Fair Allocation of Divisible Goods under Non-Linear Valuations

Haris Aziz
UNSW Sydney
Sydney, Australia
haris.aziz@unsw.edu.au

Zixu He
UNSW Sydney
Sydney, Australia
tao.he@student.unsw.edu.au

Xinhang Lu
UNSW Sydney
Sydney, Australia
xinhang.lu@unsw.edu.au

Kaiyang Zhou
UNSW Sydney
Sydney, Australia
kaiyang.zhou@student.unsw.edu.au

ABSTRACT

We study the problem of dividing homogeneous divisible goods among agents with non-linear valuations. Specifically, the value that an agent gains from a given good depends only on the amount of the good they receive, and is not necessarily linear with respect to the amount. For instance, under one-breakpoint piecewise-constant valuations, each agent specifies a threshold for each good such that this agent receives utility zero (resp., full utility of the good) when getting an amount below (resp., at least) the threshold. Given non-linear valuations that are additive across the goods, we focus on designing fair allocation algorithms and consider two well-known fairness properties: the maximin share (MMS) guarantee and envy-freeness (EF). For MMS, we devise an algorithm which always produces a $\frac{1}{2n-1}$-MMS allocation for n agents with arbitrary non-decreasing valuations. It is worth noting that this algorithmic result is almost tight as we give an impossibility of guaranteeing more than $1/n$ approximation to MMS, even when agents have one-breakpoint piecewise-constant valuations. For $n \leq 3$ agents, we show the ratio $1/n$ is *tight*. For EF, we show it is NP-hard to check the existence of an EF and Pareto optimal (PO) allocation for n agents and at least three goods, even when agents have one-breakpoint piecewise-constant valuations. We complement the hardness result by considering the case with a single divisible good, and devising a polynomial-time algorithm to check whether an EF and PO allocation exists or not for agents with piecewise-linear valuations.

KEYWORDS

Fair division; Divisible good; Non-linear valuation; MMS; Envy-free

ACM Reference Format:
Haris Aziz, Zixu He, Xinhang Lu, and Kaiyang Zhou. 2025. Fair Allocation of Divisible Goods under Non-Linear Valuations. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

1 INTRODUCTION

The allocation of scarce resources among multiple agents with different preferences is a fundamental issue that arises frequently in our society, for example, when dividing cloud computing resources such as processing time, memory, and communication bandwidth, or when handing out research grants. We often want to ensure that the allocation is *fair* to the agents, and possibly ideal in terms of other desiderata such as computation tractability, economic efficiency, etc. When the resource to be allocated is divisible and *heterogeneous*, the problem is commonly known as *cake cutting* [36], with the cake serving as a metaphor for heterogeneous divisible resource, and has been extensively studied by mathematicians, economists, political scientists, and more recently computer scientists [15, 29, 33, 35, 38].

The rich literature of cake cutting provides several ways to capture fairness, with the two most prominent notions are *proportionality* (a fair-share-based concept) and *envy-freeness* (a comparison-based fairness concept). An allocation is said to be proportional if each of the n agents receives a utility of at least $1/n$ of her total value for all resources [36], and envy-free (EF) if every agent values her own bundle the most in the allocation [24]. Both notions can always be satisfied [36, 37], admit discrete and bounded protocols in the Robertson-Webb query model [5, 35], and have been examined together with economic efficiency considerations such as being Pareto optimal (PO) or maximizing social welfare [7, 12, 20, 34].

As is common in the cake cutting literature, agents are assumed to have additive valuations across different pieces of the cake, yet the valuation function within each piece can be highly complicated. Despite being a versatile model for fair division of divisible resources such as land, time, money, and computational resources, it fails to capture the natural scenario where agents care only about how much each divisible resource they receive rather than which part, and potentially do not have linear valuations in proportion to the amount they receive, as discussed in a few papers [9, 17, 18, 23]. More specifically, consider the following real-world applications.

- *Computational resources.* Given computational resources such as CPU time, memory and communication bandwidth to be divided between various users, each user requires at least a certain amount of resources to achieve meaningful performance in their computational tasks. Some user may have CPU-intensive tasks and thus benefit more when receiving more CPU-time. Another user may have a small file, so she gradually gets benefits in proportion to the size of her file being included in the memory but will stop getting more benefits once the file is fully included.
- *Grant money.* Agents' utilities are typically assumed to be proportional with respect to the amount of money they receive. However, this may not be the case if, for example, a funding agency has specific guidelines on how much each budget category can

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA.

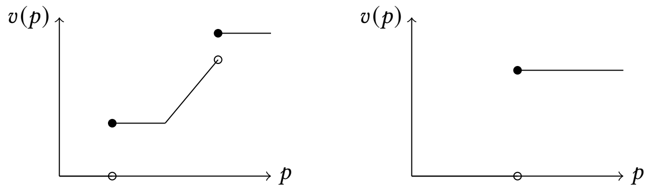

(a) Piecewise-linear utility function. **(b) One-breakpoint piecewise-constant utility function.**

Figure 1: Examples of non-linear valuations, where $v(p)$ specifies the value of a fraction p of the respective good under v.

be spent. While some research groups may require equipment for experiments and be satisfied with receiving a large amount of funding in the "Equipment" category, other theory groups may find that most of the funding in this category is irrelevant.

- *Space.* Lastly, consider the scenario where a hall is shared among various community groups. One community group needs at least half the hall to be able to organize a dinner event, and does not gets any additional benefit until it gets to book the whole hall in which case the full music system can also be played.

Given that agents have valuations depending on the amount of each homogeneous divisible resource they receive and the valuations are non-linear*, how shall we fairly allocate the resources among agents?*

1.1 Our Contributions

In this paper, we study the fair allocation of divisible resources and consider a model deviating from the cake-cutting literature. In more detail, there are m homogeneous divisible goods to be allocated among n agents. Each agent has a *non-decreasing* utility function, which specifies the value this agent receives from a given good: the value depends only on the amount of the good she receives and is not necessarily linear with respect to the amount. Agents' utilities are assumed to be additive across different goods. It is worth noting that our model captures the setting of indivisible-goods allocation with additive utilities, by letting each agent have any value only if the agent gets a good in its entirety. The *maximin share (MMS) guarantee* is a natural relaxation of proportionality when allocating indivisible goods [16], and has been extensively studied in various fair division settings [see, e.g., 2].

The majority of our results concern structured utility functions such as being piecewise-linear (see Figure 1(a)) or piecewise-constant with one breakpoint (see Figure 1(b)). The central fairness concepts in this paper are *MMS* and *envy-freeness*. In Section 3, we focus on MMS, which is satisfied if agents get bundles worth at least their own maximin share—the largest value an agent can guarantee for herself if she partitions the goods into n parts and gets the worst part. We present a polynomial-time approximation scheme (PTAS) to compute approximate MMS values for a constant number of agents with one-breakpoint piecewise-constant valuations.

When allocating indivisible goods to agents with additive utilities, a *constant* multiplicative approximation to MMS can always be achieved [28], with the state-of-the-art factor of $\left(\frac{3}{4} + \frac{3}{3836}\right)$ due to [1] and an upper bound of $\frac{39}{40}$ due to [22]. In our model, however, it is impossible to guarantee more than $1/n$ approximation to MMS, even for one-breakpoint piecewise-constant valuations. We complement this negative result by devising an algorithm that always produces a $\frac{1}{2n-1}$-MMS allocation for agents with *arbitrary non-decreasing* valuations, asymptotically matching the upper bound. We then turn to special cases involving up to three agents with one-breakpoint piecewise-constant valuations, where we show the approximation ratio $1/n$ to MMS is *tight*. Indeed, our motivating examples demonstrate that one-breakpoint piecewise-constant valuations naturally capture a wide range of real-world applications. Moreover, a good few fair division applications including dividing resources between different faculties within an institution or assets between founding members of a company often involves a small number of participants, and quite a few prominent fair division works deal exclusively with up to four agents [e.g., 3, 6, 11, 13, 14, 19, 26, 27].

In Section 4, we focus on envy-freeness, which can be satisfied trivially and vacuously by dividing each good equally among the agents. We thus also aim to achieve economic efficiency, which is not guaranteed by dividing each good equally. We show it is NP-hard to check the existence of an EF and Pareto optimal (PO) allocation, even for *three* goods and one-breakpoint piecewise-constant valuations. For a single good and piecewise-linear valuations, we devise a polynomial-time algorithm that (i) finds an allocation being EF and PO among all EF allocations, and (ii) checks the existence of an EF and PO allocation.

1.2 Related Work

Besides cake cutting, the fair allocation of indivisible goods has received extensive attention [2, 31, 39]. More recently, the fair allocation of resources of mixed nature has also been explored [30].

Perhaps the works most closely related to ours are the papers by Caragiannis et al. [18] and Bei et al. [9]. Caragiannis et al. studied the fair allocation of homogeneous divisible goods in which each agent's value depends only on the amount of each good they receive and valuations are additive across different goods. They mainly focused on randomized algorithms that achieve ex-ante envy-freeness and ex-ante approximate-PO among all envy-free lotteries; note that each deterministic allocation in the support is only required to be feasible. They also considered more general valuations, and thus adopted a query model to elicit agents' valuations and focused on query complexity. In contrast, we take a deeper dive into more structured valuations and focus on deterministic allocations which has guaranteed fairness ex post. To this end, we also examine MMS that has not been explored for this model.

Bei et al. [9] studied a fair division model with *subjective divisibility*, in which each good is either completely indivisible to some agents or completely divisible to other agents (i.e., the agents' valuation functions are linear with respect to the fraction of each good). They showed an impossibility of guaranteeing more than $\frac{2}{3}$-MMS, and this is *tight* for 2 and 3 agents. A $\frac{1}{2}$-MMS allocation always exists for any number of agents. They also adapted an envy-freeness relaxation called "EFM" [8] and investigated its compatibility with economic efficiency concepts. Our model generalizes theirs as we allow a broader class of valuations. Nevertheless, we still manage to give (asymptotically) tight approximation to MMS in various cases.

There have been works examining the division of a *single* homogeneous divisible good with non-linear valuations [17, 23]. Feige and Tennenholtz [23] presented a randomized allocation mechanism that gives each agent at least 1/2 of her *fair share* (i.e., the expected utility that she would get if she could choose the allocation rule that maximizes her expected utility). Buermann et al. [17] assumed that the available amount of the single divisible good is given by a probability distribution, and studied the trade-off between social welfare and envy-freeness as well as showed computational intractability of optimizing ex-ante social welfare subject to ex-ante EF (where randomness comes from the amount of the good).

When allocating resources of different types (e.g., computing resources) to agents with heterogeneous demands for each resource, a typical and classic assumption is that agents demand the resources in fixed proportions, known as *Leontief preferences* [25, 32] in the economics literature. This assumption requires that the resources are divisible and agents receive utilities in proportion to the resources allocated to them. Parkes et al. [32] made a more practical assumption by requiring a minimum, *indivisible* bundle of resources for the agents to receive utilities, and they conceptualized it as a step function. Their negative results (on the incompatibility between PO, strategyproofness and fairness) hold with a single resource, meaning that these continue to hold in our setting. The notable difference between their and our settings is that we assume additive utilities between the resources, instead of having the fixed proportion relations between the resources.

2 PRELIMINARIES

For any positive integer t, let $[t] \coloneqq \{1, \dots, t\}$. Our model includes a set of n agents $N = [n]$ to whom we allocate a set of m homogeneous divisible goods $M = \{g_1, g_2, \dots, g_m\}$. We assume in our work that each agent's value derived from each good depends on the fraction of the good allocated to the agent and that this derived value is not necessarily in proportion to the fraction allocated. Precisely, each agent $i \in N$ has a utility function $v_i \colon M \times [0,1] \to \mathbb{R}_{\geq 0}$ such that $v_i(g, p)$ specifies the value of a fraction p of item $g \in M$ allocated to agent i. Throughout the paper, we assume normalization, meaning that $v_i(g, 0) = 0$ for all $i \in N$ and $g \in M$, and monotonicity (a.k.a., "free disposal"), meaning that for all $i \in N$ and $g \in M$, $v_i(g, p) \leq v_i(g, p')$ if $p \leq p'$. For simplicity, we will write $v_i(g) = v_i(g, 1)$.

A *bundle* of goods M is represented by an m-dimensional vector $\mathbf{x} = (x_1, x_2, \dots, x_m)$, where each coordinate $x_j \in [0, 1]$ denotes the fraction of good g_j in bundle $\mathbf{x}$. For ease of expression, we will use "$x_j \cdot g_j$" to denote the part of good g_j allocated to a bundle. The bundle $\mathbf{x}$ is said to be *integral* if all goods in the bundle are included entirely, i.e., $x_j \in \{0, 1\}$ for all $j \in [m]$. In what follows, we will make it explicitly clear if we refer to an integral bundle.

We assume that agents' valuations across different goods are *additive*, i.e., given a bundle $\mathbf{x}$, for all $i \in N$, $v_i(\mathbf{x}) \coloneqq \sum_{j \in [m]} v_i(g_j, x_j)$. Denote by $A = (A_{i,g})_{i \in N, g \in M}$ the *allocation* of goods M among the agents, where $A_{i,g}$ specifies the fraction of item g allocated to i, $\sum_{i \in N} A_{i,g} = 1$ for each $g \in M$ is the feasibility constraint, and $A_i = (A_{i,g})_{g \in M}$ denotes agent i's bundle. Each agent i derives a utility of $v_i(A_i) \coloneqq \sum_{g \in M} v_i(g, A_{i,g})$ under allocation A.

Under this setting, we consider a few interesting sub-classes of general utility functions for the agents, namely piecewise-linear, piecewise-constant, and one-breakpoint piecewise-constant functions. Note that each of the latter is a proper sub-class of the former classes. A utility function v supported on $[0, 1]$ is said to be

- *piecewise-linear* if, for any $g \in M$, $v(g, \cdot)$ can be written as a collection of linear functions; see Figure 1(a) for an illustration. More formally, for some positive integer d:
$$v(g,p) = \begin{cases} a_1 \cdot p + b_1 & \text{if } p \in [c_0, c_1) \\ a_2 \cdot p + b_2 & \text{if } p \in [c_1, c_2) \\ \dots & \\ a_d \cdot p + b_d & \text{if } p \in [c_{d-1}, c_d], \end{cases}$$
where, $c_0 = 0$, $c_d = 1$, and for each $j \in [d]$, $a_j, b_j \in \mathbb{R}$ are the coefficients of the corresponding linear function in segment $[c_{j-1}, c_j]$. We say $c_1, c_2, \dots, c_{d-1}$ are the *breakpoints* of the utility function. Due to the normalization assumption, we have $b_1 = 0$. Moreover, since all utility functions are assumed to be non-decreasing, we must have $a_j \geq 0$ for each $j \in [d]$ and $a_j \cdot c_j + b_j \leq a_{j+1} \cdot c_j + b_{j+1}$ for each $j \in [d-1]$.
- *piecewise-constant* if, for any $g \in M$, v is piecewise-linear, and moreover, $a_j = 0$ for all $j \in [d]$.
- *one-breakpoint piecewise-constant* if, for any $g \in M$, v is piecewise-constant, and furthermore, $d = 2$; see, e.g., Figure 1(b).

We will refer to one-breakpoint piecewise-constant valuations frequently later. To avoid ambiguity, we give an alternative form (that we will use more extensively): for each $i \in N$ and $g \in M$,
$$v_i(g,p) = \begin{cases} v_i(g) & \text{if } p \geq c_{i,g} \\ 0 & \text{otherwise,} \end{cases}$$
where $c_{i,g} \in (0, 1]$ denotes the threshold of agent i for good g such that she starts getting positive utility. The valuation can be succinctly represented by "$v_i(g) \cdot \mathbb{1}_{\{p \geq c_{i,g}\}}$", where $\mathbb{1}_{\{\cdot\}}$ is the indicator function that is 1 when agent i receives a fraction of good g at least the threshold $c_{i,g}$ and 0 otherwise. It is worth noting that when $c_{i,g} = 1$ for all $i \in N$ and $g \in M$, our setting becomes exactly the indivisible-goods allocation in which agents have additive utilities.

3 MAXIMIN SHARE GUARANTEE

In this section, we investigate the well-known share-based fairness notion—the *maximin share (MMS) guarantee*, which was first proposed for indivisible goods [16], and has also been investigated in settings with a mix of both divisible and indivisible goods [9, 10].

Definition 3.1 (α-MMS). Let $\Pi_n(M)$ be the set of all n-partitions of M. The *maximin share (MMS)* of $i \in N$ is
$$\text{MMS}_i(n, M) \coloneqq \max_{(P_1, \dots, P_n) \in \Pi_n(M)} \min_{j \in [n]} v_i(P_j).$$
Any partition for which the maximum is attained is called an *MMS partition* of agent i.

An allocation A is said to satisfy the *α-MMS*, for some $\alpha \in [0, 1]$, if for every agent $i \in N$, $v_i(A_i) \geq \alpha \cdot \text{MMS}_i(n, M)$.

We demonstrate agents' MMS values below.

Example 3.2. Consider an instance involving three agents $\{1, 2, 3\}$ and five goods $\{g_1, g_2, g_3, g_4, g_5\}$. Each agent has the one-breakpoint piecewise-constant valuation $0.4 \cdot \mathbb{1}_{\{p \geq 0.3\}}$ for exactly two goods, and the valuation $0.2 \cdot \mathbb{1}_{\{p \geq 0.8\}}$ for the remaining three goods.

Specifically, agent 1 (resp., 2 and 3) regards goods g_1, g_2 (resp., g_2, g_3 and g_3, g_4) in manner as the first valuation.

It can be verified that the MMS of each agent is 1. We illustrate an MMS partition of agent 1:

- each of the two goods with breakpoint at 0.3 (i.e., g_1 and g_2) is divided equally across three different bundles so that its value 0.4 is attained in all three bundles;
- the three remaining goods with breakpoint at 0.8 are allocated integrally to the three bundles respectively.

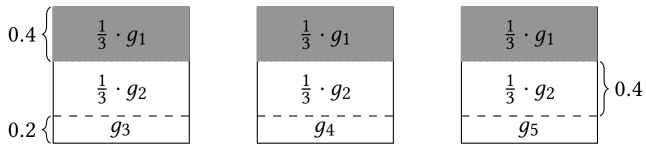

A similar MMS partition can be obtained for agents 2 and 3 by relabelling the goods according to their valuations. The following is a 0.8-MMS allocation:

- Agent 1 receives $\{g_1, 0.5 \cdot g_2, g_5\}$ and gets utility 1.
- Agent 2 receives $\{0.5 \cdot g_2, 0.5 \cdot g_3\}$ and gets utility 0.8.
- Agent 3 receives $\{0.5 \cdot g_3, g_4\}$ and gets utility 0.8.

As we mentioned in previous section, our model captures the setting of indivisible-goods allocation, in which an agent's maximin share is already NP-hard to compute and this can be seen from an immediate reduction from Partition [see, e.g., 28]. There have been *polynomial-time approximate schemes (PTASs)* designed to approximate each agent's maximin share when allocating indivisible goods [40], or a mix of heterogeneous cake and indivisible goods [10]. In our model concerning non-linear valuations over homogeneous divisible goods, we give a PTAS to approximate an agent's maximin share for constant n and one-breakpoint piecewise-constant valuations. Its detail, along with all other omitted materials, can be found in the full version of our paper [4].

Theorem 3.3. *For any constant $\varepsilon > 0$, there exists a polynomial-time algorithm that computes an MMS value with an approximation ratio of $(1-\varepsilon)^2$ for a constant number of agents with one-breakpoint piecewise-constant valuations.*

3.1 Any Number of Agents

We start with an impossibility result. Given a fair division instance, the *MMS approximation guarantee* of the instance is the maximum α such that the instance admits an α-MMS allocation. We show that the worst-case MMS approximation guarantee is at most $1/n$.

Theorem 3.4. *For n agents with one-breakpoint piecewise-constant valuations, the worst-case MMS approximation guarantee is at most $\frac{1}{n}$.*

In what follows, we provide an algorithm which always produce a $\frac{1}{2n-1}$-MMS allocation for agents with non-decreasing valuations, asymptotically matching the upper bound. The pseudocode can be found in Algorithm 1. Note that our algorithmic result works for a much broader valuation class than that being used in the impossibility result. We remark that one can adopt the *value* and *cut queries* from Caragiannis et al. [18] to access agents' non-decreasing valuations. This section is most interested in existence results, and won't discuss query or computational complexities.

Algorithm 1: $\frac{1}{2n-1}$-MMS Allocation Algorithm

Input: Agents $N = [n]$ and goods M.
Output: A $\frac{1}{2n-1}$-MMS allocation A.

```
1  foreach i ∈ N do
2      A_i ← ∅
3      Compute agent i's maximin share MMS_i(n, M).
   // Phase I: Allocate large goods.
4  while ∃g* ∈ M s.t. v_i(g*) ≥ MMS_i/(2n−1) for some i ∈ N do
5      foreach i ∈ N do
6          if v_i(g*) ≥ MMS_i/(2n−1) then
7              x_{i,g*} ← arg min_{p∈[0,1]} v_i(g*, p) ≥ MMS_i/(2n−1)
8          else
9              x_{i,g*} ← +∞
10     Let S ⊆ N denote the subset of agents such that: (i) for
       all i ∈ S and for all j ∈ N \ S, x_{i,g*} ≤ x_{j,g*}, (ii)
       ∑_{i∈S} x_{i,g*} ≤ 1, and (iii) for all j ∈ N \ S,
       ∑_{i∈S∪{j}} x_{i,g*} > 1.
11     foreach i ∈ S do A_{i,g*} ← x_{i,g*}
12     N ← N \ S, M ← M \ {g*}
   // Phase II: Bag-filling for small goods.
13 while |N| > 1 do
14     Add one good at a time to an empty bundle B until
       v_i(B) ∈ [MMS_i/(2n−1), 2·MMS_i/(2n−1)] holds for some i ∈ N.
15     A_i ← B
16     N ← N \ {i}, M ← M \ B
17 Give all remaining goods to the last agent.
18 return (A_1, A_2, ..., A_n)
```

Theorem 3.5. *For n agents with arbitrary non-decreasing valuation functions, Algorithm 1 computes a $\frac{1}{2n-1}$-MMS allocation.*

Given agents N, goods M and agents' MMS values $(\mathrm{MMS}_i)_{i\in N}$, good $g \in M$ is said to be a *large good* if there exists some agent $i \in N$ such that $v_i(g) \geq \frac{\mathrm{MMS}_i}{2n-1}$; otherwise, we will say g is a *small good*. At a high level, Algorithm 1 has two phases. Algorithm 1 starts by processing large goods in the first `while`-loop (lines 4 to 12) by allocating the goods to agents. After all large goods are allocated, the remaining goods are allocated to the remaining agents in lines 13 to 16 via *bag-filling*. We remark that although the last agent gets all remaining goods, this agent may not necessarily receive a lot more value due to her non-linear utility over each good. Note that in the following, MMS_i refers to $\mathrm{MMS}_i(n, M)$ (i.e., the maximin share computed in the original n-agent instance). Sometimes, we will refer to maximin share in reduced instances—we will make it clear.

First, we investigate instances that have only small goods.

Lemma 3.6. *Given $\langle [n], M\rangle$ in which $v_i(g) < \frac{1}{2n-1} \cdot \mathrm{MMS}_i$ for all $i \in N$ and $g \in M$, Algorithm 1 computes a $\frac{1}{2n-1}$-MMS allocation.*

Proof. Given the n-agent instance, the first `while`-condition in line 4 is evaluated as false due to the assumption in the lemma statement that for all agents $i \in N$ and goods $g \in M$, $v_i(g) < \frac{\mathrm{MMS}_i}{2n-1}$. Algorithm 1 thus executes lines 13 to 16 to process this instance. Line 14 adds one good at a time to an empty bundle B until there exists some agent $i \in N$ such that $v_i(B) \geq \frac{\text{MMS}_i}{2n-1}$. As we have assumed that agents' valuations across goods are additive, such an agent i always exists. We allocate bundle B to agent i and remove both the agent and her goods from further consideration. Clearly, agent i is satisfied with her bundle and gets a utility of at least $\frac{1}{2n-1} \cdot \text{MMS}_i$.

For all agents $j \in N$, we have $v_j(B) \leq \frac{2 \cdot \text{MMS}_j}{2n-1}$, because $v_j(g) < \frac{\text{MMS}_j}{2n-1}$ for all $g \in M$. It implies that in the reduced instance $\langle N \setminus \{i\}, M \setminus B \rangle$, for all agents $j \in N \setminus \{i\}$,

$$\text{MMS}_j(n-1, M \setminus B) \geq \left(1 - \frac{2}{2n-1}\right) \cdot \text{MMS}_j.$$

Put differently, whenever we remove one agent and their goods from consideration in Algorithm 1, the remaining agents' MMS values in the reduced instance decrease by at most $\frac{2}{2n-1}$ of their MMS values computed in the original (n-agent) instance.

By the design of the algorithm, it is clear that each removed agent in lines 13 to 16 is satisfied with their received bundle and gets a utility of at least $\frac{1}{2n-1}$ of their own maximin share (computed in the original n-agent instance). It remains to show when there is only one agent left, this last agent is satisfied with receiving all remaining goods. This can be seen from the fact that all remaining goods is worth at least $1 - \frac{2\cdot(n-1)}{2n-1} = \frac{1}{2n-1}$ of her maximin share. □

In what follows, we investigate instances having large goods.

Lemma 3.7. *Given any instance with agents $N = [n]$, goods M and $v_i(g) \geq \frac{1}{2n-1} \cdot MMS_i$ for some $i \in N$ and $g \in M$, lines 4 to 12 of Algorithm 1 computes an allocation $(A_i)_{i \in N'}$ for agents $N' \subseteq N$ and a reduced instance with agents $N'' = N \setminus N'$ and goods $M'' = M \setminus \bigcup_{i \in N'} A_i$ such that*

(i) for each $i \in N'$, $v_i(A_i) \geq \frac{MMS_i}{2n-1}$;
(ii) for each $i \in N''$ and $g \in M''$, $v_i(g) < \frac{MMS_i}{2n-1}$;
(iii) for each $i \in N''$, $MMS_i(|N''|, M'') \geq |N''| \cdot \frac{2 \cdot MMS_i}{2n-1}$.

Lemma 3.8. *Given any (reduced) instance with agents N'' and goods M'', lines 13 to 16 of Algorithm 1 computes an $\frac{1}{2n-1}$-MMS allocation for the original instance $\langle N, M \rangle$.*

The correctness of Theorem 3.5 follows from Lemmas 3.6 to 3.8.

3.2 A Small Number of Agents

We now consider cases involving a small number of agents who have one-breakpoint piecewise-constant valuations over the goods.

Our main result in this subsection is the following: For $n \leq 3$ agents with one-breakpoint piecewise-constant valuations, there always exists a $\frac{1}{n}$-MMS allocation. It is worth noting that our algorithmic results match the upper bounds stated in Theorem 3.4. For $n \leq 3$ agents, we give a tight approximation ratio for MMS.

Our algorithmic result holds trivially when $n = 1$. In what follows, we will first present an algorithm that can always output a $\frac{1}{2}$-MMS allocation for two agents in Section 3.2.1. Built up on this 2-agent-$\frac{1}{2}$-MMS algorithm, we will then present an algorithm that always outputs a $\frac{1}{3}$-MMS allocation for three agents in Section 3.2.2. To establish the algorithms, we present some auxiliary lemmas which will be proven useful. These lemmas hold for any number of agents, and connect an original instance and its reduced instance where some goods are removed. Our first two lemmas, at a high level, establish the relation between agents' MMS values in a reduced instance and those in the original instance.

Lemma 3.9. *Let $I = \langle N = [n], M \rangle$ be an instance, $M^r \subseteq M$ be a subset of goods, and $I' = \langle N, M \setminus M^r \rangle$ be an instance where goods M^r are removed from M. Then, for any agent $i \in N$,*

$$MMS_i(n, M \setminus M^r) + v_i(M^r) \geq MMS_i(n, M).$$

Lemma 3.10. *Consider any instance $\langle [n], M \rangle$ and fix $i \in [n]$. Let $M' \subseteq M$ be an integral subset of goods such that $v_i(M') \leq \alpha \cdot MMS_i(n, M)$. Then, $MMS_i(n-1, M \setminus M') \geq (1-\alpha) \cdot MMS_i(n, M)$.*

Our next lemma shows that if we remove a bundle of goods whose thresholds are greater than 0.5 (for some agent) and whose value is upper bounded, then the agent still have enough value for the remaining goods.

Lemma 3.11. *Consider an instance $\langle [n], M \rangle$ and let $i \in [n]$. If there exists $G \subseteq M$ such that $c_{i,g} > 0.5$ for all $g \in G$ and $v_i(G) < \alpha \cdot MMS_i(n, M)$ where $\alpha \geq 0$ is some non-negative real number, then $v_i(M \setminus G) \geq \frac{n-\alpha}{n} \cdot MMS_i(n, M)$.*

We now introduce the concepts of *compatible goods* and *contested goods*. Given any instance $\langle [n], M \rangle$, for each $g \in M$, good g is said to be a *compatible good* if $\sum_{i \in [n]} c_{i,g} \leq 1$; otherwise, the good g is said to be a *contested good*. In words, compatible goods are those where the thresholds for all agents sum to at most 1. Intuitively, each compatible good g can be allocated to all agents and each agent i gets a utility of $v_i(g)$. By applying Lemma 3.9 to compatible goods, we formalize below the intuition that it suffices to focus on allocating contested goods.

Given any instance $I = \langle N = [n], M \rangle$, let G^c (resp., G) be the set of all compatible (resp., contested) goods in instance I such that $M = G^c \cup G$. Consider a reduced instance $I' = \langle N, G \rangle$ with the same set of agents N and all compatible goods G^c being removed. Suppose we are given an α-MMS allocation $A' = (A'_1, A'_2, \ldots, A'_n)$ for instance I', where $\alpha \in [0, 1]$. Put differently, for all agents $i \in N$,

$$v_i(A'_i) \geq \alpha \cdot \text{MMS}_i(n, G).$$

We construct an allocation $A = (A_1, A_2, \ldots, A_n)$ as follows:

- for each $g \in G$ and $i \in N$, let $A_{i,g} = A'_{i,g}$; and
- for each $g \in G^c$, let $A_{i,g} = c_{i,g}$ for all $i \in [n-1]$ and $A_{n,g} = 1 - \sum_{i \in [n-1]} c_{i,g}$.

It can be verified that for each $i \in N$,

$$\begin{aligned} v_i(A_i) = v_i(A'_i) + \sum_{g \in G^c} v_i(g) &= v_i(A'_i) + v_i(G^c) \\ &\geq \alpha \cdot \text{MMS}_i(n, G) + v_i(G^c) \\ &\geq \alpha \cdot \left(\text{MMS}_i(n, G) + v_i(G^c)\right) \geq \alpha \cdot \text{MMS}_i(n, M), \end{aligned}$$

where the last transition is due to Lemma 3.9. As a result, allocation A is an α-MMS allocation for instance I. To summarize, we have the following lemma.

Lemma 3.12. *Let G be the set of all contested goods in $I = \langle N, M \rangle$. For any $\alpha \in [0, 1]$, if there exists an α-MMS allocation for instance $\langle N, G \rangle$, an α-MMS allocation for instance I always exists.*

Algorithm 2: 2-Agent-$\frac{1}{2}$-MMS-Alg($[2], M$)

Input: An instance $\langle N = [2], M\rangle$, where both agents have one-breakpoint piecewise-constant valuations.
Output: A $\frac{1}{2}$-MMS allocation A.

1 $A_1, A_2 \leftarrow \emptyset$
// Allocate compatible goods.
2 $G^c \leftarrow \{g \in M : c_{1,g} + c_{2,g} \leq 1\}$
3 **foreach** $g \in G^c$ **do** $A_{1,g} \leftarrow c_{1,g}; A_{2,g} \leftarrow 1 - c_{1,g}$
// Allocate the remaining contested goods.
4 $G \leftarrow M \setminus G^c$
5 **foreach** $i \in N$ **do** Compute $\text{MMS}_i \coloneqq \text{MMS}_i(2, G)$.
6 **if** $\exists g^* \in G, i \in N$ *such that* $v_i(g^*) \geq MMS_i/2$ **then**
7   **if** $v_{3-i}(g^*) \geq MMS_{3-i}/2$ **then**
8   Relabel i as the agent such that $c_{i,g^*} \leq c_{3-i,g^*}$.
9   $A_i \leftarrow A_i \cup \{g^*\}; A_{3-i} \leftarrow A_{3-i} \cup (G \setminus \{g^*\})$
10 **else**
11   **foreach** $i \in N$ **do** $G_i \leftarrow \{g \in G : c_{i,g} > 0.5\}$
12   **if** $\exists i \in N$ *such that* $v_i(G_i) \geq MMS_i$ **then**
13   Add one good integrally at a time from G_i to an empty bundle G^* until $v_i(G^*) \in [\text{MMS}_i/2, \text{MMS}_i)$.
14   Let agent $3 - i$ pick her preferred bundle out of G^* and $G \setminus G^*$, and allocate the other bundle to i.
15   **else** $A_i \leftarrow A_i \cup (G \setminus G_i); A_{3-i} \leftarrow A_{3-i} \cup G_i$
16 **return** $A = (A_1, A_2)$

3.2.1 Two Agents. We are now ready to present an algorithm which always returns a $\frac{1}{2}$-MMS allocation for two agents having one-breakpoint piecewise-constant valuations. The pseudocode can be found in Algorithm 2. In this algorithm, we first allocate all compatible goods G^c with the respective thresholds to both agents, followed by allocating the remaining contested goods G by breaking into cases based on whether there exists some (contested) good $g \in G$ and agent i such that $v_i(g) > \text{MMS}_i(2, G)/2$.

Theorem 3.13. *For $n = 2$ agents and one-breakpoint piecewise-constant valuations, Algorithm 2 returns a $\frac{1}{2}$-MMS allocation.*

Proof. Algorithm 2 starts by allocating all compatible goods G^c in line 3. By Lemma 3.12, it suffices to show that Algorithm 2 finds a 1/2-MMS allocation in the reduced instance $I' = \langle N, G = M \setminus G^c\rangle$ which contains only contested goods. For notational convenience, in the remainder of this proof, let MMS_i denote $\text{MMS}_i(N, G)$, i.e., the maximin share in the reduced instance I'.

We distinguish cases based on whether there exists good $g^* \in G$ and $i \in N$ such that $v_i(g^*) \geq \text{MMS}_i/2$. Suppose that the `if`-statement in line 6 is evaluated as true, i.e., such a good g^* and agent i exists. We first consider the case where $v_{3-i}(g^*) \geq \frac{\text{MMS}_{3-i}}{2}$, and assume without loss of generality that $c_{i,g^*} \leq c_{3-i,g^*}$ (relabel the agents if needed). Since good g^* is contested, $c_{3-i,g^*} > 0.5$, which implies that $v_{3-i}(G \setminus \{g^*\}) \geq \text{MMS}_{3-i}$. We now consider the other case where $v_{3-i}(g^*) < \frac{\text{MMS}_{3-i}}{2}$. It can be verified that

$$v_{3-i}(G \setminus \{g^*\}) = v_{3-i}(G) - v_{3-i}(g^*) \geq \text{MMS}_{3-i} - \frac{\text{MMS}_{3-i}}{2} = \frac{\text{MMS}_{3-i}}{2}.$$

In either case, allocating good g^* to agent i and all remaining goods $G \setminus \{g^*\}$ to the other agent in line 9 gives a $\frac{1}{2}$-MMS allocation for instance I', as desired.

Finally, we consider the scenario where for all $i \in N$ and $g \in G$, $v_i(g) < \text{MMS}_i/2$. For each $i \in N$, let $G_i \subseteq G$ consist of goods where agent i has thresholds larger than 0.5. There are two cases:

(i) $v_i(G_i) \geq \text{MMS}_i$ for some agent $i \in N$, and
(ii) $v_i(G_i) < \text{MMS}_i$ for both agents.

In the first case, because all goods $g \in G$ follow that $v_i(g) < \text{MMS}_i/2$, we can keep adding goods $g \in G_i$ to find a subset $G^* \subseteq G_i$ such that $v_i(G^*) \in \left[\frac{\text{MMS}_i}{2}, MMS_i\right)$. By Lemma 3.11, $v_i(G \setminus G^*) \geq \frac{1}{2} \cdot \text{MMS}_i$. Therefore, allocating the preferred bundle between G^* and $G \setminus G^*$ to agent $3 - i$ and the other bundle to agent i is a 1/2-MMS allocation for instance I'. In the second case, we have $v_i(G_i) < \text{MMS}_i$ for both agents $i \in N$. By Lemma 3.11, this implies $v_i(G \setminus G_i) \geq \text{MMS}_i/2$, where $G \setminus G_i$ is the set of goods such that agent i has thresholds at most 0.5. However, since all goods in G are contested, the sets $G \setminus G_i$ and $G \setminus G_{3-i}$ are mutually exclusive. This means that allocating $G \setminus G_i$ to agent i and all remaining goods to the other agent satisfies 1/2-MMS in instance I'.

In conclusion, Algorithm 2 finds a 1/2-MMS allocation in the reduced instance I'. Moreover, due to Lemma 3.12, Algorithm 2 finds a $\frac{1}{2}$-MMS allocation for the two agents in the original instance. □

3.2.2 Three Agents. Built upon our 2-agent-$\frac{1}{2}$-MMS allocation algorithm, we proceed to present an algorithm that finds a 1/3-MMS allocation for three agents with one-breakpoint piecewise-constant valuations. The pseudocode is given as Algorithm 3.

Algorithm 3 takes as input a set of agents $N = [3]$ and a set of goods $M = G^c \cup G$, where G^c denotes the set of compatible goods and G denotes the set of contested goods. Following Lemma 3.12, Algorithm 3 starts by allocating all compatible goods G^c with the respective thresholds to the three agents. It remains to show that our algorithm finds a $\frac{1}{3}$-MMS allocation in the reduced instance with agents N and contested goods G. Note also that in the following, for each agent $i \in N$, $\text{MMS}_i \coloneqq \text{MMS}_i(3, G)$.

Based on whether there exists some agent i and good $g \in G$ such that $v_i(g) \geq \text{MMS}_i/3$, our algorithm can be naturally broken into two components, with lines 5 to 10 processing instances with some large good and lines 11 to 31 processing remaining instances with only small goods. Comparing with the 2-agent case presented previously, our Algorithm 3 is more intricate in the sense that besides agents' valuations for a single good or a set of goods, we also need to reason about the thresholds by which agents start having positive utilities for a given good. Naturally, this is more complicated since each good can be allocated to any of the 7 non-empty subsets of agents $\{1, 2, 3\}$ based on the instance.

Theorem 3.14. *For $n = 3$ agents and one-breakpoint piecewise-constant valuations, Algorithm 3 returns a $\frac{1}{3}$-MMS allocation.*

Before showing Algorithm 3 indeed finds a 1/3-MMS allocation, we prove a few useful results. The following two lemmas allow us to form an allocation when there is some "large" good(s) in the given instance, i.e., an agent i values a good g at least $\text{MMS}_i/3$.

Lemma 3.15. *Let $i \in N$. If there exists $g^* \in G$ such that $c_{i,g^*} > 1/3$, then $v_i(G \setminus g^*) \geq MMS_i$.*

Algorithm 3: 3-Agent-$\frac{1}{3}$-MMS Allocation

Input: Instance $I = \langle N = [3], M = G^c \cup G \rangle$, where agents have one-breakpont piecewise-constant valuations.
Output: A 1/3-MMS allocation A.

1 $A_1, A_2, A_3 \leftarrow \emptyset$
2 **foreach** $g \in G^c$ **do** // Allocate compatible goods.
3 $\quad A_{i,g} \leftarrow c_{i,g} \forall i \in [n-1]$; $A_{n,g} \leftarrow 1 - \sum_{i \in [n-1]} c_{i,g}$
4 **foreach** $i \in N$ **do** Compute $\text{MMS}_i \coloneqq \text{MMS}_i(3, G)$.
// Process instances with some large goods.
5 **if** $\exists g^* \in G$ *and* $a \in N$ *such that* $v_a(g^*) \geq \frac{MMS_a}{3}$ **then**
6 $\quad$ Let $S \subsetneq N$ be the set of agents such that $v_i(g^*) \geq \frac{\text{MMS}_i}{3}$ for all $i \in S$ and $\sum_{j \in S} c_{j,g^*} \leq 1$, breaking ties in favour of larger $|S|$ and then smaller $\sum_{j \in S} c_{j,g^*}$.
7 $\quad$ **foreach** $i \in S$ **do** $A_{i,g^*} \leftarrow c_{i,g^*}$
8 $\quad$ **if** $|S| = 2$ **then** $A_k \leftarrow A_k \cup G \setminus \{g^*\}$, where $k \notin S$.
9 $\quad$ **else** Call Algorithm 2 on sub-instance $\langle N \setminus S, G \setminus \{g^*\} \rangle$
10 $\quad$ **return** $A = (A_1, A_2, A_3)$
// Allocate small goods. Note, $\forall i \in N, v_i(g) < \frac{\text{MMS}_i}{3}$.
11 **foreach** $i \in N$ **do** $G_i \leftarrow \{g \in G \colon c_{i,g} > 0.5\}$; $G'_i \leftarrow G \setminus G_i$
12 **if** $\exists i \in N$ *such that* $v_i(G_i) \geq MMS_i$ **then**
13 $\quad$ Partition goods G into three integral bundles G^*_1, G^*_2, G^*_3 such that $v_i(G^*_1), v_i(G^*_2), v_i(G^*_3) \geq \text{MMS}_i/3$.
14 $\quad$ **if** $\exists j \neq i$ *values two distinct bundles at least* $\frac{MMS_j}{3}$ **then**
15 $\quad\quad$ **return** *the preferred bundle to* $k \in N \setminus \{i, j\}$*, then the preferred bundle to* j*, and finally the last bundle to* i.
16 $\quad$ **else**
17 $\quad\quad$ Let G^*_a be the bundle s.t. $v_j(G^*_a) < \frac{\text{MMS}_j}{3} \forall j \neq i$.
18 $\quad\quad$ $A_i \leftarrow G^*_a$
19 $\quad\quad$ **return** 2-Agent-$\frac{1}{2}$-MMS-Alg$([3] \setminus \{i\}, G \setminus G^*_a)$
20 $\forall i, j \in N, G'_{\{i,j\}} \leftarrow \{g \in G \colon c_{i,g} \leq 0.5 \wedge c_{j,g} \leq 0.5\}$
21 **if** $\exists G'_{\{i,j\}}$ *s.t.* $v_i(G'_{\{i,j\}}) \geq \frac{MMS_i}{3}$ *and* $v_j(G'_{\{i,j\}}) \geq \frac{MMS_j}{3}$ **then**
22 $\quad$ Relabel agents. Let $G^* \subseteq G'_{\{i,j\}}$ be s.t. $v_i(G^*) \geq \frac{\text{MMS}_i}{3}$, $v_j(G^*) \geq \frac{\text{MMS}_j}{3}$, and $v_k(G^*) < \frac{2 \cdot \text{MMS}_k}{3}$.
23 $\quad$ $\forall g \in G^*$, $A_{i,g} \leftarrow 0.5$ and $A_{j,g} \leftarrow 0.5$; $A_k \leftarrow G \setminus G^*$
24 **else**
25 $\quad$ Pick any $G'_{\{i,k\}}$ such that $v_k(G'_{\{i,k\}}) < \text{MMS}_k/3$.
26 $\quad$ Iteratively move $g \in G'_i$ to G^*_i until $v_i(G^*_i) \geq \text{MMS}_i/3$.
27 $\quad$ $\forall g \in G^*_i$, $A_{i,g} \leftarrow 0.5$ and $A_{j,g} \leftarrow 0.5$
28 $\quad$ Iteratively move $g \in G'_j$ to G^*_j until $v_j(A_j \cup G^*_j) \geq \frac{\text{MMS}_j}{3}$.
29 $\quad$ $\forall g \in G^*_j$, $A_{j,g} \leftarrow 0.5$ and $A_{k,g} \leftarrow 0.5$
30 $\quad$ Allocate everything else to k.
31 **return** $A = (A_1, A_2, A_3)$

Lemma 3.16. *Let $i \in N$. If there exists $g^* \in G$ such that $c_{i,g^*} > 1/2$, then $MMS_i(2, G \setminus \{g^*\}) \geq MMS_i(3, G)$.*

We now establish the correctness of lines 5 to 10 of Algorithm 3.

Lemma 3.17. *Given agents $N = [3]$ and (contested) goods G, if there exists some $a \in N$ and some $g^* \in G$ such that $v_a(g^*) \geq \frac{MMS_a}{3}$, lines 5 to 10 of Algorithm 3 returns a $\frac{1}{3}$-MMS allocation.*

Proof. Let $S \subseteq N$ be the set of agents such that $v_i(g^*) \geq \frac{\text{MMS}_i}{3}$ for all $i \in S$ and $\sum_{j \in S} c_{j,g^*} \leq 1$, breaking ties in favour of larger $|S|$ and then smaller $\sum_{j \in S} c_{j,g^*}$. In other words, we can allocate a fraction c_{i,g^*} of good g^* to each agent $i \in S$ and the agent receives a utility of at least $\frac{\text{MMS}_i}{3}$. Recall that we are allocating contested goods, meaning that good g^* can be allocated to at most two agents. Clearly, $|S| \geq 1$, as the `if`-condition in line 5 is evaluated as true.

Below, we distinguish two cases based on whether $|S| = 2$ or $|S| = 1$. When $|S| = 2$, since we break ties in favour of smaller $\sum_{j \in S} c_{j,g^*}$, the agent $k \notin S$ must either have $c_{k,g^*} > 1/3$, or $v_k(g^*) < \frac{\text{MMS}_k}{3}$. In the first case, Lemma 3.15 implies that $v_k(G \setminus \{g^*\}) \geq \text{MMS}_k$; and in the second case, we have $v_k(G \setminus \{g^*\}) > \text{MMS}_k - \frac{\text{MMS}_k}{3} = \frac{2\text{MMS}_k}{3}$. This implies that allocating $G \setminus \{g^*\}$ would give a 1/3-MMS allocation.

When $|S| = 1$, let j and k be the agents not in S. For any agent $a \in \{j, k\}$, we must have either $c_{a,g^*} > 1/2$ or $v_a(g^*) < \frac{\text{MMS}_a}{3}$. In the first case, Lemma 3.16 implies that under the reduced instance $I' = \langle \{j, k\}, G \setminus \{g^*\} \rangle$, $\text{MMS}_a(2, G \setminus \{g^*\}) \geq \text{MMS}_a(3, G)$; in the second case, since $v_a(g^*) < \frac{\text{MMS}_a}{3}$, Lemma 3.10 implies that $\text{MMS}_a(2, G \setminus \{g^*\}) \geq \frac{2}{3}\text{MMS}_a(3, G)$. Therefore in both cases, a 1/2-MMS allocation under instance I' must correspond to an at least 1/3-MMS allocation under the instance with all three agents and goods G. As a result, allocating $G \setminus g^*$ to j and k using Algorithm 2 would give a 1/3-MMS allocation. The conclusion follows. ◻

In the following, we are concerned with 3-agent instances with only "small" goods. Formally, for all agents $i \in N$ and all goods $g \in G$, we have $v_i(g) < \text{MMS}_i/3$. For each agent $i \in N$, let G_i (resp., G'_i) denote the set of goods g such that $c_{i,g} > 0.5$ (resp., $c_{i,g} \leq 0.5$).

First, we observe that when $v_i(G_i) \geq \text{MMS}_i$ for some agent i, then we can partition goods G into three integral bundles such that all three bundles are worth at least $\text{MMS}_i/3$.

Lemma 3.18. *Let $i \in [3]$ and suppose that $v_i(g) < MMS_i/3$ for all goods $g \in G$. If $v_i(G_i) \geq MMS_i$, then G can be partitioned into three integral bundles each of which is worth at least $MMS_i/3$.*

We now establish the correctness of lines 11 to 31 of Algorithm 3.

Lemma 3.19. *Given agents $N = [3]$ and (contested) goods G, if for all goods $g \in G$ and agents $i \in N$, $v_i(g) < MMS_i/3$, lines 11 to 31 Algorithm 3 returns a 1/3-MMS allocation.*

The proof of Theorem 3.14 follows from Lemmas 3.17 and 3.19.

4 EF AND PARETO OPTIMAL ALLOCATIONS

In this section, we focus on *envy-freeness*. To set the stage, we first define several relevant concepts.

Definition 4.1 (EF). An allocation A is said to be *envy-free (EF)* if, for every pair of agents $i, j \in N$, $v_i(A_i) \geq v_i(A_j)$.

Definition 4.2 (PO). Given an allocation A, another allocation $A' = (A'_i)_{i \in N}$ *Pareto dominates* A if $v_i(A'_i) \geq v_i(A_i)$ for all $i \in N$ and $v_j(A'_j) > v_j(A_j)$ for some $j \in N$. An allocation is said to be *Pareto optimal (PO)* if no other allocation Pareto dominates it.

Definition 4.3. An allocation A is said to be *EF-constrained-PO* if A is EF and no other envy-free allocation A' exists such that $v_i(A'_i) \geq v_i(A_i)$ for all $i \in N$ and $v_i(A'_j) > v_i(A_j)$ for some $j \in N$.

A PO allocation is guaranteed to exist for divisible goods, regardless of the utility function. While an EF allocation of multiple divisible goods also always exists in our context since each item can be equally divided among agents, such an allocation is often not PO. This raises the issue of achieving both EF and PO simultaneously, and for what instances does an EF and PO allocation always exists. Under the indivisible goods setting, it is well established that determining whether an EF allocation exists for multiple goods is NP-hard, and the problem of finding an allocation that satisfies both EF and PO is even more computationally complex [21]. In our model, we show even when there are 3 goods, determining whether an EF and PO allocation exists is NP-hard.

Theorem 4.4. *The existence of an EF and PO allocation is NP-hard for $m \geq 3$ and one-breakpoint piecewise-constant valuations.*

In the reminder of this section, we examine the single-good case, which has been explored in related but different contexts [17, 22].

4.1 Single Good Allocation

We present an algorithm that efficiently computes an EF-constrained PO allocation for a single good, which can then be used to determine whether an EF and PO allocation exists. For ease of notation and since there is only $m = 1$ good concerned, we denote $v_i(p) = v_i(g, p)$ as a function of allocated proportion in this section only. Similarly, we denote $A_i = A_{i,g}$ as the proportion allocated to agent i under A.

Let $\underline{s_i}(b)$ represent the minimum proportion s such that $v_i(s) = v_i(b)$; let $B = \{b_1, \ldots, b_k\}$ denotes all breakpoints across all agents. The proposed algorithm operates in three stages to ensure EF-constrained PO in the allocation:

(1) Identify the largest breakpoint b_j such that $\sum_{i \in N} \underline{s_i}(b_j) \leq 1$. Allocate to each agent i a portion of the good with $A_i = \underline{s_i}(b_j)$.
(2) Let ϵ denote a small remainder of the good, and define L as the set of agents such that, for all $i \in L$, $\underline{s_i}(b_j + \epsilon) > b_j$. Adjust the allocation for each agent $i \in L$ as:

$$A_i = \min\left(b_{j+1} - \epsilon, b_j + \frac{1 - \sum_{k \in N} \underline{s_k}(b_j)}{|L|}\right),$$

while keeping the allocation for other agents unchanged.
(3) Allocate the remaining portion of the good in such a way that no agent gets b_{j+1} proportion of the good.

Since there are polynomially many breakpoints, the algorithm runs in polynomial time. We now show that the algorithm returns an EF-constrained PO allocation.

Lemma 4.5. *The allocation found by Algorithm 4 is EF.*

Lemma 4.6. *No EF allocation Pareto dominates the allocation found by Algorithm 4.*

Combining Lemmas 4.5 and 4.6, we show that the allocation found by Algorithm 4 is EF-constrained PO, and thus the following theorem follows.

Theorem 4.7. *An EF-constrained PO allocation always exist and can be found in polynomial time for a single divisible good with piecewise-linear utility functions.*

Algorithm 4: EF-constrained PO allocation

Input: An instance $I = (N, M)$.
Output: An EF-constrained PO allocation A

$j \leftarrow 0$
while $\sum_{i \in N} \underline{s_i}(b_{j+1}) \leq 1$ **do** $j = j + 1$
foreach $i \in N$ **do** $A_i \leftarrow \underline{s_i}(b_j)$
$L \leftarrow \emptyset$
foreach $i \in N$ **do**
  if $\underline{s_i}(b_j + \epsilon) > b_j$ **then** $L \leftarrow L \cup i$
foreach $i \in L$ **do** $A_i \leftarrow \min\left(b_{j+1} - \epsilon, b_j + \frac{1 - \sum_{i \in N} \underline{s_i}(b_j)}{|L|}\right)$
while $\sum_{i \in N} A_i < 1$ **do**
  find i such that $A_i < b_{j+1} - \epsilon$
  $A_i \leftarrow \min(b_{j+1} - \epsilon, 1 + A_i - \sum_{j \in N} A_j)$
return A

Further, one can easily show that any EF and PO allocation must itself be EF-constrained PO and PO. However, it turns out a stronger result holds for $m = 1$ good: an EF and PO allocation exists if and only if the EF-constrained PO allocation found by Algorithm 4 is also PO. This is proved in the following theorem.

Theorem 4.8. *The existence of an EF and PO allocation can be checked in polynomial time for a single divisible good with piecewise-linear utility functions.*

5 DISCUSSION

In this paper, we have studied a fair division model where a set of homogeneous divisible goods are allocated among agents who have non-linear valuations in the sense that each agent's value depends only on the amount of each good she receive. We focus on the fair-share-based notion of MMS and the comparison-based fairness notion of envy-freeness. In more detail, we give tight or asymptotically tight approximation to MMS in various cases. For envy-freeness, we explore computational complexity of checking the existence of an envy-free and PO allocation. Since one-breakpoint piecewise-constant valuations capture a wide range of real-world applications and can be easily elicited from agents in practice, it would be intriguing to fully understand whether we can guarantee $\frac{1}{n}$-MMS for n agents—we conjecture that the answer is affirmative, and the computational complexity of checking the existence of an EF and PO allocation for exactly two goods.

In future research, since complementarity and substitutability are common in practice, it would be interesting to consider compact but more expressive valuations that relax our assumption of values across different goods being additive. As one possibility, we could let each agent has a utility function $v((p_j)_{j \in [m]})$ that specifies the value for getting fraction p_1 of item g_1, p_2 of item g_2, and so on.

ACKNOWLEDGMENTS

This work was partially supported by the NSF-CSIRO grant on "Fair Sequential Collective Decision-Making" (Grant No. RG230833) and the ARC Laureate Project FL200100204 on "Trustworthy AI."

REFERENCES

[1] Hannaneh Akrami and Jugal Garg. 2024. Breaking the 3/4 Barrier for Approximate Maximin Share. In *Proceedings of the 35th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 74–91.

[2] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. 2023. Fair Division of Indivisible Goods: Recent Progress and Open Questions. *Artificial Intelligence* 322 (2023), 103965.

[3] Georgios Amanatidis, George Christodoulou, John Fearnley, Evangelos Markakis, Christos-Alexandros Psomas, and Eftychia Vakaliou. 2018. An Improved Envy-Free Cake Cutting Protocol for Four Agents. In *Proceedings of 11th International Symposium on Algorithmic Game Theory (SAGT)*. 87–99.

[4] Haris Aziz, Zixu He, Xinhang Lu, and Kaiyang Zhou. 2025. Fair Allocation of Divisible Goods under Non-Linear Valuations. *CoRR* (2025).

[5] Haris Aziz and Simon Mackenzie. 2016. A Discrete and Bounded Envy-Free Cake Cutting Protocol for Any Number of Agents. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 416–427.

[6] Haris Aziz and Simon Mackenzie. 2016. A Discrete and Bounded Envy-Free Cake Cutting Protocol for Four Agents. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*. 454–464.

[7] Xiaohui Bei, Ning Chen, Xia Hua, Biaoshuai Tao, and Endong Yang. 2012. Optimal Proportional Cake Cutting with Connected Pieces. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*. 1263–1269.

[8] Xiaohui Bei, Zihao Li, Jinyan Liu, Shengxin Liu, and Xinhang Lu. 2021. Fair Division of Mixed Divisible and Indivisible Goods. *Artificial Intelligence* 293 (2021), 103436.

[9] Xiaohui Bei, Shengxin Liu, and Xinhang Lu. 2023. Fair Division with Subjective Divisibility. In *Proceedings of the 19th Conference on Web and Internet Economics (WINE)*. 677.

[10] Xiaohui Bei, Shengxin Liu, Xinhang Lu, and Hongao Wang. 2021. Maximin Fairness with Mixed Divisible and Indivisible Goods. *Autonomous Agents and Multi-Agent Systems* 35, 2 (2021), 34:1–34:21.

[11] Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. 2022. Almost Full EFX Exists for Four Agents. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. 4826–4833.

[12] Steven J. Brams, Michal Feldman, John K. Lai, Jamie Morgenstern, and Ariel D. Procaccia. 2012. On Maxsum Fair Cake Divisions. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*. 1285–1291.

[13] Steven J. Brams and Peter C. Fishburn. 2000. Fair Division of Indivisible Items Between Two People with Identical Preferences: Envy-freeness, Pareto-optimality, and Equity. *Social Choice and Welfare* 17, 2 (2000), 247–267.

[14] Steven J. Brams, Marc Kilgour, and Christian Klamler. 2014. Two-Person Fair Division of Indivisible Items: An Efficient, Envy-Free Algorithm. *Notices of the AMS* 61, 2 (2014), 130–141.

[15] Steven J. Brams and Alan D. Taylor. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.

[16] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.

[17] Jan Buermann, Enrico H. Gerding, and Baharak Rastegari. 2020. Fair Allocation of Resources with Uncertain Availability. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 204–212.

[18] Ioannis Caragiannis, Vasilis Gkatzelis, Alexandros Psomas, and Daniel Schoepflin. 2022. Beyond Cake Cutting: Allocating Homogeneous Divisible Goods. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 208–216.

[19] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. 2024. EFX Exists for Three Agents. *J. ACM* 71, 1 (2024), 4:1–4:27.

[20] Yuga J. Cohler, John K. Lai, David C. Parkes, and Ariel D. Procaccia. 2011. Optimal Envy-Free Cake Cutting. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*. 626–631.

[21] Bart de Keijzer, Sylvain Bouveret, Tomas Klos, and Yingqian Zhang. 2009. On the Complexity of Efficiency and Envy-Freeness in Fair Division of Indivisible Goods with Additive Preferences. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT)*. 98–110.

[22] Uriel Feige, Ariel Sapir, and Laliv Tauber. 2021. A Tight Negative Example for MMS Fair Allocations. In *Proceedings of the 17th Conference on Web and Internet Economics (WINE)*. 355–372.

[23] Uriel Feige and Moshe Tennenholtz. 2014. On Fair Division of a Homogeneous Good. *Games and Economic Behavior* 87 (2014), 305–321.

[24] Duncan Karl Foley. 1967. Resource Allocation and the Public Sector. *Yale Economics Essays* 7, 1 (1967), 45–98.

[25] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. 2011. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)*. 323–336.

[26] Alexandros Hollender and Aviad Rubinstein. 2023. Envy-Free Cake-Cutting for Four Agents. In *Proceedings of the 64th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. 113–122.

[27] D. Marc Kilgour and Rudolf Vetschera. 2018. Two-player Fair Division of Indivisible Items: Comparison of Algorithms. *European Journal of Operational Research* 271, 2 (2018), 620–631.

[28] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2018. Fair Enough: Guaranteeing Approximate Maximin Shares. *J. ACM* 65, 2 (2018), 8:1–8:27.

[29] Claudia Lindner and Jörg Rothe. 2024. Cake-Cutting: Fair Division of Divisible Goods. In *Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division* (2nd ed.), Jörg Rothe (Ed.). Springer Cham, Chapter 8, 507–603.

[30] Shengxin Liu, Xinhang Lu, Mashbat Suzuki, and Toby Walsh. 2024. Mixed Fair Division: A Survey. *Journal of Artificial Intelligence Research* 80 (2024), 1373–1406.

[31] Trung Thanh Nguyen and Jörg Rothe. 2023. Complexity Results and Exact Algorithms for Fair Division of Indivisible Items: A Survey. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*. 6732–6740.

[32] David C. Parkes, Ariel D. Procaccia, and Nisarg Shah. 2015. Beyond Dominant Resource Fairness: Extensions, Limitations, and Indivisibilities. *ACM Transactions on Economics and Computation* 3, 1 (2015), 3:1–3:22.

[33] Ariel D. Procaccia. 2016. Cake Cutting Algorithms. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, Chapter 13, 311–329.

[34] Johannes H. Reijnierse and Jos A. M. Potters. 1998. On Finding an Envy-free Pareto-optimal Division. *Mathematical Programming* 83, 1 (1998), 291–311.

[35] Jack Robertson and William Webb. 1998. *Cake-Cutting Algorithm: Be Fair If You Can*. A K Peters/CRC Press.

[36] Hugo Steinhaus. 1949. Sur La Division Pragmatique. *Econometrica* 17 (1949), 315–319.

[37] Francis Edward Su. 1999. Rental Harmony: Sperner's Lemma in Fair Division. *The American Mathematical Monthly* 106, 10 (1999), 930–942.

[38] Warut Suksompong. 2021. Constraints in Fair Division. *ACM SIGecom Exchanges* 19, 2 (2021), 46–61.

[39] Warut Suksompong. 2025. Weighted Fair Division of Indivisible Items: A Review. *Inform. Process. Lett.* 187 (2025), 106519.

[40] Gerhard J. Woeginger. 1997. A Polynomial-time Approximation Scheme for Maximizing the Minimum Machine Completion Time. *Operations Research Letters* 20, 4 (1997), 149–154.