptyset) = 0$. u' is the average value function of v' . The quota Q requires every agent to receive exactly m items.

(\Rightarrow) Suppose EF-EXISTENCE is a YES instance, and A^* is an envy-free allocation under V . Then we show that AEF-1-EXISTENCE with quotas is a YES instance. Let A' be an allocation in the AEF-1-EXISTENCE with quotas instance where each agent i gets all the items in A_i^* and fills up the quota with items in D . It follows from the definition that A' satisfies Q . Now we show that A' is AEF-1. Note that for any agents i, h , $u'_i(A'_h) = \frac{1}{m}v_i(A_h^*)$. Since A^* is envy-free, $v_i(A_i^*) \geq v_i(A_h^*)$. Therefore, $u'_i(A'_i) \geq u'_i(A'_h)$, and A' is an AEF (thus AEF-1) allocation.

(\Leftarrow) Suppose AEF-1-EXISTENCE with quotas is a YES instance, and A' is an AEF-1 allocation satisfying Q . We first show that A must also be an AEF allocation. Suppose this is not the case, and agent i envies agent h . We show i envies h even after removing one item. From binary valuation, we have $v'_i(A'_i) \leq v'_i(A'_h) - 1$. If agent i removes one item from A_h , the average value of A'_h is not smaller than A'_i , and A'_h has fewer items than A'_i . Therefore, A'_h still has a higher average value than A'_i , and i envies h . If agent i removes one item from A'_i , the average value will be no more than $\frac{v'_i(A'_h) - 1}{m-1}$. This value is strictly less than $u_i(A'_h) = \frac{v'_i(A'_h)}{m}$ for $v'_i(A'_h) < m$. If $v'_i(A'_h) = m$, then $v'_i(A'_i) = 0$ since M' contains at most m valuable items for i . Therefore, i still envies h after removing any item. This is a contradiction. Therefore, A' must also be AEF allocation.

Now we show that EF-EXISTENCE with binary value is also a YES instance. Let A^* be a allocation in EF-EXISTENCE instance such that $A_i^* = A'_i \cap M$ for every $i \in N$. Similarly, with relationship $u'_i(A'_h) = \frac{1}{m}v_i(A_h^*)$, the AEF-ness of A' implies the envy-freeness of A^* . \square

Due to the hardness of the problem, we turn to consider AEF-1 allocation with a *fixed* number of agents n . We show that, for binary valuations, AEF-1-EXISTENCE with quotas for a fixed number of agents is in P , in contrast with the hardness in the variable n case.

THEOREM 3. *There exists a polynomial-time algorithm that, given any instance of AEF-1-EXISTENCE with quotas for a fixed number of agents and binary valuations, decides if there exists an AEF-1 allocation satisfying the quota, and outputs an allocation if there exists.*

Aziz et al. [4] proposes a pseudo-polynomial time dynamic programming algorithm to find an EF-1 allocation maximizing social welfare given a constant number of agents. We apply their technique and propose Algorithm 1 to compute AEF-1-EXISTENCE with quotas for binary valuations.

State. A state in Algorithm 1 is a triplet (W, H, k) . Suppose A is a partial allocation of $M = \{g_1, g_2, \dots, g_m\}$ where g_1, g_2, \dots, g_k has been allocated. W is a n -vector that records the number of items each agent is allocated, i.e. $W_i = |A_i|$ for each i . H is a $n \times n$ -matrix that records the additive value of each agent toward each bundle, i.e. $H(i, h) = v_i(A_h)$. W and H together record each agent's average value on each bundle. For any pair of agents i, h , $u_i(A_h) = \frac{H(i, h)}{W_h}$. $k = 0, 1, 2, \dots, m$ indicates that item g_1, g_2, \dots, g_k has been allocated while other items are not. $k = 0$ means no item has been allocated yet. For each state, we maintain two values. $Vld(W, H, k) \in \{0, 1\}$ indicates whether this state is reached, which stands for there exists a partial allocation of whom the state is (W, H, k) . $Prev(W, H, k) \in N$ records the agent that item g_k is allocated to reach the current state. For any allocation, it is sufficient to judge whether it is AEF-1 and satisfies Q from its corresponding state.

State Transition. For a given state (W, H, k) with $k < m$ and $Vld(W, H, k) = 1$, we enumerate the agent to whom item g_{k+1} is allocated. For each agent i , we find the updated state $(W', H', k+1)$ after g_{k+1} is allocated to i and set $Vld(W', H', k+1) = 1$ and $Prev(W', H', k+1) = i$. The algorithm start from $(0^n, 0^{n \times n}, 0)$ and iterated for $k = 0, 1, \dots, m-1$. The search space of W is $\mathcal{W} = \{0, 1, \dots, m\}^n$, and the search space H is $\mathcal{H} = \{0, 1, \dots, m\}^{n \times n}$. Finally, the algorithm finds if there is a state $Vld(W, H, m) = 1$ that is AEF-1 and satisfies Q . If so, the algorithm outputs YES and constructs the allocation backward with $Prev$. Otherwise, the algorithm outputs NO.

The technique of enumerating all possible values in Algorithm 1 can be extended to valuations where a bundle has at most $Poly(m)$ different (additive) values. However, for general valuation, a bundle can have exponentially many values. In fact, we show that AEF-1-EXISTENCE with quotas with fixed $n \geq 3$ is NP-complete.

THEOREM 4. *AEF-1-EXISTENCE with quotas with fixed $n \geq 3$ is NP-complete.*

PROOF SKETCH. We propose a reduction from the computation problem of EQUAL-CARDINALITY PARTITION, a variation of PARTITION that requires equal size between two subsets and is also NP-complete [20]. An instance of EQUAL-CARDINALITY PARTITION consists of a multiset $X = \{x_1, x_2, \dots, x_{2k}\}$ where $x_i \in \mathbb{N}$. The goal is to determine whether X can be partitioned into two subsets Y and $X \setminus Y$ with equal size k and equal sum T .

Given an EQUAL-CARDINALITY PARTITION instance, we construct a AEF-1-EXISTENCE with quotas instance with three agents and $3k+6$ items. (If $n > 3$, we add agents that value all items as 0 and are required to receive no items by the

Algorithm 1 DP for AEF-1 with quota with binary valuation

Require: Agent set N , Item set M , binary valuation profile V , and quota Q .

Ensure: An AEF-1 allocation satisfying Q if it exists.

```

1: Initialization:  $Vld(0^n, 0^{n \times n}, 0) \leftarrow 1$ .
2: for  $k = 0, 1, \dots, m$  do
3:   for  $W \in \mathcal{W}$  and  $H \in \mathcal{H}$  such that  $Vld(W, H, k) = 1$  do
4:     for  $i = 1, 2, \dots, n$  do
5:       Update  $W', H'$  after assigning  $g_{k+1}$  to  $i$ .
6:        $Vld(W', H', k+1) \leftarrow 1$ .
7:        $Prev(W, H', k+1) \leftarrow i$ .
8:   for  $W \in \mathcal{W}, H \in \mathcal{H}$  such that  $Vld(W, H, m) = 1$  do
9:     if  $(W, H, m)$  is AEF-1 and satisfies  $Q$  then
10:    Construct the allocation from  $Prev$  backwards.
11:    return the allocation.
12: return NO

```

quota.) Agents share the same valuation function v and average value function u . The quota Q requires every agent to be allocated exactly $k + 2$ items. The item set $M = M_1 \cup M_2 \cup M_3$ consists of three parts:

- $M_1 = \{g_1, g_2, \dots, g_{2k}\}$, where $v(g_j) = x_j + k^2 T^2$.
Let $T' = \frac{1}{2} \sum_{g \in M_1} v(g) = T + k^3 T^2$.
- M_2 contains $k + 1$ copies of b with $v(b) = \frac{(k+2)T'}{(k+1)^2}$.
- M_3 contains five copies of 0 with $v(0) = 0$.

We state that the value of b is smaller than any item in M_1 . (Proof in the full version of this paper.)

LEMMA 1. *For any $g \in M_1$, $v(b) < v(g)$.*

(\Rightarrow) If EQUAL-CARDINALITY PARTITION is a YES instance, and Y and $(X \setminus Y)$ are an equal-size and equal-sum partition, we show the following allocation A is AEF-1 and satisfies Q .

- (1) $A_1 = \{g_j \mid x_j \in Y\} \cup \{0, 0\}$.
- (2) $A_2 = \{g_j \mid x_j \in (X \setminus Y)\} \cup \{0, 0\}$.
- (3) $A_3 = M_2 \cup \{0\}$.

It's not hard to verify that each agent gets exactly $k + 2$ items, $u(A_1) = u(A_2) = \frac{T'}{k+2}$, and $u(A_3) = \frac{T'}{k+1}$. Agent 1 and 2 does not envy each other, and agent 3 does not envy agent 1 and 2. When comparing with agent 3, agent 1 and agent 2 can remove an item 0 in their own bundle, and $u(A_1 \setminus \{0\}) = u(A_2 \setminus \{0\}) = u(A_3) = \frac{T'}{k+1}$. Therefore, A is AEF-1 and satisfies Q .

(\Leftarrow) If AEF-1-EXISTENCE with quotas is a YES instance, and A is an AEF-1 allocation satisfying Q . We show that EQUAL-CARDINALITY PARTITION is a YES instance in three steps.

First, no agent can have more than two item 0 in their bundles. Otherwise, the agent get at least three 0 envies the agent get at most one 0 even after removing one item.

Second, the agent with exactly one item 0 (agent 3, with loss of generality) must have all the item b . Otherwise, since $v(b) < v(g)$ for any $g \in M_1$, the average value of A_3 will exceed $\frac{T'}{k+1}$, and one of A_1 and A_2 will have average value strictly less than $\frac{T'}{k+2}$. Then the owner of this bundle will envy agent 3 even after removing one item.

Finally, agent 1 and agent 2's bundles must derive a equal-cardinality partition of X . Otherwise, the average value of the less-valuable bundle will be strictly less than $\frac{T'}{k+2}$. With the same reasoning as the second step, the owner of this

bundle will envy agent 3 (with $u(A_3) = \frac{T'}{k+1}$) even after removing one item. Therefore, EQUAL-CARDINALITY PARTITION is a YES instance. The full proof is available in the full version of this paper. \square

6 APPROXIMATION ON AEF-1 WITH QUOTAS

The NP-hardness on AEF-1-EXISTENCE with quotas urges us to look into approximation results. A natural idea is to round the value of items so that a bundle can have at most $\text{Poly}(m)$ different values and apply the procedure of Algorithm 1.

For simplicity of calculation, we assume $\max_{i,g} v_i(g) = 1$. We propose two approximation notions of AEF-1 based on the additive error and multiplicative ratio respectively.

DEFINITION 7 (ε -ERROR AEF-1). *Given $\varepsilon \geq 0$, an allocation A is ε -error AEF-1 if for any pairs of agents $i, h \in N$, there exists an item $g \in A_i \cup A_h$ such that $u_i(A_i \setminus \{g\}) \geq u_i(A_h \setminus \{g\}) - \varepsilon$.*

DEFINITION 8 (α -AEF-1). *Given $0 < \alpha \leq 1$, an allocation A is α -AEF-1 if for any pairs of agents $i, h \in N$, there exists an item $g \in A_i \cup A_h$ such that $u_i(A_i \setminus \{g\}) \geq \alpha \cdot u_i(A_h \setminus \{g\})$.*

PROPOSITION 2. *Given the assumption that $\max_{i,g} v_i(g) = 1$, if an allocation is α -AEF-1, then it is $(1 - \alpha)$ -error AEF-1.*

PROOF. From α -AEF-1 we know that $u_i(A_i \setminus \{g\}) \geq \alpha \cdot u_i(A_h \setminus \{g\})$. Therefore,

$$\begin{aligned} u_i(A_i \setminus \{g\}) &\geq \alpha \cdot u_i(A_h \setminus \{g\}) \\ &= u_i(A_h \setminus \{g\}) - (1 - \alpha) \cdot u_i(A_h \setminus \{g\}) \\ &\geq u_i(A_h \setminus \{g\}) - (1 - \alpha). \end{aligned}$$

The last inequality comes from $u_i(A_h \setminus \{g\}) \leq 1$. \square

Proposition 2 tells us that α -AEF-1 implies ε -error AEF-1 given bounded valuations. However, ε -error AEF-1 does not guarantee α -AEF-1, as shown in the followings (Example 4). Our goal is to find an approximation algorithm that returns an α -AEF-1 with α close to 1 if possible. We first introduce our rounding scheme.

Rounding. Given the rounding parameter $r \in \mathbb{N}^+$ and the upper bound $a > 0$, we divide $[0, a]$ into $r + 1$ intervals $\{0\}, (0, \frac{a}{r}], (\frac{a}{r}, \frac{2a}{r}], \dots, (\frac{(r-1)a}{r}, a]$. For $k = 1, 2, \dots, r$, a positive value $\frac{(k-1)a}{r} < x \leq \frac{ka}{r}$ is rounded to $\frac{ka}{r}$. 0 is rounded to 0.

If we apply the same rounding scheme to each $v_i(g)$ and directly apply Algorithm 1, we will be able to find an $\frac{2a}{r}$ -error AEF-1 allocation, because an AEF-1 allocation in the original valuations implies an $\frac{a}{r}$ -error AEF-1 allocation in the rounded valuations, and an $\frac{a}{r}$ -error AEF-1 allocation in the rounded valuations in turn implies an $\frac{2a}{r}$ -error AEF-1 allocation in the original valuations. However, there is no guarantee on α -AEF-1, because the value of an item can be rounded from arbitrarily small to $\frac{a}{r}$. Example 4 shows a case where an AEF-1 allocation in the rounded valuation turns out to be a poor approximation in the original valuation.

EXAMPLE 4. *Given any rounding parameters a, r (assuming $a < r$), consider an instance with more than two agents and more than three items. The table below describes an allocation where the first two agents 1 and 2 are allocated the first three items g_1, g_2, g_3 .*

For item g_3 , $\alpha \in (0, 1)$ is an arbitrarily small positive value. It is not hard to verify that A is not AEF-1, as agent 2 envies agent 1 even after removing an item. Moreover, the allocation is no better than α -AEF-1. However, after rounding, the value

		g_1	g_2	g_3
		$\frac{a}{r}$	$\frac{a}{r}$	$\frac{\alpha a}{r}$
1	$\frac{a}{r}$	$\frac{a}{r}$	$\frac{\alpha a}{r}$	
	$\frac{a}{r}$	$\frac{a}{r}$	$\frac{\alpha a}{r}$	
2				

Table 3. Allocation A where rounding leads to poor approximation.

of g_1 and g_2 is unchanged while the value of g_3 is rounded to $\frac{a}{r}$. Then A is AEF-1 in the rounded valuation. Therefore, if Algorithm 1 finds (the state of) A , it returns an α -AEF-1 allocation where α can be arbitrarily small.

Although the instance has an AEF-1 allocation A' where agent 1 gets g_1 and g_3 and agent 2 gets g_2 , the algorithm may not find A' . Note that after rounding g_2 and g_3 both have value of $\frac{a}{r}$. This means that A and A' share the same state in Algorithm 1. Which allocation is constructed depends on the *Prev* record. By carefully manipulating the order of items, we can let the algorithm returns A rather than A' .

Therefore, we need a more refined rounding and searching scheme that can distinguish between A and A' to ensure a closer approximation ratio between the original and the rounded valuation. The rounding of each agent should be proportional to their valuations so that the rounding error is not too large compared with the value of their own bundles. Bu et al. [12] proposes a bi-criteria approximation algorithm to maximize EF-1 ratio and social welfare simultaneously. We follow their techniques to enumerate all items being removed in the envy comparisons, i.e. the “1” in “AEF-1”.

Removing matrix. A removing matrix R is a $n \times n$ matrix recording the items to remove when agents compare bundles with each other. For every pair of agents i, h , $R(i, h) = (g, l) \in (M \cup \{\emptyset\}) \times \{i, h\}$. The first value g is the item to remove when agent i compares their bundle with h 's bundle, and the second value l indicates whether g belongs to i or h . $g = \emptyset$ means i does not remove any item when comparing with h . In this case, l makes no difference. Specifically, $R(i, i) = (\emptyset, i)$ for each agent i . A removing matrix is valid if it derives a partial allocation of M . That is, it does not contain two entries (g, l_1) and (g, l_2) such that $l_1 \neq l_2$.

Our algorithm runs in four steps. We enumerate all valid removing matrices R . For each R , we first allocate the items that have been pre-allocated by R . Next, we round the values of the unallocated items based on each agent's valuation. Then, we run the dynamic programming search on the unallocated items under the rounded valuation to validate all possible states. Finally, we search states where an agent will not envy another agent by their rounding error under the rounded valuation after removing one item. If such a state exists, the algorithm returns the corresponding allocation. Otherwise, the algorithm returns NO. A detailed description of the algorithm is available in the full version of this paper.

Step 1: pre-allocation. Given an removing matrix R , let M^R be the set of item g such there there exists an entry $(g, i) \in R$. We allocate items in M^R according to R . Let W_0^R and H_0^R be the vector of bundle sizes and the valuation matrix after M^R has been allocated.

Step 2: rounding. For each agent i , let M_i^R be the set of i 's removing items, and Let $M'_i = M \setminus M_i^R$. For each agent i , we set the rounding upper-bound $a = \max_{g \in M'_i} v_i(g)$. $r = m^2 n^2$ is the same for all agents. We create the rounded valuation v_i^R by rounding the values for all items in M'_i . That is, for any agent i , $v_i^R(g) = v_i(g)$ if $g \in M_i^R$, and $v_i^R(g)$ is rounded to the closest larger $\frac{ka_i}{r}$ if $g \in M'_i$. Let u_i^R be the average value function of v_i^R . Given a fixed M_i^R , each bundle will have at most $\text{Poly}(m)$ possible values. Precisely, v_i^R takes value from $\{0, \frac{a_i}{r}, \frac{2a_i}{r}, \dots, |M'_i| \cdot a_i\}$.

Step 3: dynamic programming. The dynamic programming follows the same procedure of Algorithm 1 that enumerates and iteratively validates states. The difference is that we run the dynamic programming only on $M' = M \setminus M^R$, i.e. the unallocated items in step 1. The start state is $(W_0^R, H_0^R, 0)$ which represent the state where all items in M^R and no items in M' has been allocated.

Step 4: searching. Finally, we search if there is a state that satisfies two conditions. (1) The state satisfies the quota Q . (2) For any agent i and h , i does not envy h by more than $\frac{a_i}{r}$ under the rounded valuation u_i^R , after removing one item, i.e. $u(A_i \setminus \{g\}) \geq u(A_i \setminus \{g\}) - \frac{a_i}{r}$ for some g . If such a state exists, the algorithm constructs the allocation backward and returns it. If such allocation does not exist for any state and any R , the algorithm returns NO. The reason for searching a bounded envy allocation rather than an AEF-1 allocation is to guarantee that the algorithm will always return NO if there does not exist an AEF-1 allocation (Theorem 5).

The following lemma shows that, by rounding the value function of each agent based on the most valuable item in M'_i , the average value of an agent's bundle is lower bound by the average value of M'_i .

LEMMA 2. *Given any removing matrix R and allocation A , and for any agent i , if i does not envy any other agent by more than $\varepsilon > 0$ under R and u^R , then $u_i^R(A_i) \geq \frac{1}{n}u_i^R(M'_i) - \varepsilon$.*

PROOF SKETCH. Suppose this is not true, and agent i has a bundle A_i with average value less than $\frac{1}{n}u_i^R(M'_i) - \varepsilon$. We consider the agent h that takes the share of M'_i with the largest average value under i 's valuation. Then i envies h by more than ε even after removing one item, which is a contradiction. The full proof is available in the full version of this paper. \square

Lemma 2 guarantees that the rounding error is small compared with the bundle's average value. Note that M'_i contains at most m items, and the largest item has a value of a_i . Therefore, (in i 's valuation,) the average value of M'_i is at least $\frac{a_i}{m}$, and the average value of A_i is at least $\frac{a_i}{mn} - \varepsilon$. On the other hand, the rounding error is $\frac{a_i}{r} = \frac{a_i}{m^2n^2}$. With a lower bound of average value and a bounded error, a reasonably good approximation ratio can be guaranteed, as shown in Theorem 5.

THEOREM 5. *Given any instance of AEF-1-EXISTENCE with quotas (\mathcal{I}, Q) ,*

- (1) *if the algorithm returns NO, then (\mathcal{I}, Q) does not have an AEF-1 allocation satisfying Q .*
- (2) *if the algorithm returns YES, it gives a $(1 - \frac{4}{mn})$ -AEF-1 allocation satisfying Q .*

PROOF SKETCH. (NO case) We turn to prove the equivalent statement that if (\mathcal{I}, Q) exists an AEF-1 allocation satisfying Q , then the algorithm always returns YES. Suppose A is an AEF-1 allocation satisfying Q , and R is a removing matrix of A that achieves AEF-1. We show that any agent i will not envy another agent h by more than $\frac{a_i}{r}$ under u^R after removing one item. For the original valuation u , we have $u_i(A_i \setminus \{g\}) \geq u_i(A_h \setminus \{g\})$ for any i, h and some g . For the rounded valuation, we have $v_i(g) \leq v_i^R(g) \leq v_i(g) + \frac{a_i}{r}$. Therefore, the average value of any bundle will neither decrease nor increase more than $\frac{a_i}{r}$ in the rounded valuation. Therefore, $u_i^R(A_i \setminus \{g\}) - u_i^R(A_h \setminus \{g\}) \geq -\frac{a_i}{r}$, and agent i does not envy agent h by more than $\frac{a_i}{r}$. In the process of the algorithm, the state of A will be validated in Step 3 and found in Step 4, and the algorithm will return YES.

(YES case) We show that if an allocation A satisfies that any agent i will not envy another agent h by more than $\frac{a_i}{r}$ under u^R after removing one item, then A will be $(1 - \frac{4}{mn})$ -AEF-1 under u . Consider any pair of agent i and h and suppose i still envies h even after removing one item under the original valuation u . (If such a pair does not exist, then A is an AEF-1 allocation, and the statement holds.) For simplicity, let A'_i and A'_h be the bundles after agent i

removing item g . Since envy is bounded by $\frac{a_i}{r}$ in the rounded valuation u^R , it is also bounded in the original valuation: $u_i(A'_i) \geq u_i(A'_h) - \frac{2a_i}{r}$. On the other hand, since i envies h , $u_i(A'_h) > u_i(A'_i)$. Consider two subcases.

(1) $u_i(A'_i) \geq u_i(A_i)$. In this case agent i either remove an item in A_h or an item with a value lower than average in A_i . In this case, we have $u_i(A'_h) > u_i(A_i)$. With the rounding, we know $u_i(A_i) \geq u_i^R(A_i) - \frac{a_i}{r}$. And from Lemma 2, $u_i^R(A_i) \geq \frac{1}{n}u_i^R(M'_i) - \frac{a_i}{r} \geq \frac{a_i}{mn} - \frac{a_i}{r}$. Aggregating all these inequalities, we get

$$\begin{aligned} u_i(A'_i) &\geq u_i(A'_h) - \frac{2a_i}{r} \\ &= (1 - \frac{1}{4mn})u_i(A'_h) + \frac{1}{4mn}u_i(A'_h) - \frac{2a_i}{r} \\ &\geq (1 - \frac{1}{4mn})u_i(A'_h) + \frac{1}{4mn} \cdot \frac{a_i}{mn} - \frac{4a_i}{r} \\ &= (1 - \frac{1}{4mn})u_i(A'_h). \end{aligned}$$

(2) $u_i(A'_i) < u_i(A_i)$. In this case, i removes an item in A_i with a value higher than average. With a similar reasoning, we turn to show that $u_i(A_i) \geq (1 - \frac{1}{4mn})u_i(A_h)$.

Therefore, A is an $(1 - \frac{4}{mn})$ -AEF-1 allocation under the original valuation u . \square

7 CONCLUSION AND FUTURE WORK

In this paper, we propose average envy-freeness where envy is defined on the average value of a bundle. AEF provides a fairness criterion for allocation problems in scenarios where agents have different entitlements and the entitlements depend on the allocation itself. We study the existence and complexity of AEF and AEF-1. While deciding the existence of an AEF allocation is NP-hard, an AEF-1 allocation always exists and can be computed in polynomial time. We also study the complexity of AEF-1 with quotas. While AEF-1 with quotas is NP-complete to decide, we provide polynomial-time algorithms for instances with a constant number of agents to find AEF-1 allocation under binary valuation and approximated AEF-1 allocation under general valuation.

The notion of average envy-freeness can be extended in multiple aspects. One extension is scenarios with multiple copies. For example, in a paper review scenario, a paper should be reviewed by multiple papers, but a reviewer cannot review a paper multiple times. Another extension is scenarios where items bring different entitlements to agents. We expect the entitlement of agents to be the sum of entitlements of items in their bundles, and envy is defined on the sum of values divided by the entitlement of the agent. It is also an intriguing direction to find relaxations of AEF other than AEF-1. An interesting observation is that “AEF-X” is an even stronger notion than AEF since agents should not envy each other even if their bundle’s average value is decreased by removing the most valuation item.

ACKNOWLEDGMENTS

We thank all anonymous reviewers for their helpful comments and suggestions. LX acknowledges NSF, #2007994, and #2106983, and #2303000 for support. The research of Biaoshuai Tao was supported by the National Natural Science Foundation of China (Grant No. 62102252).

REFERENCES

[1] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. 2022. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi-Agent Systems* 36 (2022), 1–21.

- [2] Haris Aziz, Serge Gaspers, Simon Mackenzie, and Toby Walsh. 2015. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227 (2015), 71–92.
- [3] Haris Aziz, Xin Huang, Nicholas Mattei, and Erel Segal-Halevi. 2019. The constrained round robin algorithm for fair and efficient allocation. *arXiv preprint arXiv:1908.00161* (2019).
- [4] Haris Aziz, Xin Huang, Nicholas Mattei, and Erel Segal-Halevi. 2022. Computing welfare-Maximizing fair allocations of indivisible goods. *European Journal of Operational Research* (2022). <https://doi.org/10.1016/j.ejor.2022.10.013>
- [5] Haris Aziz, Hervé Moulin, and Fedor Sandomirskiy. 2020. A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation. *Operations Research Letters* 48, 5 (2020), 573–578. <https://doi.org/10.1016/j.orl.2020.07.005>
- [6] Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2021. Fair-Share Allocations for Agents with Arbitrary Entitlements. In *Proceedings of the 22nd ACM Conference on Economics and Computation* (Budapest, Hungary). Association for Computing Machinery, New York, NY, USA, 127. <https://doi.org/10.1145/3465456.3467559>
- [7] Moshe Babaioff, Noam Nisan, and Inbal Talgam-Cohen. 2019. Fair Allocation through Competitive Equilibrium from Generic Incomes. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA). Association for Computing Machinery, New York, NY, USA, 180. <https://doi.org/10.1145/3287560.3287582>
- [8] Gerdus Benadè, Daniel Halpern, Alexandros Psomas, and Paritosh Verma. 2023. On the Existence of Envy-Free Allocations Beyond Additive Valuations. *arXiv:2307.09648 [cs.GT]*
- [9] Umang Bhaskar, A. R. Sricharan, and Rohit Vaish. 2021. On Approximate Envy-Freeness for Indivisible Chores and Mixed Resources. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2021) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 207)*, Mary Wootters and Laura Sanità (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 1:1–1:23.
- [10] Arpita Biswas and Siddharth Barman. 2018. Fair Division under Cardinality Constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence* (Stockholm, Sweden). AAAI Press, 91–97.
- [11] Arpita Biswas and Siddharth Barman. 2019. Matroid Constrained Fair Allocation Problem. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence* (Honolulu, Hawaii, USA). AAAI Press, Article 1259, 2 pages. <https://doi.org/10.1609/aaai.v33i01.33019921>
- [12] Xiaolin Bu, Zihao Li, Shengxin Liu, Jiaxin Song, and Biaoshuai Tao. 2022. On the Complexity of Maximizing Social Welfare within Fair Allocations of Indivisible Goods. *arXiv preprint arXiv:2205.14296* (2022).
- [13] Kristóf Bérczi, Erika R. Bérczi-Kovács, Endre Boros, Fekadu Tolessa Gedefà, Naoyuki Kamiyama, Telikepalli Kavitha, Yusuke Kobayashi, and Kazuhisa Makino. 2020. Envy-free Relaxations for Goods, Chores, and Mixed Items. *arXiv:2006.04428 [econ.TH]*
- [14] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. 2021. Weighted Envy-Freeness in Indivisible Item Allocation. *ACM Transactions on Economics and Computation* 9, 3, Article 18 (August 2021), 39 pages. <https://doi.org/10.1145/3457166>
- [15] Mithun Chakraborty, Ulrike Schmidt-Kraepelin, and Warut Suksompong. 2021. Picking sequences and monotonicity in weighted fair division. *Artificial Intelligence* 301 (2021), 103578. <https://doi.org/10.1016/j.artint.2021.103578>
- [16] Mithun Chakraborty, Erel Segal-Halevi, and Warut Suksompong. 2022. Weighted Fairness Notions for Indivisible Items Revisited. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 5 (Jun. 2022), 4949–4956. <https://doi.org/10.1609/aaai.v36i5.20425>
- [17] Amitay Dror, Michal Feldman, and Erel Segal-Halevi. 2021. On Fair Division under Heterogeneous Matroid Constraints. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 6 (May 2021), 5312–5320. <https://doi.org/10.1609/aaai.v35i6.16670>
- [18] Alireza Farhadi, Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Sébastien Lahaie, David Pennock, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2019. Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research* 64 (2019), 1–20.
- [19] Jiarui Gan, Bo Li, and Xiaowei Wu. 2021. Approximately envy-free budget-feasible allocation. *arXiv preprint arXiv:2106.14446* (2021).
- [20] Michael Garey and David Johnson. 1979. *Computers and Intractability*. W. H. Freeman and Company.
- [21] Mohammad Ghodsi, Mohammadtaghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2018. Fair Allocation of Indivisible Goods: Improvements and Generalizations. In *Proceedings of the 2018 ACM Conference on Economics and Computation* (Ithaca, NY, USA). Association for Computing Machinery, New York, NY, USA, 539–556.
- [22] Mohammad Ghodsi, Mohammadtaghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2022. Fair allocation of indivisible goods: Beyond additive valuations. *Artificial Intelligence* 303 (2022), 103633.
- [23] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, Heju Wang, and Lirong Xia. 2020. Fair Division Through Information Withholding. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 02 (April 2020), 2014–2021. <https://doi.org/10.1609/aaai.v34i02.5573>
- [24] Bo Li, Yingkai Li, and Xiaowei Wu. 2022. Almost (Weighted) Proportional Allocations for Indivisible Chores. In *Proceedings of the ACM Web Conference 2022* (Virtual Event, Lyon, France). Association for Computing Machinery, New York, NY, USA, 122–131. <https://doi.org/10.1145/3485447.3512057>
- [25] Justin Payan and Yair Zick. 2022. I Will Have Order! Optimizing Orders for Fair Reviewer Assignment. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, Lud De Raedt (Ed.). International Joint Conferences on Artificial Intelligence Organization, 440–446. <https://doi.org/10.24963/ijcai.2022/63>
- [26] Cristina Roadevin and Harry Hill. 2021. How can we decide a fair allocation of healthcare resources during a pandemic? *Journal of Medical Ethics* 47, 12 (2021), e84–e84.
- [27] Warut Suksompong and Nicholas Teh. 2022. On maximum weighted Nash welfare for binary valuations. *Mathematical Social Sciences* 117 (2022), 101–108. <https://doi.org/10.1016/j.mathsocsci.2022.03.004>

- [28] Wei Wang, Ben Liang, and Baochun Li. 2015. Multi-Resource Fair Allocation in Heterogeneous Cloud Computing Systems. *IEEE Transactions on Parallel and Distributed Systems* 26, 10 (2015), 2822–2835. <https://doi.org/10.1109/TPDS.2014.2362139>
- [29] Xiaowei Wu, Cong Zhang, and Shengwei Zhou. 2023. Weighted EF1 Allocations for Indivisible Chores. *Forthcoming on the 23rd ACM Conference on Economics and Computation* (2023).