

Ordinal Maximin Share Approximation for Goods

Hadi Hosseini

Pennsylvania State University, University Park

HADI@PSU.EDU

Andrew Searns

Johns Hopkins University Applied Physics Laboratory

ANDREW.SEARNS@JHUAPL.EDU

Erel Segal-Halevi

Ariel University, Ariel

ERELSGL@GMAIL.COM

Abstract

In fair division of indivisible goods, ℓ -out-of- d maximin share (MMS) is the value that an agent can guarantee by partitioning the goods into d bundles and choosing the ℓ least preferred bundles. Most existing works aim to guarantee to all agents a constant fraction of their 1-out-of- n MMS. But this guarantee is sensitive to small perturbation in agents' cardinal valuations. We consider a more robust approximation notion, which depends only on the agents' *ordinal* rankings of bundles. We prove the existence of ℓ -out-of- $\lfloor(\ell + \frac{1}{2})n\rfloor$ MMS allocations of goods for any integer $\ell \geq 1$, and present a polynomial-time algorithm that finds a 1-out-of- $\lceil\frac{3n}{2}\rceil$ MMS allocation when $\ell = 1$. We further develop an algorithm that provides a weaker ordinal approximation to MMS for any $\ell > 1$.

1. Introduction

Fair division is the study of how to distribute a set of items among a set of agents in a fair manner. Achieving fairness is particularly challenging when items are *indivisible*. Computational and conceptual challenges have motivated researchers and practitioners to develop a variety of fairness concepts that are applicable to a large number of allocation problems.¹ One of the most common fairness concepts, proposed by Budish (2011), is Maximin Share (MMS), that aims to give each agent a bundle that is valued at a certain threshold. The MMS threshold, also known as 1-out-of- d MMS, generalizes the guarantee of the *cut-and-choose* protocol. It is the value that an agent can secure by partitioning the items into d bundles, assuming it will receive the least preferred bundle. The MMS value depends on the number of partitions, d . When all items are goods (i.e., have non-negative values), the 1-out-of- d MMS threshold is (weakly) monotonically decreasing as the number of partitions (d) increases.

When allocating goods among n agents, a natural desirable threshold is satisfying 1-out-of- n MMS for all agents. Unfortunately, while this value can be guaranteed for $n = 2$ agents through the cut-and-choose protocol, a 1-out-of- n MMS allocation of goods may not exist in general for $n \geq 3$ (Procaccia & Wang, 2014; Kurokawa, Procaccia, & Wang, 2018). These negative results have given rise to *multiplicative approximations*, wherein each agent is guaranteed at least a constant fraction of its 1-out-of- n MMS. While there have been

1. See Bouveret, Chevaleyre, and Maudet (2016), Lang and Rothe (2016), Markakis (2017), Moulin (2019) for detailed surveys and discussions.

many attempts in developing algorithms that improve the bound to close to 1, the best currently known fraction is $\frac{3}{4} + \frac{1}{12n}$ (Garg & Taki, 2020).

Despite numerous studies devoted to their existence and computation, there is a conceptual and practical problem with the multiplicative approximations of MMS: they are very sensitive to agents' precise cardinal valuations. To illustrate, suppose $n = 3$ and there are four goods g_1, g_2, g_3, g_4 that Alice values at 30, 39, 40, 41 respectively. Her 1-out-of-3 MMS is 40, and thus a $\frac{3}{4}$ fraction guarantee can be satisfied by giving her the bundle $\{g_1\}$ or a bundle with a higher value. But if her valuation of good g_3 changes slightly to $40 + \varepsilon$ (for any $\varepsilon > 0$), then $\frac{3}{4}$ of her 1-out-of-3 MMS is larger than 30, the bundle $\{g_1\}$ is no longer acceptable for her. Thus, the acceptability of a bundle (in this example $\{g_1\}$) might be affected by an arbitrarily small perturbation in the value of an *irrelevant* good (i.e. g_3).

In the microeconomics literature, it is common to measure agents' preferences as *ordinal rankings* of the bundles; even when utility functions are used, it is understood that they only represent rankings. From this viewpoint, the set of acceptable bundles should only depend on the ranking of the bundles, and should not be affected by changes in valuations that—similar to the ε change in the value of g_3 —do not affect this ranking. According to this principle, Budish (2011) suggested the 1-out-of- $(n + 1)$ MMS as a relaxation of the 1-out-of- n MMS. In the above example, 1-out-of-4 MMS fairness can be satisfied by giving Alice $\{g_1\}$ or a better bundle; small inaccuracies or noise in the valuations do not change the set of acceptable bundles. Hence, this notion provides a more robust approach in evaluating fairness of allocations.

To date, it is not known if 1-out-of- $(n + 1)$ MMS allocations are guaranteed to exist. We aim to find allocations of goods that guarantee 1-out-of- d MMS for some integer $d > n$. A 1-out-of- d MMS allocation guarantees to each agent a bundle that is at least as good as the worst bundle in the best d -partition.

The aforementioned guarantee can be naturally generalized to ℓ -out-of- d MMS (Babaioff, Nisan, & Talgam-Cohen, 2021), that guarantees to each agent the value obtained by partitioning the goods into d bundles and selecting the ℓ least-valuable ones. Therefore, we further investigate the ℓ -out-of- d MMS generalization that allows us to improve the fairness thresholds. The notion of ℓ -out-of- d MMS fairness is *robust* in the sense that, a fair allocation remains fair even when each agent's utility function goes through an arbitrary monotonically-increasing transformation. Given these notions, we ask the following questions:

In the allocation of indivisible goods, (a) For what combinations of integers ℓ and d , can ℓ -out-of- d MMS allocations be guaranteed? and (b) For what integers ℓ and d can ℓ -out-of- d MMS allocations be computed in polynomial time?

1.1 Our Contributions

We investigate the existence and computation of ordinal MMS approximations and make several contributions.

In **Section 4**, we prove the existence of ℓ -out-of- d MMS allocation of goods when $d \geq \lfloor (\ell + \frac{1}{2})n \rfloor$ (Theorem 1). In particular, 1-out-of- $\lfloor 3n/2 \rfloor$ MMS, 2-out-of- $\lfloor 5n/2 \rfloor$ MMS, 3-out-of- $\lfloor 7n/2 \rfloor$ MMS, and so on, are all guaranteed to exist. This finding generalizes the previously known existence result of 1-out-of- $\lceil 3n/2 \rceil$ MMS (Hosseini & Searns, 2021).

The proof uses an algorithm which, given lower bounds on the ℓ -out-of- d MMS values of the agents, returns an ℓ -out-of- d MMS allocation. The algorithm runs in polynomial time given the agents' lower bounds. However, computing the exact ℓ -out-of- d MMS values is NP-hard. In the following sections we propose two solutions to this issue.

In **Section 5**, we present polynomial-time algorithms that find an ℓ -out-of- $(d + o(n))$ MMS-fair allocation, where $d = (\ell + \frac{1}{2})n$. Specifically, for $\ell = 1$, we present a polynomial-time algorithm for finding a 1-out-of- $\lceil 3n/2 \rceil$ MMS allocation (Theorem 2); this matches the existence result for 1-out-of- $\lceil 3n/2 \rceil$ MMS up to an additive gap of at most 1. For $\ell > 1$, we present a different polynomial-time algorithm for finding a 1-out-of- $\lceil (\ell + \frac{1}{2})n + O(n^{2/3}) \rceil$ MMS allocation (Theorem 3).

In **Appendix A**, we conduct simulations with valuations generated randomly from various distributions. For several values of ℓ , we compute a lower bound on the ℓ -out-of- $\lfloor (\ell + \frac{1}{2})n \rfloor$ MMS guarantee using a simple greedy algorithm. We compare this lower bound to an upper bound on the $(\frac{3}{4} + \frac{1}{12n})$ -fraction MMS guarantee, which is currently the best known worst-case multiplicative MMS approximation.² We find that, for any $\ell \geq 2$, when the number of goods is at least $\approx 20n$, the lower bound on the ordinal approximation is better than the upper bound on the multiplicative approximation. This implies that, in practice, the algorithm of Section 4 can be used with these lower bounds to attain an allocation in which each agent receives a value that is significantly better than the theoretical guarantees.

1.2 Techniques

At first glance, it would seem that the techniques used to attain $2/3$ approximation of MMS should also work for achieving 1-out-of- $\lceil 3n/2 \rceil$ MMS allocations, since both guarantees approximate the same value, namely, the $\frac{2}{3}$ approximation of the “proportional share” ($\frac{1}{n}$ of the total value of all goods). In **Appendix B** we present an example showing that this is not the case, and thus, achieving ordinal MMS approximations requires new techniques. In this section, we briefly describe the techniques that we utilize to achieve ordinal approximations of MMS.

Lone Divider. To achieve the existence result for any $\ell \geq 1$, we use a variant of the *Lone Divider* algorithm, which was first presented by Kuhn (1967) for finding a proportional allocation of a divisible good (also known as a “cake”). Recently, it was shown that the same algorithm can be used for allocating indivisible goods too. When applied directly, the Lone Divider algorithm finds only an ℓ -out-of- $((\ell+1)n-2)$ MMS allocation (Aigner-Horev & Segal-Halevi, 2022), which for small ℓ is substantially worse than our target approximation of ℓ -out-of- $\lfloor (\ell + \frac{1}{2})n \rfloor$. We overcome this difficulty by adding constraints on the ways in which the ‘lone divider’ is allowed to partition the goods, as well as arguing on which goods are selected to be included in each partition (see Section 4).

Bin Covering. To develop a polynomial-time algorithm when $\ell = 1$, we extend an algorithm of Csirik, Frenk, Labbè, and Zhang (1999) for the *bin covering* problem—a dual of

2. In general, ordinal and multiplicative approximations are incomparable from the theoretical standpoint—each of them may be larger than the other in some instances (see Appendix A). Therefore, we compare them through simulations using synthetic data.

the more famous *bin packing* problem (Johnson, 1973). In this problem, the goal is to fill as many bins as possible with items of given sizes, where the total size in each bin must be above a given threshold. This problem is NP-hard, but Csirik et al. (1999) presents a polynomial-time $2/3$ approximation. This algorithm cannot be immediately applied to the fair division problem since the valuations of goods are *subjective*, meaning that agents may have different valuations of each good. We adapt this technique to handle subjective valuations.

2. Related Work

2.1 Maximin Share

The idea of using the highest utility an agent could obtain if all other agents had the same preferences as a benchmark for fairness, originated in the economics literature (Moulin, 1990, 1992). It was put to practice in the context of course allocation by Budish (2011), where he introduced the ordinal approximation to MMS, and showed a mechanism that guarantees 1-out-of- $(n + 1)$ MMS to all agents by adding a small number of excess goods. In the more standard fair division setting, in which adding goods is impossible, the first non-trivial ordinal approximation was 1-out-of- $(2n - 2)$ MMS (Aigner-Horev & Segal-Halevi, 2022). Hosseini and Searns (2021) studied the connection between guaranteeing 1-out-of- n MMS for $2/3$ of the agents and the ordinal approximations for *all* agents. The implication of their results is the existence of 1-out-of- $\lceil 3n/2 \rceil$ MMS allocations and a polynomial-time algorithm for $n < 6$. Whether or not 1-out-of- $(n + 1)$ MMS can be guaranteed without adding excess goods remains an open problem to date.

The generalization of the maximin share to arbitrary $\ell \geq 1$ was first introduced by Babaioff, Nisan, and Talgam-Cohen (2019), Babaioff et al. (2021), and further studied by Segal-Halevi (2020). They presented this generalization as a natural fairness criterion for agents with different entitlements. The implication relations between ℓ -out-of- d MMS-fairness guarantees for different values of ℓ and d were characterized by Segal-Halevi (2019). Recently, the maximin share and its ordinal approximations have also been applied to some variants of the *cake-cutting* problem (Elkind, Segal-Halevi, & Suksompong, 2021c, 2021b, 2021a; Bogomolnaia & Moulin, 2022).

2.2 Multiplicative MMS Approximations

The multiplicative approximation to MMS originated in the computer science literature (Procaccia & Wang, 2014). The non-existence of MMS allocations (Kurokawa et al., 2018) and its intractability (Bouveret & Lemaître, 2016; Woeginger, 1997) have given rise to a number of approximation techniques.

These algorithms guarantee that each agent receives an approximation of their maximin share threshold. The currently known algorithms guarantee $\beta \geq 2/3$ (Kurokawa et al., 2018; Amanatidis, Markakis, Nikzad, & Saberi, 2017; Garg, McGlaughlin, & Taki, 2018) and $\beta \geq 3/4$ (Ghodsi, HajiAghayi, Seddighin, Seddighin, & Yami, 2018; Garg & Taki, 2020) in general, and $\beta \geq 7/8$ (Amanatidis et al., 2017) as well as $\beta \geq 8/9$ (Gourvès & Monnot, 2019) when there are only three agents. There are also MMS approximation algorithms for settings with constraints, such as when the goods are allocated on a cycle and each agent

must get a connected bundle (Truszczyński & Lonc, 2020). McGlaughlin and Garg (2020) showed an algorithm for approximating the maximum Nash welfare (the product of agents' utilities), which also attains a fraction $1/(2n)$ of the MMS.

Recently, Nguyen, Nguyen, and Rothe (2017) gave a Polynomial Time Approximation Scheme (PTAS) for a notion defined as *optimal-MMS*, that is, the largest value, β , for which each agent receives at least a fraction β of its MMS. Since the number of possible partitions is finite, an optimal-MMS allocation always exists, and it is an MMS allocation if $\beta \geq 1$. However, an optimal-MMS allocation may provide an arbitrarily bad ordinal MMS guarantee. Searns and Hosseini (2020), Hosseini and Searns (2021) show that for every n , there is an instance with n agents in which under *any* optimal-MMS allocation only a constant number of agents (≤ 4) receive their MMS value.

2.3 Fairness Based on Ordinal Information

An advantage of the ordinal MMS approximation is that it depends only on the ranking over the bundles. Other fair allocation algorithms with this robustness property are the Decreasing Demands algorithm of Herreiner and Puppe (2002), the Envy Graph algorithm of Lipton, Markakis, Mossel, and Saberi (2004), and the UnderCut algorithm of Brams, Kilgour, and Klamler (2012).

Amanatidis, Birmpas, and Markakis (2016), Halpern and Shah (2021) study an even stronger robustness notion, where the agents report only a ranking over the *goods*. Their results imply that, in this setting, the highest attainable multiplicative approximation of MMS is $\Theta(1/\log n)$.

Menon and Larson (2020) define a fair allocation algorithm as *stable* if it gives an agent the same value even if the agent slightly changes his cardinal valuations of goods, as long as the ordinal ranking of goods remains the same. They show that most existing algorithms are not stable, and present an approximately-stable algorithm for the case of two agents.

Finally, robustness has been studied also in the context of *fair cake-cutting*. Aziz and Ye (2014) define an allocation *robust-fair* if it remains fair even when the valuation of an agent changes, as long as its ordinal information remains unchanged. Edmonds and Pruhs (2011) study cake-cutting settings in which agents can only cut the cake with a finite precision.

3. Preliminaries

3.1 Agents and Goods

Let $N = [n] := \{1, \dots, n\}$ be a set of agents and M denote a set of m indivisible goods. We denote the value of agent $i \in N$ for good $g \in M$ by $v_i(g)$. We assume that the valuation functions are *additive*, that is, for each subset $G \subseteq M$, $v_i(G) = \sum_{g \in G} v_i(g)$, and $v_i(\emptyset) = 0$.³

An *instance* of the problem is denoted by $I = \langle N, M, V \rangle$, where $V = (v_1, \dots, v_n)$ is the valuation profile of agents. We assume all agents have a non-negative valuation for each good $g \in M$, that is, $v_i(g) \geq 0$. An *allocation* $A = (A_1, \dots, A_n)$ is an n -partition of M that allocates the bundle of goods in A_i to each agent $i \in N$.

3. In Appendix D we complement our results with a non-existence result for the more general class of *responsive* preferences.

It is convenient to assume that the number of goods is sufficiently large. Particularly, some algorithms implicitly assume that $m \geq n$, while some algorithms implicitly assume that $m \geq \ell \cdot n$. These assumptions are without loss of generality, since if m in the original instance is smaller, we can just add dummy goods with a value of 0 to all agents.

3.2 The Maximin Share

For every agent $i \in N$ and integers $1 \leq \ell < d$, the ℓ -out-of- d *maximin share* of i from M , denoted $\text{MMS}_i^{\ell\text{-out-of-}d}(M)$, is defined as

$$\text{MMS}_i^{\ell\text{-out-of-}d}(M) := \max_{\mathbf{P} \in \text{PARTITIONS}(M, d)} \min_{Z \in \text{UNION}(\mathbf{P}, \ell)} v_i(Z)$$

where the maximum is over all partitions of M into d subsets, and the minimum is over all unions of ℓ subsets from the partition. We say that an allocation A is an ℓ -out-of- d -MMS allocation if for all agents $i \in N$, $v_i(A_i) \geq \text{MMS}_i^{\ell\text{-out-of-}d}(M)$.

Obviously $\text{MMS}_i^{\ell\text{-out-of-}d}(M) \leq \frac{\ell}{d} v_i(M)$, and the equality holds if and only if M can be partitioned into d subsets with the same value. Note that $\text{MMS}_i^{\ell\text{-out-of-}d}(M)$ is a weakly-increasing function of ℓ and a weakly-decreasing function of d .

The value $\text{MMS}_i^{\ell\text{-out-of-}d}(M)$ is at least as large, and sometimes larger than, $\ell \cdot \text{MMS}_i^{1\text{-out-of-}d}(M)$. For example, suppose $\ell = 2$, there are $d - 1$ goods with value 1 and one good with value $\varepsilon < 1$. Then $\text{MMS}_i^{2\text{-out-of-}d}(M) = 1 + \varepsilon$ but $2 \cdot \text{MMS}_i^{1\text{-out-of-}d}(M) = 2\varepsilon$.

The maximin-share notion is scale-invariant in the following sense: if the values of each good for an agent, say i , are multiplied by a constant c , then agent i 's MMS value is also multiplied by the same c , so the set of bundles that are worth for i at least $\text{MMS}_i^{\ell\text{-out-of-}d}(M)$ does not change.

3.3 The Lone Divider Algorithm

A general formulation of the Lone Divider algorithm, based on Aigner-Horev and Segal-Halevi (2022), is shown in Algorithm 1. It accepts as input a set M of items and a threshold value t_i for each agent i . These values should satisfy the following condition for each agent $i \in N$.

Definition 1 (Reasonable threshold). Given a set M , a value function v_i on M , and an integer $n \geq 2$, a *reasonable threshold* for v_i is a real number $t_i \in \mathbb{R}$ satisfying the following condition: for every integer $k \in \{0, \dots, n - 1\}$ and any k disjoint subsets $B_1, \dots, B_k \subseteq M$, if

$$\forall c \in [k] : v_i(B_c) < t_i,$$

then there exists a partition of $M \setminus \cup_{c \in [k]} B_c$ into $M_1 \cup \dots \cup M_{n-k}$, such that

$$\forall j \in [n - k] : v_i(M_j) \geq t_i.$$

Informally, if any k unacceptable subsets are given away, then i can partition the remainder into $n - k$ acceptable subsets. In particular, the case $k = 0$ implies that agent i can partition the original set M into n acceptable subsets.

ALGORITHM 1: The Lone Divider algorithm. Based on Kuhn (1967).

Input: An instance $\langle N, M, V \rangle$ where N is the set of agents, M is the set of items, V is the vector of agents' valuations; and a reasonable threshold vector $(t_i)_{i=1}^n$ as denoted in Definition 1.

Output: A partition $M = A_1 \cup \dots \cup A_n$ such that $v_i(A_i) \geq t_i$ for all $i \in [n]$.

- 1: Some arbitrary agent $a \in N$ is asked to partition M into $|N|$ disjoint subsets, $(Y_j)_{j \in N}$, with $\forall j \in N : v_a(Y_j) \geq t_a$.
- 2: Define a bipartite graph G with the agents of N on the one side and the set $Y := \{Y_1, \dots, Y_{|N|}\}$ on the other side. Add an edge (i, Y_j) whenever $v_i(Y_j) \geq t_i$.
- 3: Find a maximum-cardinality envy-free matching⁵ in G . Give each matched element in Y to the agent paired to it in N .
- 4: Let $N \leftarrow$ the unallocated agents and $M \leftarrow$ the unallocated objects. If $N \neq \emptyset$ go back to Step 1.

Given an instance $I = \langle N, M, V \rangle$ with $N = [n]$, a vector $(t_i)_{i=1}^n$ of real numbers is called a *reasonable threshold vector for I* if t_i is a reasonable threshold for v_i for all $i \in N$.⁴

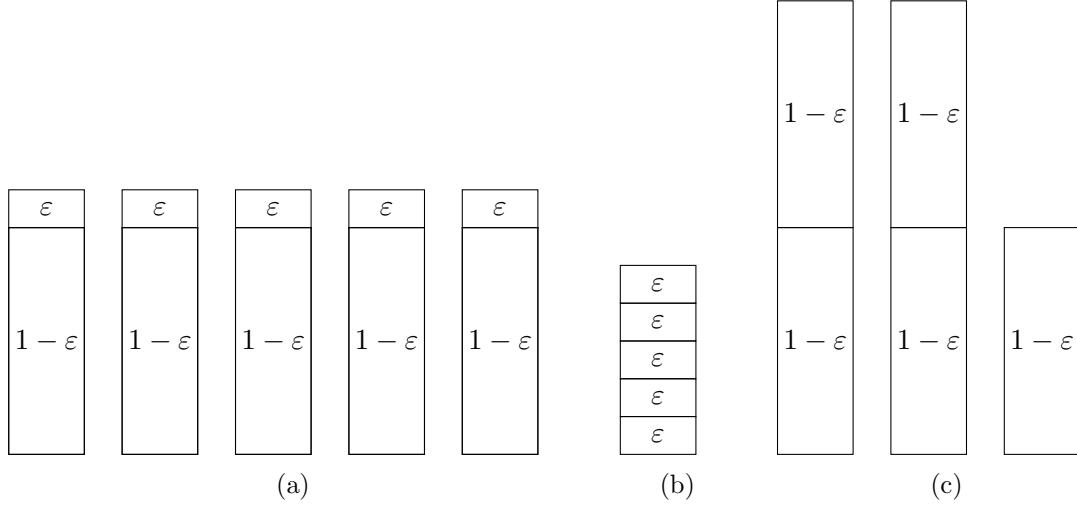
Example 1 (Reasonable threshold). Suppose M is perfectly divisible (e.g. a cake), and let $t_i := v_i(M)/n$. This threshold is reasonable, since if some k bundles with value less than t_i are given away, the value of the remaining cake is more than $(n - k)t_i$. Since the cake is divisible, it can be partitioned into $n - k$ acceptable subsets. This does not necessarily hold when M is a set of indivisible items; hence, finding reasonable thresholds for indivisible items setting is more challenging. ■

Algorithm Description Algorithm 1 proceeds in the following way: in each step, a single remaining agent is asked to partition the remaining goods into acceptable bundles—bundles whose values are above the divider's threshold. Then, all agents point at those bundles that are acceptable for them, and the algorithm finds an *envy-free matching* in the resulting bipartite graph.⁵ The matched bundles are allocated to the matched agents, and the algorithm repeats with the remaining agents and goods. It is easy to see that, if all threshold values t_i are reasonable, then Lone Divider guarantees agent i a bundle with a value of at least t_i . For example, when M is a cake, $t_i = v_i(M)/n$ is a reasonable threshold for every i , so Lone Divider can be used to attain a *proportional cake-cutting* (Kuhn, 1967).

When M is a set of indivisible goods, $t_i = \text{MMS}_i^{\ell\text{-out-of-}[(\ell+1)n-2]}(M)$ is a reasonable threshold for every $\ell \geq 1$ (Aigner-Horev & Segal-Halevi, 2022), so these ordinal approximations can all be computed directly through the Lone Divider algorithm. However, directly applying the Lone Divider algorithm cannot guarantee a better ordinal approximation, as we show next.

Example 2 (Execution of Algorithm 1). For simplicity, we present an example for $\ell = 1$. We show that, while Algorithm 1 can guarantee 1-out-of- $(2n - 2)$ MMS, it cannot

4. While we use the Lone Divider algorithm for allocating indivisible goods, it is a more general scheme that can also be used to divide chores or mixed items, divisible or indivisible. See Aigner-Horev and Segal-Halevi (2022) for details.
5. An *envy-free matching* in a bipartite graph $(N \cup Y, E)$ is a matching in which each unmatched agent in N is not adjacent to any matched element in Y . The bipartite graph generated by the Lone Divider algorithm always admits a nonempty envy-free matching, and a maximum-cardinality envy-free matching can be found in polynomial time (Aigner-Horev & Segal-Halevi, 2022).

Figure 1: An illustration of the goods' values in Example 2, for $n = 4$.

- (a) The $2n - 3 = 5$ MMS bundles of some agent.
- (b) The unacceptable bundle taken by the first divider.
- (c) The remaining goods, which cannot be combined into $n - 1 = 3$ acceptable bundles.

guarantee 1-out-of- $(2n - 3)$ MMS. Suppose that there are $4n - 6$ goods, and that all agents except the first divider value some $2n - 3$ goods at $1 - \varepsilon$ and the other $2n - 3$ goods at ε (see Figure 1). Then the 1-out-of- $(2n - 3)$ MMS of all these agents is 1.

However, it is possible that the first divider takes an unacceptable bundle containing all $2n - 3$ goods of value ε . Then, no remaining agent can partition the remaining goods into $n - 1$ bundles of value at least 1. In this instance, it is clear that while $\text{MMS}_i^{\ell\text{-out-of-}(2n-2)}(M)$ is a reasonable threshold, $\text{MMS}_i^{\ell\text{-out-of-}(2n-3)}(M)$ is not. ■

4. Ordinal Approximation of MMS for Goods

In this section we prove the following theorem.

Theorem 1. *Given an additive goods instance, an ℓ -out-of- d MMS allocation always exists when $d = \lfloor (\ell + \frac{1}{2})n \rfloor$.*

The proof is constructive: we present an algorithm (Algorithm 2) for achieving the above MMS bound. Since the algorithm needs to know the exact MMS thresholds for each agent (which is NP-hard to compute), its run-time is not polynomial. In Section 5 we present a different algorithm to compute ℓ -out-of- d MMS allocation when $\ell = 1$ in polynomial-time.

Algorithm 2 starts with two normalization steps, some of which appeared in previous works and some are specific to our algorithm. For completeness, we describe the normalization steps in Sections 4.1 and 4.2. The algorithm applies to the normalized instance an adaptation of the Lone Divider algorithm, in which the divider in each step must construct a *balanced* partition. We explain this notion in Section 4.3.

4.1 Scaling

We start by scaling the valuations such that $\text{MMS}_i^{\ell\text{-out-of-}d}(M) = \ell$ for each agent i . The scale-invariance property implies that such rescalings do not modify the set of bundles that are acceptable for i . Then, for each i we perform an additional scaling as follows.

- Consider a particular d -partition attaining the maximum in the $\text{MMS}_i^{\ell\text{-out-of-}d}(M)$ definition. Call the d bundles in this partition the *MMS bundles* of agent i .
- Denote the total value of the $\ell - 1$ least-valuable MMS bundles by x_i (or just x , when i is clear from the context). By definition, the value of the ℓ -th MMS bundle must be exactly $\ell - x$, while the value of each of the other $d - \ell$ MMS bundles is at least $\ell - x$.
- For each MMS bundle with value larger than $\ell - x$, arbitrarily pick one or more goods and decrease their value until the value of the MMS bundle becomes exactly $\ell - x$. Note that this does not change the MMS value.

After the normalization, the sum of values of all goods is

$$\begin{aligned} v_i(M) &= (d - \ell + 1) \cdot (\ell - x_i) + x_i \\ &= \ell + (d - \ell)(\ell - x_i) = d + (d - \ell)(\ell - 1 - x_i). \end{aligned}$$

Since $d = \lfloor (\ell + \frac{1}{2})n \rfloor \geq (\ell + \frac{1}{2})n - \frac{1}{2} = \ell n + n/2 - 1/2$,

$$\begin{aligned} v_i(M) &\geq (\ell n + n/2 - 1/2) + (\ell n + n/2 - 1/2 - \ell)(\ell - 1 - x_i) \\ &= n \cdot \ell + (n - 1) \cdot \ell(\ell - 1 - x_i) + (n - 1) \cdot (\ell - x_i)/2. \end{aligned} \tag{1}$$

The goal of the algorithm is to give each agent i a bundle A_i with $v_i(A_i) \geq \ell$. We say that such a bundle is *acceptable* for i .

Example 3 (Scaling). To illustrate the parameter x , consider the following two instances with $n = 5, \ell = 3$ and $d = \lfloor (\ell + \frac{1}{2})n \rfloor = 17$.

1. There are 17 goods with the value of 1.
2. There are 16 goods valued 1.2 and one good with the value of 0.6.

Here, each MMS bundle contains a single good. In both cases, the value of every 3 goods is at least 3. In the first case $x = 2$ and the total value is $5 \cdot 3 + 4 \cdot 3 \cdot 0 + 4 \cdot 1/2 = 17$. In the second case, $x = 1.8$ and the total value is $5 \cdot 3 + 4 \cdot 3 \cdot 0.2 + 4 \cdot 1.2/2 = 19.8$. ■

4.2 Ordering the Instance

As in previous works (Bouveret & Lemaître, 2016; Barman & Krishna Murthy, 2017; Garg et al., 2018; Huang & Lu, 2021), we apply a preliminary step in which the instance is *ordered*, i.e., $v_i(g_1) \geq \dots \geq v_i(g_m)$ for each agent $i \in N$. Ordering is done as follows:

- Index the goods in M arbitrarily g_1, \dots, g_m .

ALGORITHM 2: Finding an ℓ -out-of- $\lfloor(\ell + \frac{1}{2})n\rfloor$ MMS allocation.

Input: An instance $\langle N, M, V \rangle$ and an integer $\ell \geq 1$.

Output: An ℓ -out-of- $\lfloor(\ell + \frac{1}{2})n\rfloor$ MMS allocation.

- 1: Scale the valuations of all agents as explained in Section 4.1.
- 2: Order the instance as explained in Section 4.2.
- 3: Run the Lone Divider algorithm (Algorithm 1) with threshold values $t_i = \ell$ for all $i \in N$, with the restriction that, in each partition made by the lone divider, all bundles must be ℓ -balanced (Definition 2).

- Tell each agent i to adopt, for the duration of the algorithm, a modified value function that assigns, to each good g_j , the value of the j -th most valuable good according to i . For example, the new $v_i(g_1)$ should be the value of i 's most-valuable good; the new $v_i(g_m)$ should be the value of i 's least-valuable good; etc. Ties are broken arbitrarily.

During the execution of the algorithm, each agent answers all queries according to this new value function. For example, an agent asked whether the bundle $\{g_1, g_4, g_5\}$ is acceptable, should answer whether the bundle containing his best good, 4th-best good and 5th-best good is acceptable. Once the algorithm outputs an allocation, it can be treated as a *picking sequence* in which, for example, an agent who receives the bundle $\{g_1, g_4, g_5\}$ has the first, fourth and fifth turns. It is easy to see that such an agent receives a bundle that is at least as good as the bundle containing her best, 4th-best and 5th-best goods. Hence, if the former is acceptable then the latter is acceptable too.

Clearly, given an *unordered* instance, its corresponding ordered instance can be generated in polynomial time (for each agent $i \in [n]$, we need $O(m \log m)$ steps for ordering the valuations). Given an allocation for the ordered instance, one can compute the allocation for the corresponding unordered instance in time $O(n)$, using the picking-sequence described above.

4.3 Restricted Lone Divider

In Section 3.3 we illustrated the limitations of the plain Lone Divider algorithm (Algorithm 1). We can improve its performance by restricting the partitions that the lone divider is allowed to make in Step 1 of Algorithm 1. Without loss of generality, we may assume (by adding dummy goods if needed) that $m \geq n \cdot \ell$.

For every $l \in \{1, \dots, \ell\}$, denote $G_l^n := \{g_{(l-1)n+1}, \dots, g_{ln}\}$. In other words, G_1^n contains the n most-valuable goods; G_2^n contains the n next most-valuable goods; and so on. Since the instance is ordered, these sets are the same for all agents.

Definition 2 (ℓ -balanced bundle). Given an ordered instance and an integer $\ell \geq 1$, a nonempty bundle $B \subseteq M$ is called ℓ -balanced if

- B contains exactly one good from G_1^n .
- If $|B| \geq 2$, then B contains exactly one good from G_2^n .
- If $|B| \geq 3$, then B contains exactly one good from G_3^n .

- ...

- If $|B| \geq \ell$, then B contains exactly one good from G_ℓ^n .

Note that an ℓ -balanced bundle contains at least ℓ goods. The definition of ℓ -balanced bundles only constrains the allocation of the first ℓn goods; there may be arbitrarily many additional goods in $M \setminus \bigcup_{i=1}^\ell G_i^n$, and they may be allocated arbitrarily.

Algorithm 2 requires the lone divider to construct a partition in which all n bundles are ℓ -balanced.

Example 4 (ℓ -balanced bundles). Suppose there are five agents ($n = 5$) and $m = 20$ goods, where the value of each good $j \in [20]$ is precisely j for all agents. Then, a 1-balanced bundle must contain a good $j \in \{20, 19, 18, 17, 16\}$; a 2-balanced bundle must contain a good from $\{20, 19, 18, 17, 16\}$ and a good from $\{15, 14, 13, 12, 11\}$; a 3-balanced bundle must contain, in addition to these, a good from $\{10, 9, 8, 7, 6\}$; and so on. ■

4.4 Construction for a Single Divider

In order to prove the correctness of Algorithm 2, it is sufficient to prove that the threshold value $t_i = \ell$ is a *reasonable threshold* (see Definition 1) for each agent i , with the additional restriction that all bundles should be ℓ -balanced.

To do this, it is sufficient to consider a single divider, Alice. We denote her normalized ordered value measure by v , and the sum of her $\ell - 1$ least-valuable MMS bundles by x . We consider a particular MMS partition for Alice, and refer to the bundles in this partition as the *MMS bundles*.

Assume that k unacceptable bundles $(B_c)_{c=1}^k$ have already been given to other agents and that all these bundles are ℓ -balanced. Therefore, for each $c \in [k]$, it must be that $v(B_c) < \ell$.⁶ We have to prove that Alice can use the remaining goods to construct $n - k$ acceptable bundles that are also ℓ -balanced. Particularly, we prove below that Alice can construct $n - k$ acceptable bundles, each of which contains exactly 1 remaining good from each of G_1^n, \dots, G_ℓ^n .

4.5 Main Idea: Bounding the Waste

Given a bundle B_a , denote its *waste* by $w(B_a) := v(B_a) - \ell$. This is the value the bundle contains beyond the acceptability threshold of ℓ . Note that the waste of acceptable bundles is positive and that of unacceptable bundles is negative. The total initial value for Alice is given by (1). The total waste she can afford in her partition is therefore

$$v(M) - n \cdot \ell = (n - 1) \cdot (\ell - x)/2 + (n - 1) \cdot \ell(\ell - 1 - x).$$

The first term implies that she can afford an average waste of $(\ell - x)/2$ for $n - 1$ bundles; the second term implies that she can afford an average waste of $\ell(\ell - 1 - x)$ for $n - 1$ bundles.

6. Recall that the Lone Divider algorithm allocates bundles using an envy-free matching. This means that all bundles allocated before Alice's turn are unacceptable to Alice.

Example 5 (Bounding the waste). Consider Example 3. In case (1), the total value is 17 and we need 5 bundles with a value of 3, so the affordable waste is 2. The average over 4 bundles is $0.5 = (3 - 2)/2 + 3 \cdot 0$. In case (2), the total value is 19.8, so the affordable waste is 4.8. The average over 4 bundles is $1.2 = (3 - 1.8)/2 + 3 \cdot (0.2)$. In both cases, if there are 4 acceptable bundles with that amount of waste, then the remaining value is exactly 3, which is sufficient for an additional acceptable bundle. ■

The following lemma formalizes this observation.

Lemma 1. *Suppose there exists a partition of M into*

- *Some $t \geq 0$ bundles with an average waste of at most $(\ell - x)/2 + \ell(\ell - 1 - x)$;*
- *A subset S of remaining goods, with $v(S) < \ell$.*

Then $t \geq n$.

Proof. For brevity, we denote $w := (\ell - x)/2 + \ell(\ell - 1 - x)$. The *total* value of the bundles equals their number times their *average* value. So the total value of the t bundles is at most $t \cdot \ell + t \cdot w$. After adding $v(S) < \ell$ for the remaining goods, the sum equals $v(M)$, so

$$\begin{aligned} (t + 1) \cdot \ell + t \cdot w &> v(M) \\ \geq n \cdot \ell + (n - 1) \cdot w &\quad (\text{by (1)}). \end{aligned}$$

Therefore, at least one of the two terms in the top expression must be larger than the corresponding term in the bottom expression. This means that either $(t + 1)\ell > n\ell$, or $tw > (n - 1)w$. Both options imply $t \geq n$. □

Remark 1. The value of each MMS bundle is at most $\ell - x$. Therefore, any bundle that is the union of exactly ℓ such MMS bundles has a waste of at most $\ell(\ell - x) - \ell = \ell(\ell - 1 - x)$ and thus it satisfies the upper bound of Lemma 1. In particular, this is satisfied for every bundle with at most ℓ goods.

Below we show how Alice can find a partition in which the average waste is upper bounded as in Lemma 1. This partition will consist of the following bundles:

- The k previously-allocated bundles, with waste < 0 (since they are unacceptable);
- Some newly-constructed bundles with exactly ℓ goods and waste $\leq \ell(\ell - 1 - x)$ (by Remark 1);
- Some newly-constructed bundles with waste at most $(\ell - x)/2$;
- Some pairs of bundles, where the waste in one is larger than $(\ell - x)/2$ but the waste in the other is smaller than $(\ell - x)/2$, such that the average is at most $(\ell - x)/2$.

4.6 Step 0: Bundles with Exactly ℓ Goods

Recall that before Alice's turn, some k bundles have been allocated, with a value of less than ℓ . Hence, their waste is less than 0. Since these bundles are ℓ -balanced, they contain exactly one good from each of G_1^n, \dots, G_ℓ^n (and possibly some additional goods).

Therefore, exactly $n - k$ goods are available in each of G_1^n, \dots, G_ℓ^n .

Next, Alice checks all the ℓ -tuples containing one good from each of G_1^n, \dots, G_ℓ^n (starting from the highest-valued goods in each set). If the value of such an ℓ -tuple is at least ℓ , then it is acceptable and its waste is at most $\ell(\ell - 1 - x)$ by Remark 1.

After Step 0, there are some $k' \geq k$ bundles with a waste of at most $\ell(\ell - 1 - x)$, each of which contains exactly one good from each of G_1^n, \dots, G_ℓ^n . Of these, k are previously-allocated bundles, and $k' - k$ are newly-constructed acceptable bundles. In each of G_1^n, \dots, G_ℓ^n , there remain exactly $n - k'$ goods. The total value of each ℓ -tuple of remaining goods from G_1^n, \dots, G_ℓ^n is less than ℓ . Alice will now construct from them some $n - k'$ bundles with an average waste of at most $(\ell - x)/2 + \ell(\ell - 1 - x)$. Lemma 1 implies the following lemma on the remaining goods (the goods not in these k' bundles):

Lemma 2. *Suppose there exists a partition of the remaining goods into*

- Some $t \geq 0$ bundles with an average waste of at most $(\ell - x)/2 + \ell(\ell - 1 - x)$;
- A subset S of remaining goods, with $v(S) < \ell$.

Then $t \geq n - k'$.

Alice's strategy branches based on the number of *high-value goods*.

4.7 High-Value Goods

We define high-value goods as goods g with $v(g) > (\ell - x)/2$. Denote by h the number of high-valued goods in M . Since the instance is ordered, goods g_1, \dots, g_h are high-valued. All MMS bundles are worth at most $\ell - x$, and therefore may contain at most one high-value good each. Since the number of MMS bundles is $\ell n + n/2$, we have $h \leq \ell n + n/2$.

For each $j \in [h]$, we denote

- $M_j :=$ the MMS bundle containing g_j . Since the value of all MMS bundles is at most $(\ell - x)$, each MMS bundle contains at most one high-value good, so the M_j are all distinct.
- $R_j :=$ the *remainder set* of g_j , i.e., the set $M_j \setminus \{g_j\}$.
- $r_j := v(R_j)$.

We consider three cases, based on the number of high-value goods.

Case #1: $h \leq \ell n$. This means that all high-value goods are contained in $G_1^n \cup \dots \cup G_\ell^n$, so after removing $(B_c)_{c=1}^{k'}$, at most $\ell n - \ell k'$ high-value goods remain—at most $n - k'$ in each of G_1^n, \dots, G_ℓ^n . Alice constructs the required bundles by *bag-filling*—a common technique in MMS approximations (e.g. Garg et al. (2018)).

- Repeat at most $n - k'$ times:

- Initialize a bag with a good from each of G_1^n, \dots, G_ℓ^n (Step 0 guarantees that the total value of these goods is less than ℓ).
- Fill the bag with goods from outside G_1^n, \dots, G_ℓ^n . Stop when either no such goods remain, or the bag value raises above ℓ .

Since all goods used for filling the bag have a value of at most $(\ell - x)/2$, the waste of each constructed bundle is at most $(\ell - x)/2$. By construction, all these bundles are acceptable except the last one. Apply Lemma 2 with S being the set of goods remaining in the last bag, and t being the number of acceptable bundles constructed by bag-filling. The lemma implies that $t \geq n - k'$.

Case #2: $k' \geq n/2$. Alice uses bag-filling as in Case #1.

Here, the waste per constructed bundle might be more than $(\ell - x)/2$. However, since the value of a single good is at most $\ell - x$, the waste of each constructed bundle is at most $\ell - x$.

In each of the k' bundles of Step 0, the waste is at most $\ell(\ell - 1 - x)$. Since $k' \geq n/2 \geq n - k'$, the *average* waste per bundle is at most $\ell(\ell - 1 - x) + (\ell - x)/2$. Hence Lemma 2 applies, and at least $n - k'$ acceptable bundles are constructed.

Case #3: $h > \ell n$ and $k' < n/2$. In this case, Alice will have to construct some bundles with waste larger than $(\ell - x)/2$. However, she will compensate for it by constructing a similar number of bundles with waste smaller than $(\ell - x)/2$, such that the average waste per bundle remains at most $(\ell - x)/2$.

After removing $(B_c)_{c=1}^{k'}$, exactly $h - \ell k'$ high-value goods remain. They can be partitioned into two subsets:

- $H_+ :=$ the $(n - k')\ell$ top remaining goods — those contained in $G_1^n \cup \dots \cup G_\ell^n$; exactly $n - k'$ in each of G_1^n, \dots, G_ℓ^n . By assumption, $n - k' > n/2$.
- $H_- :=$ the other $h - \ell n$ high-value goods — those not contained in $G_1^n \cup \dots \cup G_\ell^n$. Since $h \leq \ell n + n/2$, the set H_- contains at most $n/2$ goods.

This is the hardest case; to handle this case, we proceed to Step 1 below.

4.8 Step 1: Bundling High-value Goods.

Alice constructs at most $|H_-|$ bundles as follows.

- Repeat while H_- is not empty:
 - Initialize a bag with the lowest-valued remaining good from each of G_1^n, \dots, G_ℓ^n (Step 0 guarantees that their total value is less than ℓ).
 - Fill the bag with goods from H_- , until the bag value raises above ℓ .

Note that $|H_-| \leq n/2 < n - k' = |G_1^n| = \dots = |G_\ell^n|$, so as long as H_- is nonempty, each of G_1^n, \dots, G_ℓ^n is nonempty too, and Alice can indeed repeat. By construction, all filled bags except the last one are valued at least ℓ ; it remains to prove that the number of these bags is sufficiently large.

Let s be the number of acceptable bundles constructed once H_- becomes empty. Note that, in addition to these bundles, there may be an *incomplete bundle* — the last bundle, whose construction was terminated while its value was still below ℓ .

Let $P_+ \subseteq H_+$ be the set of $s\ell$ goods from H_+ in the acceptable bundles. Denote by $P_- \subseteq H_-$ the set of s goods from H_- that were added *last* to these s bundles (bringing their value from less-than- ℓ to at-least- ℓ). Note that the waste in each of these s bundles might be larger than $(\ell - x)/2$, but it is at most the value of a single good from P_- , so the total waste is at most $\sum_{j \in P_-} v(g_j)$.

After this step, besides the s acceptable bundles, there are some $(n - k' - s)\ell$ high-value goods remaining in H_+ (some ℓ of these goods are possibly in the incomplete bundle, if such a bundle exists). Alice now has to construct from them some $n - k' - s$ acceptable bundles.

4.9 Step 2: Using the Remainders.

Alice now constructs acceptable bundles by bag-filling. She initializes each bag with the incomplete bundle from Step 1 (if any), or with an ℓ -tuple of unused goods from G_1^n, \dots, G_ℓ^n . Then, she fills the bag with low-value goods from the following *remainder sets*:

- There are ℓk remainder sets that correspond to the ℓk goods allocated within the k unacceptable bundles $(B_c)_{c=1}^k$. We denote them by $R_{c,1}^U, \dots, R_{c,\ell}^U$ and their values by $r_{c,1}^U, \dots, r_{c,\ell}^U$ for $c \in [k]$.
- There are $\ell s + s$ remainder sets that correspond to the ℓs high-value goods in P_+ and the s high-value goods in P_- . We denote them by R_j^P and their values by r_j^P for $j \in P_+ \cup P_-$.

By definition, the total value of all these remainders is:

$$\begin{aligned} \text{TOTAL-REMAINDER-VALUE} &= v \left(\bigcup_{j \in P_+ \cup P_-} R_j^P \cup \bigcup_{c=1}^k \bigcup_{l=1}^{\ell} R_{c,l}^U \right) \\ &= \sum_{j \in P_+ \cup P_-} r_j^P + \sum_{c=1}^k \sum_{l=1}^{\ell} r_{c,l}^U. \end{aligned}$$

For each remainder-set R_j , denote by R'_j , the subset of R_j that remains after removing the at most k unacceptable bundles $(B_c)_{c=1}^k$ with more than ℓ goods.⁷ Each unacceptable bundle B_c contains, in addition to the ℓ high-value goods $g_{c,l}$ for $l \in \{1, \dots, \ell\}$, some low-value goods with a total value of less than $\sum_{l=1}^{\ell} r_{c,l}^U$ (since the total value of the unacceptable bundle is less than ℓ). Therefore, the total value of low-value goods included in these unacceptable bundles is at most $\sum_{c=1}^k \sum_{l=1}^{\ell} r_{c,l}^U$ (equality holding iff $k = 0$). Therefore, the

7. Bundles B_c with at most ℓ goods do not consume anything from the remainder-sets R_j , since they contain only high-value goods from G_1^n, \dots, G_ℓ^n . From the same reason, the $k' - k$ acceptable bundles constructed in Step 0 do not consume anything from the remainder sets.

total remaining value satisfies

$$\begin{aligned}
 \text{TOTAL-REMAINDER-VALUE} &= v \left(\bigcup_{j \in P_+ \cup P_-} R'_j{}^P \cup \bigcup_{c=1}^k \bigcup_{l=1}^{\ell} R'_{c,l}{}^U \right) \\
 &\geq \left(\sum_{j \in P_+ \cup P_-} r_j{}^P + \sum_{c=1}^k \sum_{l=1}^{\ell} r_{c,l}{}^U \right) - \sum_{c=1}^k \sum_{l=1}^{\ell} r_{c,l}{}^U \\
 &= \sum_{j \in P_+ \cup P_-} r_j{}^P.
 \end{aligned} \tag{2}$$

The bag-filling proceeds as follows.

1. Initialize $a := 1$.
2. Initialize a bag with either the incomplete bundle from Step 1 (if any), or some ℓ unused top goods. We denote the ℓ top goods used for initializing bag a by $g_{j[a,1]} \in G_1^n, \dots, g_{j[a,\ell]} \in G_{\ell}^n$.
3. Add to the bag the remainder-sets $R'_j{}^P$ and $R'_{c,l}{}^U$ in an arbitrary order. Stop when either no such remainder-sets remain, or the bag value raises above ℓ .
4. If there are still some unused remainder-sets and high-value goods, let $a := a + 1$ and go back to Step 2.

The bag-filling stops when either there are no more high-value goods, or no more remainder-sets. In the former case, Alice has all $n - k'$ required bundles (s from Step 1 and $n - k' - s$ from Step 2), and the construction is done. We now analyze the latter case.

By construction, we go to the next bag only after the current bag becomes at least ℓ . Therefore, all bags except the last one are valued at least ℓ . Our goal now is to prove that the number of these “all bags except the last one” is sufficiently large.

Let t be the number of bundles constructed with a value of at least ℓ . For each $a \in [t]$, The a -th bag contains the high-value goods $g_{j[a,1]}, \dots, g_{j[a,\ell]}$ and some remainder-sets. How much remainder-sets should it contain? Suppose it contains remainder-sets with a total value of $\sum_{l=1}^{\ell} r_{j[a,l]}$. Then, the total bundle value is $\sum_{l=1}^{\ell} v(M_{j[a,l]})$. By assumption, the total value of every ℓ MMS bundles is at least ℓ , so the bundle value is at least ℓ . Therefore, to make bundle a acceptable, it is sufficient to add to it a value of $\sum_{l=1}^{\ell} r_{j[a,l]}$.

Denote by $j[a, *]$ the index of the last remainder-set added to bag a (bringing its value from less-than- ℓ to at-least- ℓ). The total value of remainder-sets in the bag is thus less than $r_{j[a,*]} + \sum_{l=1}^{\ell} r_{j[a,l]}$.

The total value of remainder-sets in the unfilled $(t+1)$ -th bag is less than $\sum_{l=1}^{\ell} r_{j[t+1,l]}$, where $j[t+1, 1], \dots, j[t+1, \ell]$ are indices of some remaining high-value goods. Therefore,

the total value of remainder-sets in all $t + 1$ bags together satisfies

$$\begin{aligned}
 v \left(\bigcup_{j \in P_+ \cup P_-} R_j'^P \cup \bigcup_{c=1}^k \bigcup_{l=1}^{\ell} R_{c,l}^U \right) &< \sum_{a=1}^t \left(r_{j[a,*]} + \sum_{l=1}^{\ell} r_{j[a,l]} \right) + \left(\sum_{l=1}^{\ell} r_{j[t+1,l]} \right) \\
 &= \left(\sum_{a=1}^{t+1} \sum_{l=1}^{\ell} r_{j[a,l]} \right) + \left(\sum_{a=1}^t r_{j[a,*]} \right). \tag{3}
 \end{aligned}$$

Combining (2) and (3) gives

$$\left(\sum_{a=1}^{t+1} \sum_{l=1}^{\ell} r_{j[a,l]} \right) + \left(\sum_{a=1}^t r_{j[a,*]} \right) > \sum_{j \in P_+ \cup P_-} r_j^P. \tag{4}$$

In the left-hand side there are $\ell(t + 1) + t = (\ell + 1)t + \ell$ terms, while in the right-hand side there are $(\ell + 1)s$ terms — $\ell + 1$ for each bundle constructed in Step 1. We now show that each term in the left-hand side is equal or smaller than a unique term in the right-hand side. Since the left-hand side is overall larger than the right-hand side, this indicates that the left-hand side must have more terms, that is, $(\ell + 1)t + \ell > (\ell + 1)s$. This implies that $t \geq s$, i.e., Alice has successfully constructed from the remainder-sets some s acceptable bundles.

- Consider first the $\ell(t + 1)$ terms $r_{j[a,l]}$, and compare them to r_j^P for $j \in P_+$. Since the bundles in Step 1 were constructed in ascending order of value, starting at the lowest-valued available goods in each of G_1^n, \dots, G_{ℓ}^n , every index $j[a, l]$ is smaller than any index $j \in G_l^n$. Therefore, every term $r_{j[a,l]}$ is smaller than some unique term r_j^P for $j \in G_l^n$, for every $l \in [\ell]$.
- Consider now the t terms $r_{j[a,*]}$, and compare them to r_j^P for $j \in P_-$. Each of the indices $j[a, *]$ is an index of some unique remainder-set, so it is either equal to some unique index $j \in P_+ \cup P_-$, or to some unique index c, l (the index some remainder-set $R_{c,l}^U$ of some unacceptable bundle B_c). All indices c, l are in $\{1, \dots, \ell n\}$, so they are smaller than the indices $j \in P_-$. Therefore, every $r_{j[a,*]}$ is either equal or smaller than some unique term r_j^P for $j \in P_-$.

So Alice has s new acceptable bundles. The waste of each of these is $r_{j[a,*]}$, which — as mentioned above — is equal to or smaller than some unique term r_j^P for $j \in P_-$. Therefore, the total waste of all these s bundles is at most the following sum of s terms: $\sum_{j \in P_-} r_j^P$.

Recall that the waste of each of the s acceptable bundles from Step 1 was at most $v(g_j)$ for some $j \in P_-$. Therefore, the total waste of the $2s$ acceptable bundles constructed so far

is at most

$$\begin{aligned}
 & \sum_{j \in P_-} r_j^P + \sum_{j \in P_-} v(g_j) \\
 &= \sum_{j \in P_-} (r_j^P + v(g_j)) \\
 &= \sum_{j \in P_-} v(M_j) \\
 &\leq \sum_{j \in P_-} (\ell - x) \quad \text{by the normalization (Section 4.1)} \\
 &= |P_-| \cdot (\ell - x) \\
 &= s \cdot (\ell - x).
 \end{aligned}$$

Therefore, the average waste per bundle is at most $s(\ell - x)/(2s) = (\ell - x)/2$.

4.10 Step 3: Plain Bag-Filling.

At this stage, there are no more high-value goods outside H_+ . Therefore, Alice can construct the remaining bundles by plain bag-filling, initializing each bag with some ℓ -tuple of unused goods remaining in H_+ , and filling it with some low-value goods outside H_+ . Since the waste in each bundle is at most $(\ell - x)/2$, Lemma 2 implies that the total number of constructed bundles is $n - k'$.

This completes the proof that ℓ is a reasonable threshold for Algorithm 2. Therefore, the algorithm finds the allocation promised in Theorem 1.

4.11 Limits of Algorithm 2

To illustrate the limitation of Algorithm 2, we show that it cannot guarantee 1-out-of- $((\ell + \frac{1}{2})n - 2)$ MMS. For simplicity we assume that n is even so that $(\ell + \frac{1}{2})n$ is an integer.

Example 6 (Tight bound for our technique). Suppose that in the first iteration all agents except the divider have the following MMS bundles:

- $\ell n - 1$ bundles are made of two goods with values $1 - \varepsilon$ and ε .
- One bundle is made of two goods with values $1 - \ell n \varepsilon$ and $\ell n \varepsilon$.
- $n/2 - 2$ bundles are made of two goods with values $1/2, 1/2$.

So their 1-out-of- $(\ell n + n/2 - 2)$ MMS equals ℓ . However, it is possible that the first divider takes an unacceptable bundle containing $\ell - 1$ goods of value $1 - \varepsilon$, the good of value $1 - \ell n \varepsilon$, and the $\ell n - 1$ goods of value ε . Note that this bundle is ℓ -balanced. All remaining goods have a value of less than 1, so an acceptable bundle requires at least $\ell + 1$ goods. However, the number of remaining goods is only $\ell n - \ell + 1 + n - 4 = (\ell + 1)(n - 1) - 2$: $\ell n - \ell$ goods of value $1 - \varepsilon$, one good of value $\ell n \varepsilon$ and $n - 4$ goods of value $1/2$. Hence, at most $n - 2$ acceptable bundles can be constructed. ■

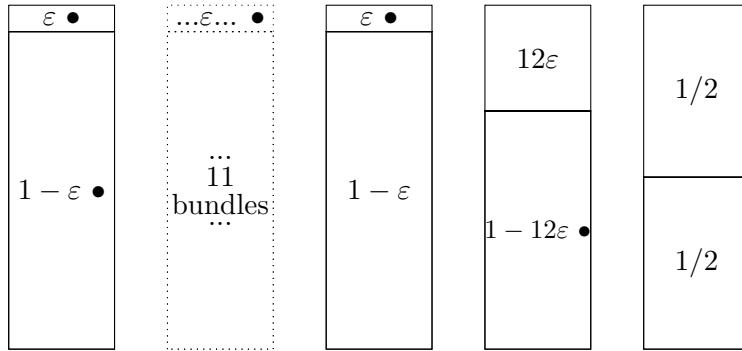


Figure 2: An illustration of the goods' values in Example 6, for $n = 6$ and $\ell = 2$. Here, $(\ell + 1/2)n - 2 = 13$. The 2-out-of-13 MMS is 2. Each rectangle represents an MMS bundle containing two goods. The first divider takes the goods marked by a bullet. Note that this is a 2-balanced bundle. The next divider cannot construct 5 bundles of value 2 from the remaining goods.

5. Ordinal Approximation for Goods in Polynomial Time

Algorithm 2 guarantees that each agent receives an ℓ -out-of- d MMS allocation for $d \geq \lfloor (\ell + \frac{1}{2})n \rfloor$. However, the algorithm requires exact MMS values to determine whether a given bundle is acceptable to each agent. Since computing an exact MMS value for each agent is NP-hard, Algorithm 2 does not run in polynomial-time even for the case of $\ell = 1$. The objective of this section is to develop polynomial-time approximation algorithms for computing ℓ -out-of- d MMS allocations.

We utilize optimization techniques used in the *bin covering* problem. This problem was presented by Assmann, Johnson, Kleitman, and Leung (1984) as a dual of the more famous *bin packing* problem. In the bin covering problem, the goal is to fill bins with items of different sizes, such that the sum of sizes in each bin is at least 1, and subject to this, the number of bins is maximized. This problem is NP-hard, but several approximation algorithms are known. These approximation algorithms typically accept a bin-covering instance I as an input and fill at least $a \cdot (OPT(I) - b)$ bins, where $a < 1$ and $b > 0$ are constants, and $OPT(I)$ is the maximum possible number of bins in I . Such an algorithm can be used directly to find an ordinal approximation of an MMS allocation when all agents have *identical* valuations. Our challenge is to adapt them to agents with *different* valuations.

5.1 The case when $\ell = 1$

For the case when $\ell = 1$, we adapt the algorithm of Csirik et al. (1999), which finds a covering with at least $\frac{2}{3} \cdot (OPT(I) - 1)$ bins (an approximation with $a = \frac{2}{3}$ and $b = 1$). Algorithm 3 generalizes the aforementioned algorithm to MMS allocation of goods. Thus, the algorithm of Csirik et al. (1999) corresponds to a special case of Algorithm 3 wherein

- All agents have the same v_i (describing the item sizes); and

ALGORITHM 3: Bidirectional bag-filling

Input: An instance $\langle N, M, V \rangle$ and threshold values $(t_i)_{i=1}^n$.**Output:** At most n subsets A_i satisfying $v_i(A_i) \geq t_i$.

- 1: Order the instance in **descending order** of value as in Section 4.2, so that for each agent i , $v_i(g_1) \geq \dots \geq v_i(g_m)$.
- 2: **for** $k = 1, 2, \dots$ **do**
- 3: Initialize a bag with the good g_k .
- 4: Add to the bag zero or more remaining goods in **ascending order** of value, until at least one agent i values the bag at least t_i .
- 5: Give the goods in the bag to an arbitrary agent i who values it at least t_i .
- 6: If every remaining agent i values the remaining goods at less than t_i , stop.
- 7: **end for**

- All agents have the same t_i (describing the bin size).⁸

For this case, we have the following lemma:

Lemma 3 (Lemma 4 of Csirik et al. (1999)). *When all agents have the same valuation v and the same threshold t , Algorithm 3 allocates at least $\frac{2}{3}(OPT(v, t) - 1)$ bundles, where $OPT(v, t)$ is the maximum number that can be filled.*

Note that Algorithm 3 works for any selection of the threshold values t_i , but if the thresholds are too high, it might allocate fewer than n bundles. Our challenge now is to compute thresholds for which n bundles are allocated. To compute a threshold for agent i , we simulate Algorithm 3 using n clones of i , that is, n agents with valuation v_i . We look for the largest threshold for which this simulation allocates at least n bundles.

Definition 3. The *1-out-of- n bidirectional-bag-filling-share of agent i* , denoted $BBFS_i^n$, is the largest value t_i for which Algorithm 3 allocates at least n bundles when executed with n agents with identical valuation v_i and identical threshold t_i .

The BBFS of agent i can be computed using binary search up to ε , where ε is the smallest difference between values that is allowed by their binary representation. The following lemma relates the BBFS to the MMS.

Lemma 4. *For any integer $n \geq 1$ and agent $i \in [n]$,*

$$BBFS_i^n \geq MMS_i^{1\text{-out-of-}\lceil\frac{3}{2}n\rceil}(M).$$

Proof. Let $t_i := MMS_i^{1\text{-out-of-}\lceil\frac{3}{2}n\rceil}(M)$. By definition of MMS, there is a partition of M into $\lceil\frac{3}{2}n\rceil$ bundles of size at least t_i . By Lemma 3, the Bidirectional-Bag-Filling algorithm with valuation v_i and bin-size t_i fills at least $\frac{2}{3}(\lceil\frac{3}{2}n\rceil - 1)$ bundles, which means at least n bundles

8. There is a minor difference: we initialize the first bag with only a single good from the left (g_1) before filling it with goods from the right (g_m, g_{m-1}, \dots). In contrast, Csirik et al. (1999) fill the first bag with several goods from the left (g_1, g_2, \dots while its value is less than the bin size), and only then start filling it with goods from the right. However, this difference is not substantial: their proof of the approximation ratio assumes only that each bin has at least one good from the left and one good from the right, so the same proof holds for our variant.

since the number of bundles is an integer. By definition of the BBFS, since Algorithm 3 allocates at least n bundles with threshold t_i , we have $t_i \leq \text{BBFS}_i^n$. \square

We define an allocation as *BBFS-fair* if it allocates to each agent $i \in [n]$ a bundle with a value of at least BBFS_i^n . Lemma 4 indicates that a BBFS-fair allocation is also 1-out-of- $\lceil 3n/2 \rceil$ MMS-fair, though the BBFS may be larger than 1-out-of- $\lceil 3n/2 \rceil$ MMS.

Lemma 5. *A BBFS-fair allocation always exists, and can be found in time polynomial in the length of the binary representation of the problem.*

Proof. We first show that, when Algorithm 3 is executed with threshold values $t_i = \text{BBFS}_i^n$ for all $i \in [n]$, it allocates n bundles. For each $j \geq 1$, denote:

- A_j — the bundle allocated at iteration j of Algorithm 3 with the true (different) valuations v_1, \dots, v_n .
- B_j^i — the bundle allocated at iteration j of agent i 's successful simulation with threshold $t_i = \text{BBFS}_i^n$.

We claim that, for every $k \geq 1$, the set of goods allocated before step k by the global algorithm is a subset of the goods allocated before step k during agent i 's simulation. That is, $\bigcup_{j=1}^{k-1} A_j \subseteq \bigcup_{j=1}^{k-1} B_j^i$ for any remaining agent i .

The claim is proved by induction on k . The base is $k = 1$. Before step 1, both $\bigcup_{j=1}^{k-1} A_j$ and $\bigcup_{j=1}^{k-1} B_j^i$ are empty, so the claim holds vacuously. Let $k \geq 1$. We assume the claim is true before iteration k , and prove that it is still true after iteration k . The initial goods g_1, \dots, g_k are obviously allocated in both runs. In agent i 's simulation, some additional goods g_m, \dots, g_s are allocated, for some $s \leq m$; in the global run, goods g_m, \dots, g_r are allocated, for some $r \leq m$. The induction assumption implies that $r \geq s$ (weakly fewer goods are allocated in the global run). In iteration k , both runs initialize the bag with the same good g_k . In i 's simulation, the bag is then filled with goods $g_{s-1}, \dots, g_{s'}$ for some $s' < s$, such that $v_i(\{g_k, g_{s-1}, \dots, g_{s'}\}) \geq t_i$. In the global run, the bag is filled with goods $g_{r-1}, \dots, g_{r'}$ for some $r' < r$. It is sufficient to prove that $r' \geq s'$. Indeed, if no agent takes the bag until it contains the goods $(\{g_k, g_{r-1}, \dots, g_{r'}\})$, then because $r \geq s$, the bag value is at least $v_i(\{g_k, g_{s-1}, \dots, g_{s'}\}) \geq t_i$. Therefore, it is acceptable to agent i , so the algorithm allocates it (either to i or to another agent).

This completes the proof of the claim. The claim implies that, as long as $k < n$, the goods in B_n^i are still available. This means that agent i values the remaining goods at least t_i . This is true for every remaining agent; therefore, the global algorithm continues to run until it allocates n bundles.

The binary search and the simulation runs for each agent i take time polynomial in the length of the binary representation of the valuations. Once the thresholds are computed, Algorithm 3 obviously runs in polynomial time. This completes the proof of the lemma. \square

Lemmas 4 and 5 together imply:

Theorem 2. *There is an algorithm that computes a 1-out-of- $\lceil 3n/2 \rceil$ MMS allocation in time polynomial in the length of the binary representation of the problem.* \square

Example 7 (Computing thresholds). Consider a setting with $m = 6$ goods and $n = 3$ agents with the following valuations:

	g_1	g_2	g_3	g_4	g_5	g_6	t_i
v_1	10	8	6	3	2	1	9
v_2	12	7	6	5	4	2	11
v_3	9	8	7	4	3	1	10

Each player computes a threshold via binary search on $[0, v_i(M)]$ for the maximum value t_i such that the simulation of Algorithm 3 yields three bundles. For agent 1, the simulation with $t_1 = 9$ yields bundles $\{g_1\}, \{g_2, g_6\}, \{g_3, g_4, g_5\}$. The corresponding simulation with $t_1 = 10$ yields bundles $\{g_1\}, \{g_2, g_5, g_6\}$ with $\{g_3, g_4\}$ insufficient to fill a third bundle.

After all thresholds have been determined from simulations, Algorithm 3 computes the circled allocation. Theorem 2 guarantees that this allocation is at least 1-out-of-5 MMS. Here the circled allocation satisfies 1-out-of-3 MMS. ■

Remark 2. When n is odd, there is a gap of 1 between the existence result for 1-out-of- $\lfloor 3n/2 \rfloor$ MMS, and the polynomial-time computation result for 1-out-of- $\lceil 3n/2 \rceil$ MMS.

In experimental simulations on instances generated uniformly at random, Algorithm 3 significantly outperforms the theoretical guarantee of 1-out-of- $\lceil 3n/2 \rceil$ MMS. In Appendix C, we provide detailed experimentations and compare the bidirectional bag-filling algorithm with other bag-filling methods (e.g. the unidirectional bag-filling algorithm).

5.2 The case when $\ell > 1$

So far, we could not adapt Algorithm 3 to finding an ℓ -out-of- $\lfloor (\ell + 1/2)n \rfloor$ MMS allocation for $\ell \geq 2$. Below, we present a weaker approximation to MMS, based on the following lemma.

Lemma 6. *For all integers $d > \ell \geq 1$:*

$$\ell \cdot \text{MMS}^{1\text{-out-of-}d}(M) \leq \text{MMS}^{\ell\text{-out-of-}d}(M) \leq \ell \cdot \text{MMS}^{1\text{-out-of-}(d-\ell+1)}(M).$$

Proof. For the leftmost inequality, Let A_1, \dots, A_d be the optimal d -partition in the definition of $\text{MMS}^{1\text{-out-of-}d}(M)$, and suppose w.l.o.g. that the bundles are ordered by ascending value. Then:

$$\begin{aligned} \ell \cdot \text{MMS}^{1\text{-out-of-}d}(M) &= \ell \cdot v_i(A_1) \\ &\leq v_i(A_1) + \dots + v_i(A_\ell) \\ &\leq \text{MMS}^{\ell\text{-out-of-}d}(M), \end{aligned}$$

where the last inequality follows from the existence of a d -partition in which the ℓ least-valuable bundles are A_1, \dots, A_ℓ .

For the rightmost inequality, let B_1, \dots, B_d be the optimal d -partition in the definition of $\text{MMS}^{\ell\text{-out-of-}d}(M)$, and suppose w.l.o.g. that the bundles are ordered by ascending value.

Then:

$$\begin{aligned}\text{MMS}^{\ell\text{-out-of-}d}(M) &= v_i(B_1) + \cdots + v_i(B_\ell) \\ &\leq \ell \cdot v_i(B_\ell) \\ &\leq \ell \cdot \text{MMS}^{1\text{-out-of-}(d-\ell+1)}(M),\end{aligned}$$

where the last inequality is proved by the partition with $(d - \ell + 1)$ bundles: $B_1 \cup \cdots \cup B_\ell, B_{\ell+1}, \dots, B_d$, in which the value of each bundle is at least $v_i(B_\ell)$. \square

For any positive integer d , we can approximate $\text{MMS}^{1\text{-out-of-}d}(M)$ by using an approximation algorithm for bin-covering, which we call *Algorithm JS* (Jansen & Solis-Oba, 2003).

Lemma 7 (Jansen and Solis-Oba (2003)). *For any $\varepsilon > 0$, Algorithm JS runs in time $\tilde{O}\left(\frac{1}{\varepsilon^6}m^2 + \frac{1}{\varepsilon^{8.76}}\right)$.⁹ If the sum of all valuations is at least $13t/\varepsilon^3$ (where t is the bin size), then Algorithm JS fills at least $(1 - \varepsilon) \cdot \text{OPT}(I) - 1$ bins.*

We can choose ε based on the instance, and get the following simpler guarantee.

Lemma 8. *Algorithm JS fills at least*

$$\text{OPT} - 2.35 \cdot \text{OPT}^{2/3} - 1$$

bins, and runs in time $\tilde{O}(m^4)$.

Proof. If any input value is at least t , then it can be put in a bin of its own, and this is obviously optimal. So we can assume w.l.o.g. that all input values are smaller than t .

Let s be the sum of values, and set $\varepsilon := (13t/s)^{1/3}$. The number of bins in any legal packing is at most s/t , so

$$\begin{aligned}\text{OPT} &\leq s/t \\ t/s &\leq 1/\text{OPT} \\ \varepsilon &\leq (13/\text{OPT})^{1/3} \\ &\approx 2.35/\text{OPT}^{1/3}.\end{aligned}$$

The ε is chosen such that $s = 13t/\varepsilon^3$. So by Lemma 7, the number of bins filled by Algorithm JS is at least

$$\begin{aligned}\text{OPT} - \varepsilon \cdot \text{OPT} - 1 \\ \geq \text{OPT} - 2.35 \cdot \text{OPT}^{2/3} - 1.\end{aligned}$$

Since by assumption each value is smaller than t , we have $s < mt$, so $\varepsilon > (13/m)^{1/3}$ and $1/\varepsilon \in O(m^{1/3})$. Therefore, the run-time is in

$$\begin{aligned}\tilde{O}\left(\frac{1}{\varepsilon^6}m^2 + \frac{1}{\varepsilon^{8.8}}\right) \\ \approx \tilde{O}(m^2 \cdot m^2 + m^{2.92}) \\ \approx \tilde{O}(m^4).\end{aligned}$$

\square

9. A more exact expression for the run-time is $O\left(\frac{1}{\varepsilon^6} \cdot \ln \frac{m}{\varepsilon} \cdot \max(m^2, \frac{1}{\varepsilon} \ln \ln \frac{1}{\varepsilon^3}) + \frac{1}{\varepsilon^4} \mathcal{T}_M(\frac{1}{\varepsilon^2})\right)$, where $\mathcal{T}_M(x)$ is the run-time complexity of the best available algorithm for matrix inversion, which is currently $O(x^{2.38})$. We simplified it a bit for clarity, and used \tilde{O} to hide the logarithmic factors.

Analogously to the definition of the BBF-Sum (BBFS) (Definition 5.1), we define the 1-out-of- d *JS-Sum* (JSS) of each agent i as the largest value t_i for which Algorithm JS fills at least d bins when executed with valuation v_i and bin-size t_i . The JSS can be computed up to any desired accuracy using binary search.

Clearly, $\text{MMS}_i^{1\text{-out-of-}d}(M) \geq \text{JSS}_i^d$. Analogously to Lemma 4, we have

Lemma 9. *For any integer $d \geq 1$ and agent $i \in [n]$:*

$$\text{JSS}_i^d \geq \text{MMS}_i^{1\text{-out-of-}\lceil d+15 \cdot d^{2/3} + 1 \rceil}(M).$$

Proof. Let $t_i := \text{MMS}_i^{1\text{-out-of-}\lceil d+15 \cdot d^{2/3} + 1 \rceil}(M)$. By definition of MMS, there is a partition of M into $\lceil d + 15 \cdot d^{2/3} + 1 \rceil$ bundles of size at least t_i . By Lemma 8, Algorithm JS with bin-size t_i fills at least

$$\begin{aligned} & (d + 15 \cdot d^{2/3} + 1) - 2.35 \cdot (d + 15 \cdot d^{2/3} + 1)^{2/3} - 1 \\ & \geq (d + 15 \cdot d^{2/3} + 1) - 2.35 \cdot (16d)^{2/3} - 1 \\ & \geq (d + 15 \cdot d^{2/3} + 1) - 14.92d^{2/3} - 1 \\ & \geq d \end{aligned}$$

bins. By definition of the JSS, since Algorithm JS allocates at least d bins with size t_i , we have $t_i \leq \text{JSS}_i^d$. \square

Now, each agent can participate in the algorithm of Section 4 without computing the exact MMS value. Given an integer $\ell \geq 2$, let $d := \lfloor (\ell + \frac{1}{2})n \rfloor$. Each agent i can compute the value of JSS_i^d using binary search. The search also finds a partition of M into d bundles, each of which has a value of at least JSS_i^d . The agent can now use this partition for scaling: the valuations are scaled such that the value of each bundle in the partition is exactly 1. The algorithm in Section 2 guarantees to each agent a bundle with value at least ℓ , which is at least $\ell \cdot \text{JSS}_i^d$. By Lemma 9, this value is at least $\ell \cdot \text{MMS}_i^{1\text{-out-of-}\lceil d+15d^{2/3}+1 \rceil}(M)$. By the right-hand side of Lemma 6, it is at least $\text{MMS}_i^{\ell\text{-out-of-}\lceil d+15d^{2/3}+\ell \rceil}(M)$. Thus, we have proved the following theorem.

Theorem 3. *Let $\ell \geq 2$ an integer, and $d := \lfloor (\ell + \frac{1}{2})n \rfloor$. It is possible to compute an allocation in which the value of each agent i is at least*

$$\text{MMS}_i^{\ell\text{-out-of-}\lceil d+15d^{2/3}+\ell \rceil}(M),$$

in time $\tilde{O}(n \cdot m^4)$.

6. Future Directions

The existence of tighter ordinal approximations that improve ℓ -out-of- $\lfloor (\ell + 1/2)n \rfloor$ MMS allocations is a compelling open problem. Specifically, one can generalize the open problem raised by Budish (2011) and ask, for any $\ell \geq 1$ and $n \geq 2$: does there exist an ℓ -out-of- $(\ell n + 1)$ MMS allocation?

For the polynomial-time algorithm when $\ell = 1$, we extend the bin covering algorithm of Csirik et al. (1999). We believe that the interaction between this problem and fair allocation of goods may be of independent interest, as it may open new ways for developing improved algorithms. For example, Csirik et al. (1999) also present a $3/4$ approximation algorithm for bin covering, which may potentially be adapted to yield a 1-out-of- $\lceil 4n/3 \rceil$ MMS allocation. Similarly, Csirik, Johnson, and Kenyon (2001) and Jansen and Solis-Oba (2003) present polynomial-time approximation schemes for bin covering, which may yield even better MMS approximations in future work.

Finally, it is interesting to study ordinal maximin approximation for items with non-positive valuations (i.e. chores), as well as for mixtures of goods and chores. Techniques for allocation of goods do not immediately translate to achieving approximations of MMS when allocating chores, so new techniques are needed (Hosseini, Searns, & Segal-Halevi, 2022).

Acknowledgments

Hadi Hosseini acknowledges support from NSF IIS grants #2052488 and #2107173. Erel Segal-Halevi is supported by the Israel Science Foundation (grant no. 712/20).

We are grateful to Thomas Rothvoss, Ariel Procaccia, Joshua Lin, Inuyasha Yagami, Chandra Chekuri, Neal Young, and the anonymous referees of EC 2021 and JAIR for their valuable feedback.

APPENDIX

Appendix A. Comparing Ordinal and Multiplicative Approximations

Our ordinal guarantees may be better than the best known multiplicative MMS approximation (i.e. $3/4$) when the number of goods is large compared to the number of agents.

To illustrate, consider the extreme case in which there are infinitely many goods of equal value (alternatively, suppose there are infinitely many goods with values that are independent and identically-distributed random variables). Then 1-out-of- n MMS converges to $1/n$ (with probability 1, by the law of large numbers).¹⁰ The ℓ -out-of- $(\ell + \frac{1}{2})n$ MMS converges to $2\ell/(2\ell + 1)$ of this value, which is larger than $3/4 + 1/(12n)$ for $\ell \geq 2$, and approaches 1 when $\ell \rightarrow \infty$.

In this section, we present a simple simulation experiment that compares the value of the ℓ -out-of- d MMS guaranteed by Theorem 1 (where $d = \lfloor (\ell + 1/2)n \rfloor$) with the best known multiplicative approximation of 1-out-of- n MMS, which is $\frac{3}{4} + \frac{1}{12n}$ (Garg & Taki, 2020). Our results show that the ordinal approximation for $\ell \geq 2$ is better than the multiplicative approximation already for $m \approx 20n$, when the values are sampled from some natural distributions. We note that the simulations only compare the worst-case guarantees and not the actual algorithm performance.

Since computing the exact MMS is NP-hard,¹¹ we used a lower bound for our ordinal approximation and an upper bound for the ‘‘competition’’, so that our ratio is a lower bound for the real ratio. For our ordinal approximation, we computed a lower bound using the *greedy number partitioning* algorithm (Graham, 1966, 1969). This algorithm is known to attain a reasonable approximation of the maximin share both in the worst case (Deuermeyer, Friesen, & Langston, 1982; Csirik, Kellerer, & Woeginger, 1992) and in the average case (Frenk & Kan, 1986). Given an integer d , the algorithm initializes d empty bundles. It iterates over the goods in descending order of their value, and puts the next good in the bundle with the smallest total value so far (breaking ties arbitrarily). Once all goods are allocated, the sum of values in the ℓ bundles with the smallest values is a lower bound for the ℓ -out-of- d MMS. Taking instead the smallest value times ℓ (which approximates $\ell \cdot \text{MMS}^{\text{1-out-of-}d}$) yields nearly identical results. For the multiplicative approximation, we just use the proportional share $v_i(M)/n$ as an upper bound for agent i ’s 1-out-of- n MMS.

For various values of m , we chose m random integers to use as the good values. We performed three simulations, in which the values were distributed (a) uniformly at random in $[1, 1000]$, (b) uniformly at random in $[1000, 2000]$, and (c) geometrically with mean value of 1000. We modified n between 4 and 20, and m between $4n$ and $80n$. The results for all n were very similar. While our approximation for $\ell = 1$ is generally worse than $3/4$ of the MMS, our approximation for $\ell \geq 2$ is better already for $m \approx 20n$, and it becomes better as m grows. Figure 3 illustrates these observations.¹²

10. An accurate computation of the convergence rate of the MMS to $1/n$ is beyond the scope of the present paper. We refer the interested reader to Mertens (2001), who studies a closely-related problem: the probability distribution of the smallest *difference* between the highest-valued bundle and the lowest-valued bundle in an n -partition.
11. Using integer linear programming, we could compute the exact value of 1-out-of-4 MMS for $m > 200$ goods in reasonable time. However, we were not able to scale our computations for larger values of n .
12. Source code for the experiments is available at <https://github.com/erelsgl/ordinal-maximin-share>.

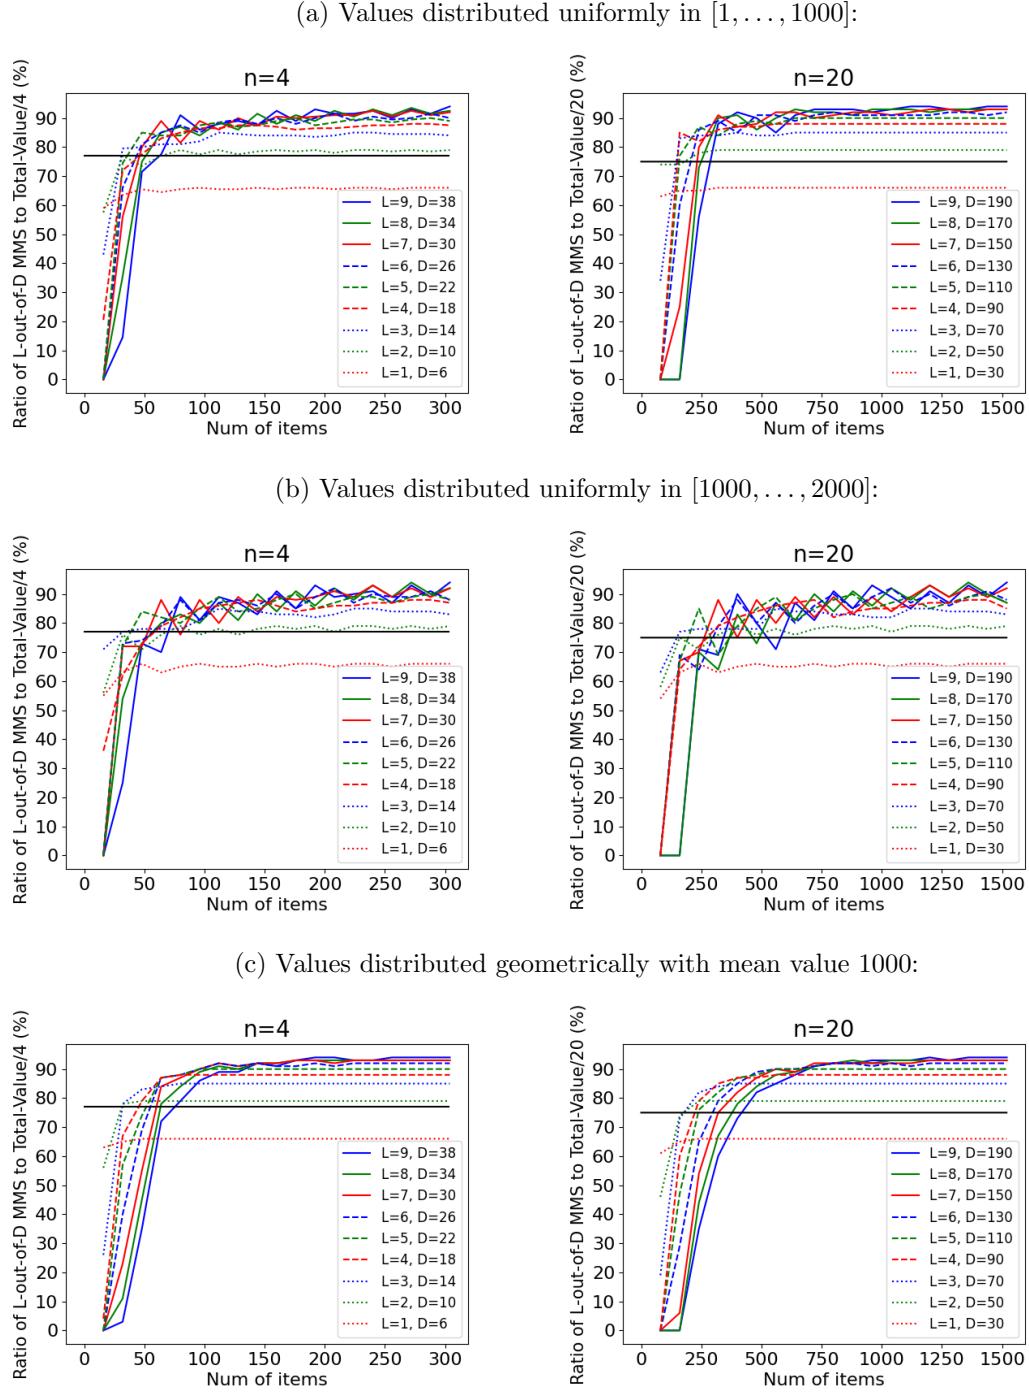


Figure 3: The value of the ℓ -out-of- d MMS (where $d = (\ell + 1/2)n$) guaranteed by our Theorem 1, as a fraction of the 1-out-of- n MMS, for different values of ℓ and n (at the left $n = 4$ and at the right $n = 20$). The horizontal black line represents $3/4 + 1/(12n)$ of the 1-out-of- n MMS

Appendix B. Failure of Some Common Techniques for Approximate-MMS Allocation

Our 1-of- $\lfloor 3n/2 \rfloor$ MMS guarantee seems very similar to multiplicative $2/3$ -MMS approximation, as both can be seen as approximations of $2/3 \cdot (v_i(M)/n)$. One could expect that the same techniques should work in both cases. To illustrate that this is not the case, we consider one such technique, recently used by Garg et al. (2018) to find a $2/3$ -MMS allocation in polynomial time. For completeness, we briefly describe their algorithm below (the detailed steps can be found in Garg et al. (2018)):

1. Scale the valuations such that all n agents value the set of all goods at $3n/2$ (this implies that their MMS value is at most $3/2$).
2. Order the instance such that $v_i(g_1) \geq \dots \geq v_i(g_m)$ to all agents i .
3. If an agent i values g_1 by at least 1, then allocate g_1 to i and recurse with the remaining goods and $n - 1$ agents.
4. If an agent i values $\{g_n, g_{n+1}\}$ by at least 1, then allocate $\{g_n, g_{n+1}\}$ to i and recurse with the remaining goods and $n - 1$ agents.
5. At this point, all agents value g_1, \dots, g_n at less than 1, and g_{n+1}, \dots, g_m at most $1/2$. Allocate the goods using bag-filling, initializing each bag $j \in [n]$ with the good g_j .

They prove that the reductions in Steps (3) and (4) above do not change the MMS value of the remaining agents. While the former reduction is valid regardless of the number of MMS bundles, the latter reduction crucially depends on the fact that there are n MMS bundles, which implies (by the pigeonhole principle) that at least one MMS bundle contains at least two goods from $\{g_1, \dots, g_{n+1}\}$. This no longer holds when there are $3n/2$ MMS bundles. Therefore, allocating goods $\{g_n, g_{n+1}\}$ is no longer a valid reduction. Moreover, if we scale the valuations as in Step (1), we may be unable to give each agent a value of at least 1.

Example 8. Suppose $n = 20$, there are 30 goods with value $1 - \varepsilon$ and one good with value 30ε . The instance is already normalized, since $v_i(M) = 30 = 3n/2$ for all $i \in N$. The value of all goods is less than 1. However, the value of $\{g_n, g_{n+1}\}$ is more than 1, and if we allocate such pairs of goods to agents, at most 15 agents will receive a bundle. This example shows that the threshold of $2v_i(M)/(3n)$ might be too high (too wasteful) for this problem. ■

Note that the challenging case in Section 4.7 (Case #3, for $\ell = 1$) is exactly the case in which more than n MMS bundles contain high-value goods, and each such good is contained in a unique MMS bundle.

Appendix C. Bidirectional vs Unidirectional Bag-Filling

Theorem 2 guarantees that the bidirectional bag-filling algorithm with a careful set up of thresholds (Algorithm 3) will compute an allocation which satisfies at least 1-of- $(3n/2 + 1)$ MMS. We empirically evaluated this algorithm by selecting different ways to set agent's thresholds. As a baseline, we use a simpler, unidirectional bag-filling algorithm, where

the goods are put into bags in decreasing order of their value. For each version we ran two sets of experiments: (i) one experiment where agents threshold values are computed individually, and (ii) one experiment where all agents' thresholds are a common fraction of their proportional share.

In each experiment, we generated 1,000 instances for each pair of $n \in [3, 20]$ and $m \in [n, 100]$ ¹³ where the valuations were uniformly distributed from [0, 1000] and then ordered.

C.1 Individual Thresholds

For each agent, we utilize binary search to find the largest individual threshold where that agent can form at least n bundles in successful simulations, as explained in Section 5. To compare the bidirectional and unidirectional bag-filling approaches, we first compute the value that each agent received, as a fraction of his proportional share $Prop_i := v_i(M)/n$. We then plot the minimum ratio any agent received over all 1,000 instances, the average ratio all agents received over 1,000 instances, and the minimum of the average ratios of all agents per instance.

If the valuations were perfectly divisible (say as $m \rightarrow \infty$), we would expect that the 1-of-($3n/2 + 1$) MMS would equate to approximately $\frac{v_i(M)}{3n/2} \approx 2/3 \frac{v_i(M)}{n} = 2/3 Prop_i$. Since proportionality implies MMS, it is an upper bound for MMS values, each agent's 1-of-($3n/2 + 1$) MMS is at most $2/3 Prop_i$.

Figure 4 and Figure 5 show that, while both the bidirectional and unidirectional algorithms exceed the $2/3$ ratio on most instances, the bidirectional algorithm averages slightly higher (about 5%), with a higher minimum average. Note that the unidirectional bag-filling algorithm is not guaranteed to give every agent 1-of-($3n/2 + 1$) MMS.

13. Observe that when $m < 3n/2 + 1$, the 1-of-($3n/2 + 1$) MMS is 0 for all agents. Thus any allocation satisfies this property.

Thresholds computed individually:

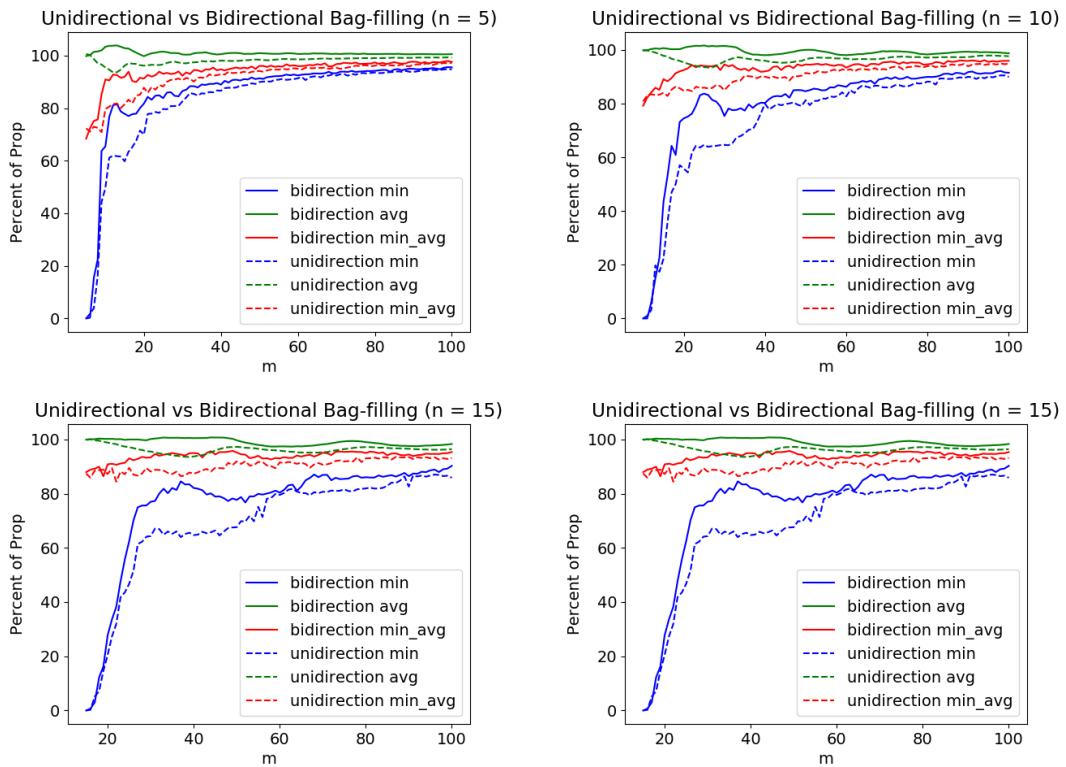


Figure 4: The minimum, average, and minimum average (minimum over all instances of the average value of agents within the instance) value received by agents in both bidirectional and unidirectional bag-filling for various n and m . Bidirectional bag-filling always exceeded its unidirectional counterpart.

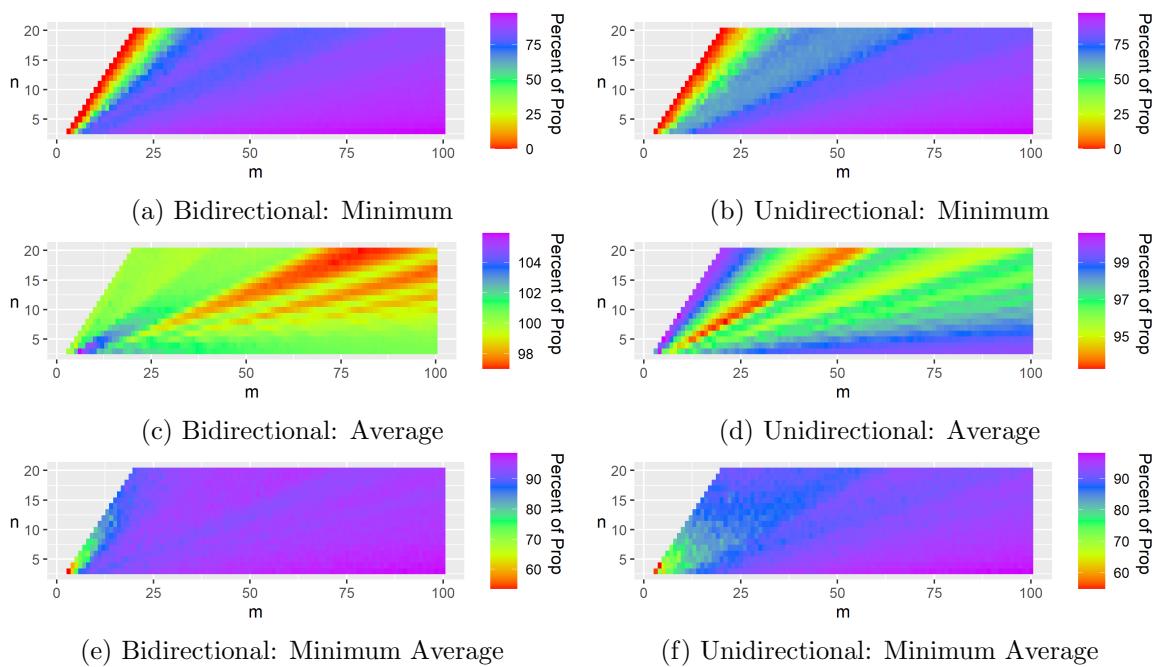


Figure 5: In each figure, the x axis is the number of goods (m) and the y axis is the number of agents (n). The color represents the fraction of proportional allocations (in percent).

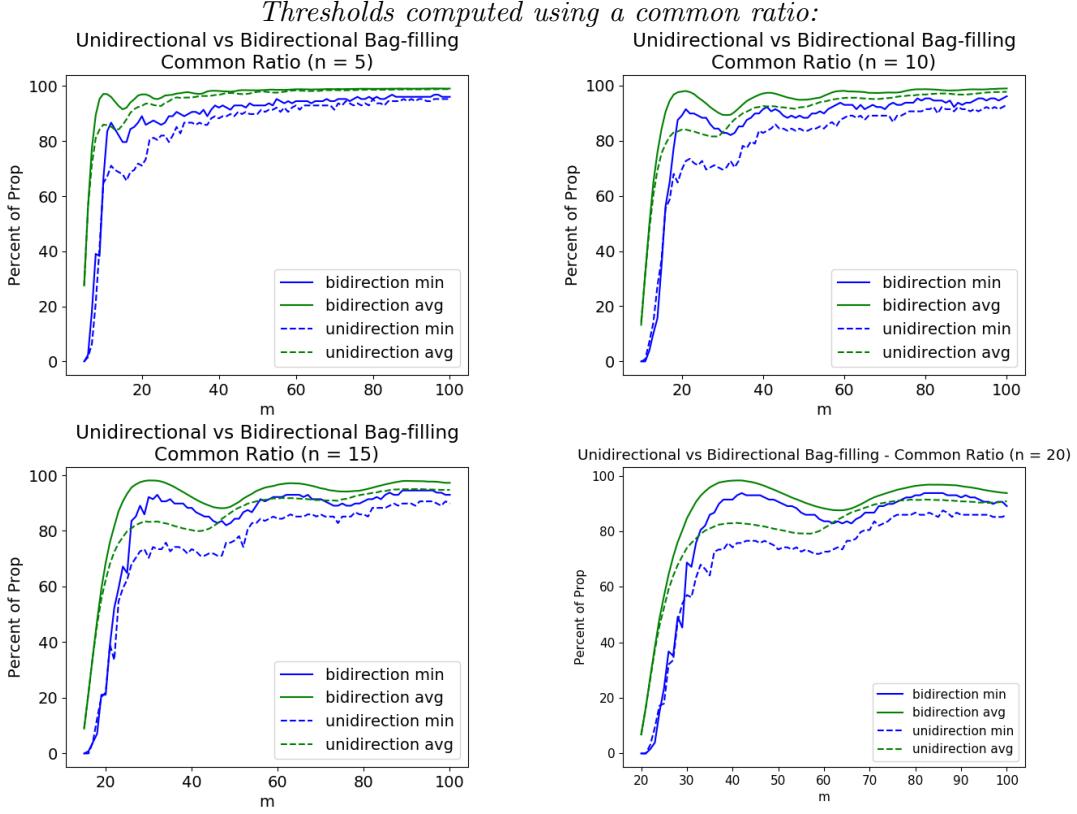


Figure 6: The minimum, average, and minimum average (minimum over all instances of the average value of agents within the instance) value received by agents in both bidirectional and unidirectional bag-filling for various n and m . Bidirectional bag-filling always exceeded its unidirectional counterpart.

C.2 Common Thresholds

We run binary search to find the largest percentage t where the bidirectional bag-filling algorithm gives each agent at least $t \frac{v_i(M)}{n}$. The ratio was computed up to an error tolerance of 0.1%. For this experiment we tracked the minimum ratio and the average ratio guaranteed to all agents across the 1,000 instances.

We observe that the bidirectional bag-filling algorithm outperforms the unidirectional bag-filling algorithm, especially when the number of goods is small relative to the number of agents ($m < 2n$). Interestingly, Figure 6 illustrates that bidirectional bag-filling achieves a higher fraction of proportionality with respect to all measures (minimum, average, and minimum average over all instances) compared to the unidirectional bag-filling algorithm.

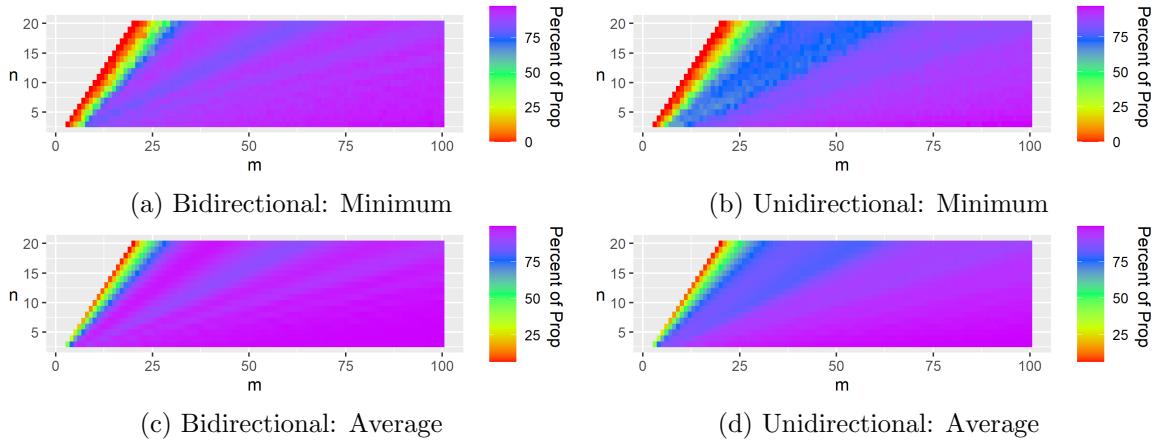


Figure 7: In each figure, the x axis is the number of goods (m) and the y axis is the number of agents (n). The color represents the fraction of proportional allocations (in percent).

Appendix D. Beyond Additive Valuations

In this section, we show that no meaningful ordinal MMS approximation is possible when generalizing additive valuations to *responsive* valuations (defined below), even for two agents. Moreover, this result implies that ordinal MMS approximations cannot be extended to *submodular* valuations.

The class of *responsive valuations* was introduced in the literature on matching markets (Roth, 1986; Alkan, 1999; Klaus & Klijn, 2005; Hatfield, 2009) and has recently been used in fair item allocation (Aziz, Biro, Lang, Lesca, & Monnot, 2019; Kyropoulou, Suksompong, & Voudouris, 2020; Babaioff et al., 2021). Formally, a valuation $v : 2^M \rightarrow \mathbb{R}_+$ is called *responsive* if for any two goods $x, y \in M$ and any subset $Z \subseteq M$:

$$v(x) \leq v(y) \iff v(Z) \leq v(Z \cup \{x\}) \leq v(Z \cup \{y\}).$$

Responsive valuations have several equivalent definitions (Aziz, Gaspers, Mackenzie, & Walsh, 2015). One of them uses the notion of *domination*. Given a valuation v on individual goods, a bundle X is *dominated by* a bundle Y , denoted $X \precsim_v Y$, if there is an injection $f : X \rightarrow Y$ such that $v(x) \leq v(f(x))$ for all $x \in X$. The valuation v is called *responsive* if

$$X \precsim_v Y \implies v(X) \leq v(Y).$$

Intuitively, responsive valuations presume that agents rank individual goods, and that their ranking of bundles is consistent with the ranking of goods.

If v is additive then it is clearly responsive, but the opposite is not necessarily true. For example, suppose there are four goods with $v(w) < v(x) < v(y) < v(z)$. Then the responsiveness condition does not imply anything about the relation between $\{z\}$ and $\{x, y\}$, since none of them dominates the other. Similarly, responsiveness does not imply anything regarding $\{w, z\}$ and $\{w, x, y\}$, since none of them dominates the other. So it is possible that $v(\{z\}) < v(\{x, y\}) < v(\{w, x, y\}) < v(\{w, z\})$. This is impossible with additive valuations.

The technique of picking-sequences, used for un-ordering an instance (see Section 4.2), works for responsive preferences too. In fact, for each agent i , the bundle picked by i during

the picking-sequence dominates the bundle allocated to i by the algorithm on the ordered instance (the picking-sequence implements the injection f). Similarly, Lemma 5, ensuring the existence of a BBFS-fair allocation, holds for responsive preferences too, since the proof only requires containment of bundles. Therefore, the class of responsive preferences initially seems like a good candidate for generalizing our results.

Unfortunately, we show below that, with responsive valuations, no meaningful ordinal approximation is possible, even for two agents. This indicates that the ordinal maximin-share approximation may be too strong for handling non-additive valuations.

Proposition 1. *For any integer $d \geq 1$, there is an instance with two agents with responsive valuations, in which no allocation guarantees both agents their 1-out-of- d maximin-share.*

Proof. We construct an instance with $m = 2^{d^2} - 1$ goods. They are ranked the same for both agents: $v_i(g_1) > \dots > v_i(g_m)$ for $i \in \{1, 2\}$.

For each $j \in [d^2]$, we denote by B_j the set of bundles that contain a majority of the goods in $\{g_1, \dots, g_{(2j)-1}\}$. Note that $B_j \subseteq 2^M$. For example:

- B_1 is the set of all bundles that contain g_1 ;
- B_2 is the set of all bundles that contain at least two goods from $\{g_1, g_2, g_3\}$;
- B_3 is the set of all bundles that contain at least four goods from $\{g_1, \dots, g_7\}$;
- B_{d^2} is the set of all bundles that contain at least 2^{d^2-1} goods from M .

Let X, Y be some bundles such that $X \in B_j$ and $Y \in 2^M \setminus B_j$ for some $j \in [d^2]$. The majority assumption implies that X cannot be dominated by Y : there cannot be an injection from a majority to a minority. Therefore, a responsive valuation may assign a larger value to X than to Y .

Before proceeding with the proof, we exemplify it for the special case $d = 2$. We define the valuations of two agents as follows.

- For agent 1, we set $v_1(X) = 1$ for any $X \in (B_1 \cap B_2) \cup (B_3 \cap B_4)$, and smaller values for other bundles. Then $\text{MMS}_1^{\text{1-out-of-}d} = 1$, since M can be partitioned into $\{g_1, g_2, g_3\} \in B_1 \cap B_2$ and $\{g_4, \dots, g_7, g_8, \dots, g_{15}\} \in B_3 \cap B_4$.
- For agent 2, we set $v_2(X) = 1$ for any $X \in (B_1 \cap B_3) \cup (B_2 \cap B_4)$, and smaller values for other bundles. Then $\text{MMS}_2^{\text{1-out-of-}d} = 1$, since M can be partitioned into $\{g_1, g_4, \dots, g_7\} \in B_1 \cap B_3$ and $\{g_2, g_3, g_8, \dots, g_{15}\} \in B_2 \cap B_4$.

Note that the agents' valuations are consistent with responsiveness. For example, consider a bundle $Y \notin (B_1 \cap B_2) \cup (B_3 \cap B_4)$. Then either $Y \notin B_1 \cup B_3$ or $Y \notin B_1 \cup B_4$ or $Y \notin B_2 \cup B_3$ or $Y \notin B_2 \cup B_4$. In any case, Y cannot dominate any bundle $X \in (B_1 \cap B_2) \cup (B_3 \cap B_4)$. So assigning to X a higher value than to Y is consistent with responsive valuations.

Suppose now that an allocation (A_1, A_2) gives agent 1 a value of at least 1. This means that either $A_1 \in B_1 \cap B_2$ or $A_1 \in B_3 \cap B_4$. If $A_1 \in B_1 \cap B_2$, then A_1 contains g_1 and a majority of the goods from $\{g_1, g_2, g_3\}$. This means that A_2 cannot contain g_1 and cannot contain a majority of $\{g_1, \dots, g_3\}$. So $A_2 \notin B_1$ and $A_2 \notin B_2$. This means that $A_2 \notin (B_1 \cap B_3) \cup (B_2 \cap B_4)$, so the value of agent 2 is less than $\text{MMS}_2^{\text{1-out-of-}d}$. Similarly,

if $A_1 \in B_3 \cap B_4$, then $A_2 \notin B_3$ and $A_2 \notin B_4$, so again the value of agent 2 is less than $\text{MMS}_2^{1\text{-out-of-}d}$. We conclude that no allocation gives both agents their 1-out-of- d MMS.

We now generalize this construction to any integer d .

- Let $v_1(X) = 1$ for any bundle satisfying

$$X \in \bigcup_{i=1}^d \left(\bigcap_{j=1}^d B_{(i-1)d+j} \right)$$

and smaller values for other bundles.

- Let $v_2(X) = 1$ for any bundle satisfying

$$X \in \bigcup_{i=1}^d \left(\bigcap_{j=1}^d B_{(j-1)d+i} \right)$$

and smaller values for other bundles.

Note that agent 1's valuations are consistent with responsiveness, since any bundle not in $\bigcup_{i=1}^d \left(\bigcap_{j=1}^d B_{(i-1)d+j} \right)$ cannot dominate a bundle from $\bigcup_{i=1}^d \left(\bigcap_{j=1}^d B_{(i-1)d+j} \right)$, and similarly for agent 2.

To compute the agents' MMS, define d^2 bundles as follows. For each $j \in [d^2]$, let

$$G_j = \{g_{2(j-1)}, \dots, g_{(2j)-1}\}.$$

For example, $G_1 = \{g_1\}$, $G_2 = \{g_2, g_3\}$, $G_3 = \{g_4, \dots, g_7\}$, and so on. Note that the G_j are pairwise-disjoint, and for any $j \in [d^2]$, $G_j \in B_j$ since it contains the majority of goods in $\{g_1, \dots, g_{(2j)-1}\}$. Then:

- $\text{MMS}_1^{1\text{-out-of-}d} = 1$, by the partition with parts $P_i := \left(\bigcup_{j=1}^d G_{(i-1)d+j} \right)$ for $i \in [d]$.
- $\text{MMS}_2^{1\text{-out-of-}d} = 1$, by the partition with parts $Q_i := \left(\bigcup_{j=1}^d G_{(j-1)d+i} \right)$ for $i \in [d]$.

Suppose now that an allocation (A_1, A_2) gives agent 1 a value of at least 1. This means that $A_1 \in \left(\bigcap_{j=1}^d B_{(i_1-1)d+j} \right)$ for some $i_1 \in [d]$. Then, $A_2 \notin B_{(i_1-1)d+j}$ for any $j \in [d]$. So $A_2 \notin \left(\bigcap_{j=1}^d B_{(j-1)d+i_2} \right)$ for any $i_2 \in [d]$. So the value of agent 2 is less than 1.

We conclude that no allocation gives both agents their 1-out-of- d MMS. The proof holds for any positive integer d . \square

Remark 3. Babaioff et al. (2021) prove that responsive preferences are a subset of submodular preferences. Every submodular preference relation can be represented by a submodular valuation function. Therefore, the impossibility in Proposition 1 extends to submodular valuations too.

References

Aigner-Horev, E., & Segal-Halevi, E. (2022). Envy-free matchings in bipartite graphs and their applications to fair division. *Information Sciences*, 587, 164–187.

Alkan, A. (1999). On the properties of stable many-to-many matchings under responsive preferences. In *Current Trends in Economics*, pp. 29–39. Springer.

Amanatidis, G., Birmpas, G., & Markakis, E. (2016). On truthful mechanisms for maximin share allocations. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 31–37.

Amanatidis, G., Markakis, E., Nikzad, A., & Saberi, A. (2017). Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms (TALG)*, 13(4), 52.

Assmann, S. F., Johnson, D. S., Kleitman, D. J., & Leung, J.-T. (1984). On a dual version of the one-dimensional bin packing problem. *Journal of algorithms*, 5(4), 502–525.

Aziz, H., Biro, P., Lang, J., Lesca, J., & Monnot, J. (2019). Efficient reallocation under additive and responsive preferences. *Theoretical Computer Science*, 790, 1–15.

Aziz, H., Gaspers, S., Mackenzie, S., & Walsh, T. (2015). Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227, 71–92.

Aziz, H., & Ye, C. (2014). Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In *International Conference on Web and Internet Economics*, pp. 1–14. Springer.

Babaioff, M., Nisan, N., & Talgam-Cohen, I. (2019). Fair allocation through competitive equilibrium from generic incomes. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 180–180.

Babaioff, M., Nisan, N., & Talgam-Cohen, I. (2021). Competitive equilibrium with indivisible goods and generic budgets. *Mathematics of Operations Research*, 46(1), 382–403.

Barman, S., & Krishna Murthy, S. K. (2017). Approximation algorithms for maximin fair division. In *Proceedings of the ACM Conference on Economics and Computation*, pp. 647–664.

Bogomolnaia, A., & Moulin, H. (2022). Guarantees in fair division: General or monotone preferences.. arXiv preprint 1911.10009.

Bouveret, S., Chevaleyre, Y., & Maudet, N. (2016). Fair allocation of indivisible items. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., & Procaccia, A. D. (Eds.), *Handbook of Computational Social Choice*, chap. 13, pp. 284–310. Cambridge University Press.

Bouveret, S., & Lemaître, M. (2016). Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2), 259–290.

Brams, S. J., Kilgour, & Klamler, C. (2012). The undercut procedure: an algorithm for the envy-free division of indivisible items. *Soc. Choice Welf.*, 39(2-3), 615–631.

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

Csirik, J., Frenk, J. B. G., Labbè, M., & Zhang, S. (1999). Two simple algorithms for bin covering. *Acta Cybernetica*, 14(1), 13–25.

Csirik, J., Johnson, D. S., & Kenyon, C. (2001). Better approximation algorithms for bin covering. In *SODA*, Vol. 1, pp. 557–566.

Csirik, J., Kellerer, H., & Woeginger, G. (1992). The exact lpt-bound for maximizing the minimum completion time. *Operations Research Letters*, 11(5), 281–287.

Deuermeyer, B. L., Friesen, D. K., & Langston, M. A. (1982). Scheduling to maximize the minimum processor finish time in a multiprocessor system. *SIAM Journal on Algebraic Discrete Methods*, 3(2), 190–196.

Edmonds, J., & Pruhs, K. (2011). Cake cutting really is not a piece of cake. *ACM Transactions on Algorithms (TALG)*, 7(4), 1–12.

Elkind, E., Segal-Halevi, E., & Suksompong, W. (2021a). Graphical cake cutting via maximin share. In *Proceedings of the International Joint Conference on Artificial Intelligence IJCAI*.

Elkind, E., Segal-Halevi, E., & Suksompong, W. (2021b). Keep your distance: Land division with separation. In *Proceedings of the International Joint Conference on Artificial Intelligence IJCAI*.

Elkind, E., Segal-Halevi, E., & Suksompong, W. (2021c). Mind the gap: Cake cutting with separation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 5330–5338.

Frenk, J. B. G., & Kan, A. R. (1986). The rate of convergence to optimality of the LPT rule. *Discrete Applied Mathematics*, 14(2), 187–197.

Garg, J., McGlaughlin, P., & Taki, S. (2018). Approximating maximin share allocations. In *2nd Symposium on Simplicity in Algorithms (SOSA 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Garg, J., & Taki, S. (2020). An improved approximation algorithm for maximin shares. In *Proceedings of the ACM Conference on Economics and Computation*, pp. 379–380.

Ghodsi, M., HajiAghayi, M., Seddighin, M., Seddighin, S., & Yami, H. (2018). Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the ACM Conference on Economics and Computation*, pp. 539–556. ACM.

Gourvès, L., & Monnot, J. (2019). On maximin share allocations in matroids. *Theoretical Computer Science*, 754, 50–64.

Graham, R. L. (1966). Bounds for certain multiprocessing anomalies. *Bell system technical journal*, 45(9), 1563–1581.

Graham, R. L. (1969). Bounds on multiprocessing timing anomalies. *SIAM Journal on Applied Mathematics*, 17(2), 416–429.

Halpern, D., & Shah, N. (2021). Fair and efficient resource allocation with partial information. In *Proceedings of the International Joint Conference on Artificial Intelligence IJCAI*, pp. 224–230.

Hatfield, J. W. (2009). Strategy-proof, efficient, and nonbossy quota allocations. *Social Choice and Welfare*, 33(3), 505–515.

Herreiner, D., & Puppe, C. (2002). A simple procedure for finding equitable allocations of indivisible goods. *Soc. Choice Welf.*, 19(2), 415–430.

Hosseini, H., & Searns, A. (2021). Guaranteeing maximin shares: Some agents left behind. In Zhou, Z.-H. (Ed.), *Proceedings of the International Joint Conference on Artificial Intelligence IJCAI*, pp. 238–244. International Joint Conferences on Artificial Intelligence Organization. Main Track.

Hosseini, H., Searns, A., & Segal-Halevi, E. (2022). Ordinal maximin share approximation for chores. In *Proceedings of the International Conference on Autonomous Agents and MultiAgent Systems*, p. forthcoming.

Huang, X., & Lu, P. (2021). An algorithmic framework for approximating maximin share allocation of chores. In *Proceedings of the ACM Conference on Economics and Computation (EC-21)*, pp. 630–631.

Jansen, K., & Solis-Oba, R. (2003). An asymptotic fully polynomial time approximation scheme for bin covering. *Theoretical Computer Science*, 306(1-3), 543–551.

Johnson, D. S. (1973). *Near-optimal bin packing algorithms*. Ph.D. thesis, Massachusetts Institute of Technology.

Klaus, B., & Klijn, F. (2005). Stable matchings and preferences of couples. *Journal of Economic Theory*, 121(1), 75–106.

Kuhn, H. W. (1967). *On games of fair division*, pp. 29–37. Princeton University Press.

Kurokawa, D., Procaccia, A. D., & Wang, J. (2018). Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM (JACM)*, 65(2), 8.

Kyropoulou, M., Suksompong, W., & Voudouris, A. A. (2020). Almost envy-freeness in group resource allocation. *Theoretical Computer Science*, 841, 110–123.

Lang, J., & Rothe, J. (2016). Fair division of indivisible goods. In *Economics and Computation*, pp. 493–550. Springer.

Lipton, R. J., Markakis, E., Mossel, E., & Saberi, A. (2004). On approximately fair allocations of indivisible goods. In *Proceedings of the ACM conference on Electronic commerce*, pp. 125–131. ACM.

Markakis, E. (2017). Approximation algorithms and hardness results for fair division with indivisible goods. In *Trends in Computational Social Choice*, pp. 231–247. AI Access.

McGlaughlin, P., & Garg, J. (2020). Improving Nash social welfare approximations. *Journal of Artificial Intelligence Research*, 68, 225–245.

Menon, V., & Larson, K. (2020). Algorithmic stability in fair allocation of indivisible goods among two agents.. arXiv preprint 2007.15203.

Mertens, S. (2001). A physicist’s approach to number partitioning. *Theoretical Computer Science*, 265(1-2), 79–108.

Moulin, H. (1990). Uniform externalities: Two axioms for fair allocation. *Journal of Public Economics*, 43(3), 305–326.

Moulin, H. (1992). Welfare bounds in the cooperative production problem. *Games and Economic Behavior*, 4(3), 373–401.

Moulin, H. (2019). Fair division in the internet age. *Annual Review of Economics*, 11(1), 407–441.

Nguyen, N.-T., Nguyen, T. T., & Rothe, J. (2017). Approximate solutions to max-min fair and proportionally fair allocations of indivisible goods. In *Proceedings of the Conference on Autonomous Agents and MultiAgent Systems*, pp. 262–271. International Foundation for Autonomous Agents and Multiagent Systems.

Procaccia, A. D., & Wang, J. (2014). Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the ACM conference on Economics and computation*, pp. 675–692. ACM.

Roth, A. E. (1986). On the allocation of residents to rural hospitals: A general property of two-sided matching markets. *Econometrica*, 54(2), 425–427.

Searns, A., & Hosseini, H. (2020). Fairness does not imply satisfaction (student abstract). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(10), 13911–13912.

Segal-Halevi, E. (2019). The maximin share dominance relation.. arXiv preprint 1912.08763.

Segal-Halevi, E. (2020). Competitive equilibrium for almost all incomes: existence and fairness. *Autonomous Agents and Multi-Agent Systems*, 34(1), 1–50.

Truszczyński, M., & Lonic, Z. (2020). Maximin share allocations on cycles. *Journal of Artificial Intelligence Research*, 69, 613–655.

Woeginger, G. J. (1997). A polynomial-time approximation scheme for maximizing the minimum machine completion time. *Operations Research Letters*, 20(4), 149–154.