

Guaranteeing Maximin Shares: Some Agents Left Behind

Hadi Hosseini¹ and Andrew Searns²

¹Pennsylvania State University

²Rochester Institute of Technology

hadi@psu.edu, abs2157@rit.edu

Abstract

The maximin share (MMS) guarantee is a desirable fairness notion for allocating indivisible goods. While MMS allocations do not always exist, several approximation techniques have been developed to ensure that all agents receive a fraction of their maximin share. We focus on an alternative approximation notion, based on the population of agents, that seeks to guarantee MMS for a fraction of agents. We show that no optimal approximation algorithm can satisfy more than a constant number of agents, and discuss the existence and computation of MMS for all but one agent and its relation to approximate MMS guarantees. We then prove the existence of allocations that guarantee MMS for $\frac{2}{3}$ of agents, and devise a polynomial time algorithm that achieves this bound for up to nine agents. A key implication of our result is the existence of allocations that guarantee MMS ^{$\lceil 3n/2 \rceil$} , *i.e.*, the value that agents receive by partitioning the goods into $\lceil \frac{3}{2}n \rceil$ bundles, improving the best known guarantee of MMS ^{$2n-2$} . Finally, we provide empirical experiments using synthetic data.

1 Introduction

Fair division deals with the allocation of a set of resources to a set of agents in a fair manner [Brams and Taylor, 1996; Foley, 1967; Hosseini *et al.*, 2020]. One of its most notable application domains deals with the allocation of *indivisible* (and non-shareable) goods. These applications arise, for example, in dividing inheritance and dispute resolution [Brams and Taylor, 1996], task assignment or course allocation [Budish, 2011], and have been popularized in recent years due to publicly accessible platforms such as *Spliddit*.

The most desirable fairness notion, envy-freeness, requires that each agent weakly prefers his own allocation to that of all other agents. A weaker notion, called proportionality, requires that each agent receives $\frac{1}{n}$ of her valuation of all goods, where n is the total number of agents. Unfortunately, in the indivisible domain neither of these notions are guaranteed to exist: consider a single good and two interested agents. Neither envy-freeness nor proportionality can be satisfied as one agent is destined to remain empty handed.

A third fairness notion, proposed by Budish [2011], is the *maximin share* (MMS) guarantee, which can be seen as a generalization of the *cut-and-choose* protocol [Brams and Taylor, 1996]. In a nutshell, the maximin share is the value that an agent can guarantee by dividing the goods into n bundles, assuming that all other agents choose a bundle before she does. There is no reason to believe that such a bound can always be satisfied for *all* agents. It turns out that an MMS allocation does not always exist [Kurokawa *et al.*, 2018], and even when it exists, computing an MMS partition is intractable [Bouveret and Lemaître, 2016]. Thus, several approximation techniques were developed to guarantee that each agent receives a fraction of her MMS. While these techniques are compelling, they may leave a large fraction of agents without their MMS guarantee.

We propose to circumvent this obstacle by taking an orthogonal direction based on the degree of fairness in the society of agents. Our goal is to find allocations that give a large fraction of agents their maximin share guarantee, possibly leaving a small fraction of agents behind (since MMS allocations do not always exist). This approach is related to ordinal approximations of MMS and is of interest in several application domains: resources in a hospital must be distributed among tasks/procedures in a manner that a subset of critical tasks can be fully accomplished; senior college students (as a fraction of all students) must be guaranteed seats in courses to avoid graduation delays.

Unfortunately, envy-freeness and proportionality cannot be approximated in this dimension. For instance, one may be interested in minimizing the fraction of envious agents. Envy-freeness (similarly proportionality) may leave almost all but one agent unsatisfied: consider n goods and n agents with identical valuations, having $1 - (n - 1)\epsilon$ value for one good and ϵ for the rest. Any distribution of goods will leave $n - 1$ agents envious. Therefore, focusing on the maximin share guarantee we ask the following questions:

What fraction of agents can be guaranteed their MMS value? And can we compute allocations that guarantee MMS for the majority of the agents?

1.1 Our Contributions

We propose a novel fairness framework, called (α, β) -MMS, wherein α fraction of agents receive β approximations of their MMS value and investigate the interplay between the

two approximation parameters. We establish the connection between (α, β) -MMS and ordinal approximations of MMS (Proposition 1), and investigate the computational boundaries of α and β as follows:

- **Optimal MMS:** We prove the existence of a family of instances where almost all agents do not receive their MMS by any optimal-MMS allocation that aims to maximize β for all agents. Our counterexample uses a *quadratic* number of goods, where $n - 3$ agents ($n - 4$ if n is odd) do not receive their MMS in any optimal-MMS allocation (Theorem 1).

- **Computing $(\frac{n-1}{n}, \beta)$ -MMS:** We devise a polynomial-time algorithm achieving (α, β) -MMS when $\alpha = \frac{n-1}{n}$ (Theorem 2). A consequence of this result is a tight approximation for $n \leq 4$ that immediately implies an algorithm for computing MMS^{n+1} for $n < 4$ (Corollary 1).

- **Existence of $(\frac{2}{3}, 1)$ -MMS:** We prove the existence of $(\frac{2}{3}, 1)$ -MMS (Theorem 3), and provide an algorithm that achieves this bound in polynomial time for $n < 9$ (Theorem 4). A key implication of our result is the existence of allocations that guarantee $\text{MMS}^{\lceil 3n/2 \rceil}$, *i.e.*, the value an agent receives by partitioning the goods into $\lceil \frac{3}{2}n \rceil$ bundles (Corollary 2). Our result significantly improves the best known guarantee of MMS^{2n-2} [Aigner-Horev and Segal-Halevi, 2019].

On the experimental front, we show that a simplified, polynomial-time variant of our algorithm satisfies a large fraction of agents on the most ‘difficult’ instances and this fraction grows as the ratio of goods to agents increases.

1.2 Related Work

On a high level, our approach is related to notions defined to measure the degree of fairness in a society of agents whether it pertains to minimizing the maximum or sum of envy [Chevalerey *et al.*, 2007; Chen and Shah, 2018; Nguyen and Rothe, 2014], minimizing envy ratio [Lipton *et al.*, 2004], balancing the amount of envy experienced in a society, or promoting fairness through equitable allocations where agents receive the same level of utility [Schneckenburger *et al.*, 2017; Freeman *et al.*, 2019]. Another closely related line of work suggests the notion of counting instances of envious pairs—among several other plausible measures—as opposed to measuring the intensity of envy [Feldman and Kirman, 1974].

On a technical level, the non-existence of MMS allocations [Procaccia and Wang, 2014] and its intractability [Bouveret and Lemaître, 2016; Woeginger, 1997] has given rise to a number of approximation techniques. These algorithms guarantee that each agent receives an approximation of their maximin share. Recently, Nguyen *et al.* [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 i receives the value of βMMS_i . Since the number of possible partitions is finite, an optimal-MMS allocation always exists, and it is an MMS allocation if $\beta \geq 1$. The current best results guarantee $\beta \geq 2/3$ [Kurokawa *et al.*, 2018; Garg *et al.*, 2018] and $\beta \geq 3/4$ [Garg and Taki, 2020; Ghodsi *et al.*, 2018] in general, and $\beta \geq 7/8$ [Amanatidis *et al.*, 2017] and $\beta \geq 8/9$ [Gourvès and Monnot, 2019] in the case of three agents.

2 Preliminaries

Let $N = \{1, \dots, n\}$ be a set of agents and M denote a set of m indivisible goods, where $m > n$. We denote the value of agent $i \in N$ for good $g \in M$ by $v_i(g) \geq 0$. We assume that the valuation functions are *additive*; that is, for each subset $G \subseteq M$, $v_i(G) = \sum_{g \in G} v_i(g)$. An **instance** of the problem is $I = \langle N, M, V \rangle$ where V is the valuation profile of agents. 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 $i \in N$.

Definition 1 (Envy-freeness). An allocation A is *envy-free* (EF) if for every pair of agents $i, j \in N$, $v_i(A_i) \geq v_i(A_j)$. An allocation A is *envy-free up to one good* (EF1) if for every pair of agents $i, j \in N$ such that $A_j \neq \emptyset$, there exists some good $g \in A_j$ such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$. An allocation A is *envy-free up to any good* (EFX) if for every pair $i, j \in N$ such that $A_j \neq \emptyset$, for any good $g \in A_j$, $v_i(A_i) \geq v_i(A_j \setminus \{g\})$. These definitions are due to Foley [1967], Budish [2011], and Caragiannis *et al.* [2016] respectively.

Definition 2 (Maximin Share Guarantee). Let $\Pi_k(M)$ denote the set of k -partitions of M . The **k -maximin share guarantee** of agent $i \in N$ on $\Pi_k(M)$ is

$$\text{MMS}_i^k(M) = \max_{(A_1, A_2, \dots, A_k) \in \Pi_k(M)} \min_{j \in [k]} v_i(A_j),$$

where $[k] = \{1, \dots, k\}$. Intuitively, this is the minimum value that can be guaranteed if agent i partitions the goods into k bundles and chooses the least valued bundle.

An allocation $A = (A_1, \dots, A_k) \in \Pi_k(M)$ is an **MMS_i^k allocation** if and only if $\forall i \in N$, $v_i(A_i) \geq \text{MMS}_i^k$. Note that $\text{MMS}_i^n(M) \leq \frac{v_i(M)}{n}$ since proportionality implies MMS. When it is clear from the context, we write MMS_i instead of $\text{MMS}_i^n(M)$ and simply refer to it as agent i 's **MMS value**.

Definition 3 (Optimal-MMS). While an MMS allocation does not always exist, an optimal relaxation of MMS guarantees that agents receive a fraction of their MMS value [Nguyen *et al.*, 2017]. Given an instance $I = \langle N, M, V \rangle$, the optimal-MMS value is defined by

$$\lambda^*(I) = \max_{(A_1, \dots, A_n) \in \Pi_n(M)} \min_i \frac{v_i(A_i)}{\text{MMS}_i^n}.$$

By definition, an allocation that gives each agent a $\lambda^*(I)$ fraction of its $\text{MMS}_i^n(M)$ value is guaranteed to exist.

Ordered instance. An instance is *ordered* when all agents agree on the linear ordering of the goods, irrespective of their valuations. Formally, I is an *ordered instance* if there exists an ordering of goods, (g_1, g_2, \dots, g_m) such that for all agents $i \in N$ we have $v_i(g_1) \geq v_i(g_2) \geq \dots \geq v_i(g_m)$. Bouveret and Lemaître [2016] showed that ordered instances are the ‘hardest’ for achieving MMS. In fact, these instances are the only known structures for which MMS does not exist [Kurokawa *et al.*, 2018]. The next lemma states that given an *unordered* instance, it is always possible to generate a corresponding ordered instance in polynomial time. Furthermore, if the ordered instance admits an MMS allocation, the original instance also admits an MMS allocation which can be computed in polynomial-time.

Lemma 1 ([Barman and Krishna Murthy, 2017]). *Let $I' = \langle N, M, V' \rangle$ be an ordered instance constructed from the original instance $I = \langle N, M, V \rangle$. Given allocation A' on I' , a corresponding allocation A on I can be computed in polynomial time such that for all $i \in N$, $v_i(A_i) \geq v'_i(A'_i)$.*

Scale invariance. The scale invariance property of MMS states that if an agent's valuations are scaled by a factor, then its MMS value scales by the same factor.

Lemma 2 ([Ghods et al., 2018]). *Let $I = \langle N, M, V \rangle$ be an instance and $c > 0$ be a real scalar. Let $I' = \langle N, M, V' \rangle$ be constructed so that $v'_i(g) = cv_i(g)$ for all $g \in M$. Then $MMS_i^{c,k}(M) = cMMS_i^k(M)$.*

Thus, an instance $I = \langle N, M, V \rangle$ and a real value k can be scaled to form a new instance $I' = \langle N, M, V' \rangle$, such that for each agent $i \in N$, $v'_i(g) = \frac{k}{v_i(M)}v_i(g)$, and $v'_i(M) = k$.

Valid reduction. We call the act of removing a set $A_i \subseteq M$ of goods and an agent i a valid reduction if the following two conditions hold: i) $v_i(A_i) \geq MMS_i^n(M)$ and ii) $\forall j \in N \setminus \{i\}$, $MMS_j^{n-1}(M \setminus A_i) \geq MMS_j^n(M)$.

Lemma 3 ([Garg et al., 2018]). *Given an ordered instance $I = \langle N, M, V \rangle$ with $|N| = n$ such that $MMS_i^n \leq 1$, if $v_i(\{g_n, g_{n+1}\}) \geq 1$, then removing $A_i = \{g_n, g_{n+1}\}$ and agent i forms a valid reduction. Similarly, the removal of $\{g_1\}$ and agent i forms a valid reduction if $v_i(\{g_1\}) \geq 1$.*

Normalized instance. An instance is normalized if all of the following properties hold: i) the instance is ordered, ii) it is scaled so that $v_i(M) = n$, and iii) it is reduced so that $v_i(g_1) < 1$ and $v_i(\{g_n, g_{n+1}\}) < 1$. By combining Lemma 1, 2, and 3, we may assume that instances are normalized. We prove this claim in Lemma 4.

3 Approximating Maximin Share

We introduce a fairness concept that allows for interpolation between two dimensions in approximating MMS pertaining to the fraction of agents α that receive a β approximation of their maximin share.

Definition 4 ((α, β) -MMS). An allocation A guarantees (α, β) -MMS if $\alpha \in (0, 1]$ fraction of agents receive at least their $\beta \in (0, 1]$ approximation of their MMS_i^n . Formally, given an instance $I = \langle N, M, V \rangle$, an allocation A guarantees (α, β) -MMS if there exists a subset $N' \subseteq N$ with $|N'| \geq \lfloor \alpha |N| \rfloor$ such that for all $i \in N'$, $v_i(A_i) \geq \beta MMS_i^n$. We say that (α, β) -MMS exists if for any instance $I = \langle N, M, V \rangle$, for every subset $N' \subseteq N$ of agents such that $|N'| = \lfloor \alpha |N| \rfloor$, there exists an allocation A such that for all $i \in N'$, $v_i(A_i) \geq \beta MMS_i^n$.

Previous MMS approximation results can be seen in this context as efforts to tighten the approximation bound for all agents ($\alpha = 1$). Notably, greedy algorithms exist to compute $(1, \frac{2}{3})$ -MMS [Barman and Krishna Murthy, 2017; Garg et al., 2018] and $(1, \frac{3}{4})$ -MMS [Ghods et al., 2018] allocations.

Remark 1. It is crucial to highlight two key distinctions: First, in contrast to previous works [Ortega, 2018; Nyman et al., 2020], the definition of (α, β) -MMS existence does not

only hold for a fixed subset of agents. Rather, it is a stronger concept that holds for every subset of $\lfloor \alpha n \rfloor$ agents. Second, (α, β) -MMS enables a social planner to pre-select the $N' \subseteq N$ of $\lfloor \alpha n \rfloor$ agents—independent of their preferences—according to some priority ordering over the agents or by selecting the agents uniformly at random. For instance, a higher priority may be given to senior college students; a practice that is already common in most course allocation procedures.

3.1 $(\alpha, 1)$ -MMS Implies MMS^k for $k \geq \lceil \frac{n}{\alpha} \rceil$

Budish [2011] showed that approximating the competitive equilibrium from equal incomes (A-CEEI) guarantees $MMS^{n+1}(M)$, i.e. adding a dummy agent and asking all agents to partition the goods into $n + 1$ bundles. However, this result does not imply the existence of MMS^{n+1} in allocating indivisible goods because allocations achieved by A-CEEI may have excess supply or excess demands. This approach can result in infeasible allocations in fair division settings that do not allow the addition of excess goods. Therefore, the existence of MMS^k for $n + 1 \leq k \leq 2n - 2$ remains an open problem. In Proposition 1 we show the relation between MMS^k with (α, β) -MMS.

Proposition 1. *The existence of $(\alpha, 1)$ -MMS implies the existence of MMS^k for $k \geq \lceil \frac{n}{\alpha} \rceil$.*

Proof. Suppose that $(\alpha, 1)$ -MMS exists. Given an instance $I = \langle N, M, V \rangle$ with n agents, we construct an instance I' from I by adding $\lceil \frac{n}{\alpha} \rceil - n$ dummy agents. Therefore, $|N'| = \lceil \frac{n}{\alpha} \rceil$. Since $(\alpha, 1)$ -MMS exists, for every subset $N'' \subseteq N'$ of size $\lfloor \alpha |N'| \rfloor$, there exists an allocation which satisfies $(\alpha, 1)$ -MMS on that set of agents. Thus, we may choose the $\lfloor \alpha |N'| \rfloor$ agents to contain exactly the well-defined original set of agents in N , that is, $N'' := N$. Hence, each agent $i \in N$ receives at least $v_i(A_i) \geq MMS_i^{|N'|}$, which implies $MMS^{\lceil \frac{n}{\alpha} \rceil}$ for agents in N . \square

3.2 Failure of Optimal MMS Algorithms

There are two primary motivations behind the approximation parameter α (the fraction of agents). First, Proposition 1 states that fixing $\beta = 1$ immediately implies an ordinal approximation of MMS by partitioning goods into $k > n$ bundles. Second, our next theorem shows that maximizing the value of β for all agents may result in only a small fraction of agents (α) receiving their MMS value. Theorem 1 shows that there exists a family of instances where any optimal-MMS allocation only gives a small constant number of agents their MMS value. Thus, an optimal-MMS algorithm (e.g. [Nguyen et al., 2017]) will result in an (α, λ^*) -MMS allocation such that α goes to zero as the number of agents, n , increases.

Theorem 1. *For every $n \geq 4$, there exists an instance with $O(n^2)$ goods where every optimal-MMS allocation guarantees at most 3 (4 if n is odd) agents their MMS value.*

The proof is inspired by constructions proposed by Kurokawa et al. [2016; 2018], but it includes a few intricate modifications by separating agents into $\lceil \frac{n}{2} \rceil$ groups and setting up valuations such that only one group can receive its

MMS value. The details and the necessary proofs are provided in the full version [Hosseini and Searns, 2021].

Theorem 1 illustrates that if the goal is to reach an optimal-MMS threshold, the fraction of agents, α , who receive their MMS guarantee, $\beta = 1$, goes to zero as the number of agents increases. Thus, we ask for what values of α , $(\alpha, 1)$ -MMS is guaranteed to exist?

4 Computing (α, β) -MMS for $n - 1$ Agents

4.1 A $(\frac{2}{3}, 1)$ -MMS Algorithm for Three Agents

It is worth noting that although MMS always exists for $n = 2$ and can be achieved through the cut-and-choose protocol [Bouveret and Lemaître, 2016], computing such an allocation remains hard [Amantidis *et al.*, 2017]. Nonetheless, for two agents an allocation that guarantees MMS^3 to each agent can be computed in polynomial time through an EF1 allocation [Aigner-Horev and Segal-Halevi, 2019]. We use this result to obtain the following result.

Proposition 2. *For three agents, $(\frac{2}{3}, 1)$ -MMS always exists and can be computed in polynomial time.*

Proof. By the definition of (α, β) -MMS, we can select any arbitrary subset of agents $N' \subset N$ such that $|N'| = 2$. Then, run an EF1 algorithm on N' , which outputs an allocation A . Both agents in N' are guaranteed to receive their MMS^3 , i.e., for each $i \in N'$ we have $v_i(A_i) \geq \text{MMS}_i^3$, implying that $\frac{2}{3}$ of agents receive their MMS. By the construction proposed by Kurokawa *et al.* [2018], an MMS allocation does not always exist for three agents; thus, $(\frac{2}{3}, 1)$ -MMS is a tight bound. \square

4.2 A General Algorithm for $n \geq 4$

To extend the analysis of approximate MMS for $n - 1$ agents, we first provide an important lemma that enables us to focus on normalized instances in the remainder of this paper. This lemma states that we can employ valid reductions repeatedly (see Section 2) for computing (α, β) -MMS allocations.

Lemma 4. *Given an instance $I = \langle N, M, V \rangle$, we can compute a normalized instance $I' = \langle N', M', V' \rangle$ in polynomial time such that any (α, β) -MMS allocation on I' implies (α, β) -MMS on I .*

We now focus attention on designing a polynomial-time algorithm that guarantees β MMS for $n - 1$ agents. The algorithm relies on removing an arbitrary agent and computing an EFX allocation for the remaining $n - 1$ agents. By Lemma 1 an ordered instance can be generated (and easily converted back) from an unordered instance. A simple variant to the envy-graph algorithm satisfies EFX on ordered instances [Barman and Krishna Murthy, 2017]. We show that applying this procedure to normalized (and thus ordered) instances satisfies $(\frac{n-1}{n}, \frac{1}{2}(\frac{n+2}{n-1}))$ -MMS. Since β depends on n , we cannot trivially extend this result to any (not normalized) instances. Nonetheless, we prove that together with Lemma 4 and Lemma 1 we can compute $(\frac{n-1}{n}, \frac{1}{2}(\frac{n+2}{n-1}))$ -MMS for any instance in polynomial-time. The full proof of the theorem, along with necessary discussions, can be found in the full version of the paper [Hosseini and Searns, 2021].

n	(α, β) -MMS
4	$(\frac{3}{4}, 1)$ -MMS
5	$(\frac{4}{5}, \frac{7}{8})$ -MMS
6	$(\frac{5}{6}, \frac{4}{5})$ -MMS
7	$(\frac{6}{7}, \frac{3}{4})$ -MMS

Table 1: Approximation bounds of (α, β) for various $n < 8$.

Theorem 2. *Given any instance of n agents, $(\frac{n-1}{n}, \frac{1}{2}(\frac{n+2}{n-1}))$ -MMS can be computed in polynomial time.*

By Proposition 1, we can add a dummy agent when $n = 3$ and select the original set of agents to obtain the following.

Corollary 1. *For $n = 3$ agents, computing an allocation satisfying $\text{MMS}_i^4, \forall i \in N$ can be done in polynomial time.¹*

Remark 2. Theorem 2 immediately illustrates an intriguing interpolation between the approximation ratios of α and β : for $n < 7$, a better approximation of MMS values (β) can be achieved in polynomial time by sacrificing only one agent. Recall that the best approximation algorithms to date guarantee only $(1, \frac{3}{4})$ -MMS [Ghods *et al.*, 2018] for general additive valuations. Table 1 shows the interpolation between α and β for $n = 4$ to $n = 7$.

5 The Existence of $(\frac{2}{3}, 1)$ -MMS Allocations

Theorem 2 optimizes the fraction of agents, α , and provides an efficient approach in achieving approximate MMS for $n - 1$ agents. Another plausible, and often practical, approach aims at maximizing the fraction of agents α who receive their MMS guarantee ($\beta = 1$).

It turns out that simple modifications to existing approximation algorithms can guarantee $(\frac{1}{2}, 1)$ -MMS allocations. A key question is whether we can improve this bound for α beyond $\frac{1}{2}$ and show the existence of such allocations. In what follows, we show that $(\frac{2}{3}, 1)$ -MMS exists and discuss an algorithm achieving this bound in polynomial-time when $n < 9$.

Our existence proof of $(\frac{2}{3}, 1)$ -MMS relies on combining techniques of bag-filling, strong normalization, and a variant of the lone divider procedure. Before discussing our main result, we briefly describe these techniques.

Bag-filling. Given a normalized instance, a good is *high-value* for agent i if $v_i(g) \geq \frac{1}{2}$; otherwise it is *low-value*. The bag-filling algorithm is a greedy approach for forming bundles in a normalized instance. An agent initializes a bag with a high-value good, and then adds low-value goods until the bag is worth at least 1. She then picks another high-value good and repeats this process.² For each bag, the total value never exceeds $1 + \frac{1}{2} = \frac{3}{2}$ since the last added good is low-value.

Remark 3. Bag-filling alone cannot guarantee $(\frac{2}{3}, 1)$ -MMS as it may ‘run out’ of low-value goods before filling $\lfloor \frac{2n}{3} \rfloor$

¹Independently, Aigner-Horev and Segal-Halevi [2019] utilized envy-free matchings to compute MMS^{n+1} for 3 agents.

²Similar approaches have been used by Garg *et al.* [2018]; Garg and Taki [2020]; Ghods *et al.* [2018] to compute $(1, \frac{2}{3})$ -MMS and $(1, \frac{3}{4})$ -MMS allocations. In their algorithms, the bag is filled until any agent values it at least 1.




Figure 1: A sample MMS partition for the divider. The tiny shapes (triangle, circles, squares, and diamonds) represent high-value goods and are ordered from g_1 to g_{10} . Note that each bundle contains no more than one high-value good. All pieces indicated in red (including the red diamonds) indicate the goods that were allocated in previous iterations. There are two such high-value goods, so $n' = \lfloor \frac{2n}{3} \rfloor - 2 = 6$. The first step is pairing the lowest $2s$ remaining high-value goods into s bundles: $\{g_6, g_{10}\}$ and $\{g_8, g_9\}$. The second phase is restricted bag-filling that only uses low-value goods from the remainder sets corresponding to already allocated high-value goods (within teal boxes). The solid borders show *all* remainder sets. Next, the remaining goods, shown in purple are used for simple bag-filling to fill the remaining 2 bundles.

bundles if the remaining value consists entirely of high-value goods. Consider $n = 9$ agents with identical valuations as follows: 5 goods of value 0.99, 5 goods of value 0.01, 1 good of value 0.95, 1 good of value 0.05, 3 goods of value 0.55, and 3 goods of value 0.45. Here, $\text{MMS}_i^n = 1$ for all agents. The high-value goods are valued more than 0.5. During bag-filling, there will be three bundles of $\{0.99, 0.45\}$, one bundle of $\{0.99, 0.05\}$, and one bundle of $\{0.99, 0.01\}$. Since only 5 bundles were filled and $\lfloor \frac{2n}{3} \rfloor = 6$, we need to fill one more bundle. Adding all remaining low-value goods to the next high-value good 0.95 yields a total value 0.99. The remaining value is tied up with the high-value goods (0.55).

Strong Normalization. We say an instance is strongly normalized if it is ordered, $v_i(M) = n$, and $\text{MMS}_i^n = 1$. Observe that the definition of strong normalization implies that each bundle of an MMS partition is valued *exactly* 1. Furthermore, this implies that $v_i(g_1) \leq 1$ and that $v_i(g_{n+1}) \leq \frac{1}{2}$.

Lemma 5. *For any additive instance $I = \langle N, M, V \rangle$, there exists another strongly normalized instance $I' = \langle N, M, V' \rangle$ such that I' is ordered and for all $i \in N$, $\text{MMS}_i^n = 1$ and $v_i(M) = n$. Furthermore, an (α, β) -MMS allocation on I' is also an (α, β) -MMS allocation on I .*

Lemma 5 shows that any instance I can be modified via a non-polynomial-time transformation into a new instance I' that is strongly normalized such that an (α, β) -MMS allocation on I' implies an (α, β) -MMS allocation on I . Hence, we can focus only on strongly normalized instances.

Algorithm description. At its core, our algorithm implements the *lone divider procedure* for indivisible goods on $n' = \lfloor \frac{2n}{3} \rfloor$ of the agents. The lone divider procedure is an extension of the proportional cake-cutting algorithm [Robertson and Webb, 1998] that leverages non-empty envy-free matchings in a bipartite graph. It proceeds as follows: given n' agents in N' , first, an arbitrary agent divides the goods into n' acceptable bundles $(A_1, \dots, A_{n'})$. Second, we construct a bipartite graph between agents in N' and the bundles $(A_1, \dots, A_{n'})$ wherein each edge connects an agent $i \in N'$ to a bundle A_k such that $v_i(A_k) \geq 1$. Notice that the divider will be adjacent to all bundles. Next, we compute a non-empty envy-free matching, where no unmatched agent is adjacent to

a bundle that was assigned to some other agent. A non-empty envy-free matching always exists if $|\mathcal{N}(N')| \geq |N'|$, where $\mathcal{N}(N')$ denotes the set of bundles adjacent to N' [Aigner-Horev and Segal-Halevi, 2019]. All matched agents receive their bundles and leave. After the removal of matched agents and goods, we repeat the same procedure over the remaining goods and agents until there are no remaining agents.³

The primary challenge in proving $(\frac{2}{3}, 1)$ -MMS existence lies in showing that in each iteration, given n' remaining agents, the divider agent can form n' bundles valued at least 1. We show that under the strong normalization assumption, the divider is always able to form n' bundles either by pairing high-value goods or through *restricted bag-filling*.

Theorem 3. *A $(\frac{2}{3}, 1)$ -MMS allocation is guaranteed to exist for any subset of $\lfloor \frac{2n}{3} \rfloor$ agents.*

Our proof relies on two technical cases based on the number of high-value goods (h). If $h \leq \lfloor \frac{2n}{3} \rfloor$, there are not many high-value goods, and thus, the simple bag-filling guarantees our desired result. The most challenging cases arise when $h > \lfloor \frac{2n}{3} \rfloor$, that is, when there are too many high-value goods. Let $s = h - \lfloor \frac{2n}{3} \rfloor$ be the number of excess high-value goods as shown in Figure 1. Here, we show that the divider agent is able to form the necessary bundles (*i.e.* the number of remaining agents) by adopting the following strategies: if $n' \leq s$, we can simply pair the remaining high-value goods since there are at least $2n'$ high-value goods available, which implies n' bundles are valued more than 1. On the other hand, if $n' > s$, we are able to form s bundles by pairing the remaining high-value goods. This means that we still need an additional $n' - s$ bundles. The key to handling this step relies on the *restricted bag-filling* procedure with exactly one high-value good per bundle but we restrict which low-value goods are used to fill bundles. In restricted bag-filling, we only use the low-value goods from bundles of the divider's MMS^n partition whose high-value goods are no longer available (since those goods have been allocated in some previous round or assigned during pairing), as highlighted in Figure 1 by teal boxes.

The main reason behind this restriction is to ensure that not too much value is ‘wasted’ during bag-filling (as discussed in

³See [Hosseini and Searns, 2021] for the technical details.

Remark 3). The restricted bag-filling relies on the strong normalization assumption - this way we can guarantee that the low-value goods used in restricted bag-filling have sufficient value when bundled with at most one high-value good.

Our proof follows by showing that if $n' \geq 2s$, it is possible to form $n' - s$ bundles through restricted bag-filling. Otherwise if $n' < 2s$, using restricted bag-filling we can form s bundles. In this case, we still need to form an additional $n' - 2s$ bundles, which is possible through simple bag filling by targeting for each bundle to have at most $\frac{3}{2}$ value.

Remark 4. The strong normalization assumption in the proof of Theorem 3 requires modifying the values of some goods depending on an MMS partition for agent i . Furthermore, the restricted bag-filling phase of Theorem 3 selects available low-value goods from specific bundles of an MMS partition for agent i . Because of these dependencies on finding MMS partitions, Theorem 3 does not imply a polynomial time algorithm for $(\frac{2}{3}, 1)$ -MMS.

We show that by relaxing the strong normalization assumption to the weaker normalization of Lemma 4, we can devise a polynomial time algorithm for $(\frac{2}{3}, 1)$ -MMS when $n < 9$. Essentially, this algorithm is a simplified variant of our previous procedure that only includes simple bag-filling and pairing to form the required bundles in the lone divider procedure. Intuitively, since the total value of low-value goods is guaranteed to be at least $1 - v_i(g_1)$, at least one bundle can be allocated during bag-filling. When $n < 9$, adding the pairing phase forms $\lfloor \frac{n}{2} \rfloor + 1 > \lfloor \frac{2n}{3} \rfloor$ bundles.

Theorem 4. For $n < 9$, $(\frac{2}{3}, 1)$ -MMS can be computed in polynomial time.

Theorem 4 only guarantees a $(\frac{2}{3}, 1)$ -MMS allocation when $n < 9$. We show by construction that this is tight (see [Hosseini and Searns, 2021] for details). We construct a family of instances with $n \geq 9$ in which a small error in computing the MMS bound causes bag-filling to stop before $\lfloor \frac{2n}{3} \rfloor$ bundles have been allocated. The challenge presented here is not unique to our algorithm. Any algorithm that satisfies $(\frac{2}{3}, 1)$ -MMS must be able to detect that $\text{MMS}_i^n < 1$ for all agents. However, this task is intractable even when agents have identical valuations [Bouveret and Lemaître, 2016].

Combining this construction with Theorem 4 implies that no smaller counter-example exists. We note that despite this bound, the algorithm is still practical as there is no bound on the number of goods. For example, more than 99% of instances in the *Spliddit* dataset deal with less than 9 agents.

The analysis in Theorem 3 and Theorem 4 are indifferent to the set of agents N' . Consequently, we may select any subset of $\lfloor \frac{2n}{3} \rfloor$ agents to be given their full MMS. An immediate consequence of this result, together with Proposition 1, is that there always exists an allocation of goods such that each agent receives at least the value of $\text{MMS}^{\lceil \frac{3n}{2} \rceil}$. Moreover, this allocation can be computed in polynomial time when $\lceil \frac{3n}{2} \rceil < 9$.

Corollary 2. An allocation satisfying $\text{MMS}^{\lceil \frac{3n}{2} \rceil}$ always exists. Moreover, such allocations can be computed in polynomial time when $n < 6$.




Figure 2: Fraction of agents, α , receiving their maximin share, i.e., $(\alpha, 1)$ -MMS.

6 Empirical Evaluations

The polynomial algorithm of Theorem 4 guarantees $(\frac{2}{3}, 1)$ -MMS only for $n < 9$. We evaluate a variant of this algorithm that does not rely on the strong normalization assumption (details relegated to the full version). It is similar to the algorithm of Theorem 4 until either no more bundles can be allocated during the lone divider procedure or all $\lfloor \frac{2n}{3} \rfloor$ agents that were initially selected have received a bundle valued at least 1. In the latter case, any remaining goods are distributed among the unselected $n - \lfloor \frac{2n}{3} \rfloor$ agents using bag-filling with priority given to agents who accept the smallest bundles. Notice that this algorithm does not guarantee that all initially selected agents receive their MMS.

We focus on ordered instances—as the most difficult instances in achieving MMS [Bouveret and Lemaître, 2016]—and generate 1,000 instances for each combination of n and m . Instances are sampled uniformly at random, ordered, and scaled such that $v_i(M) = n$ for all agents $i \in N$. Figure 2 illustrates the fraction of agents who receive their MMS for $n = 3$ to 50 agents and $m = 3$ to 200 goods. In almost all instances, the algorithm goes beyond the $\frac{2}{3}$ bound: on average across all instances, more than 90% of agents receive their MMS. Moreover, the fraction of agents receiving their MMS improves as either n or m increases. We also observe linear bands where a lower fraction of agents are satisfied due to the ratio of goods to agents. In addition, when $m < 2n$ a large fraction of the agents receive their MMS and are removed during the reduction phase of normalization.

7 Discussion

Theorem 3 proves the existence of $(\frac{2}{3}, 1)$ -MMS for any n , which implies the existence of $\text{MMS}^{\lceil \frac{3n}{2} \rceil}$. Therefore, improving the bound on $(\alpha, 1)$ -MMS for $\alpha > \frac{2}{3}$ and closing the gap between $\text{MMS}^{\lceil \frac{3n}{2} \rceil}$ and MMS^{n+1} is an intriguing future direction. Theorem 4 provides a tractable approach for computing $(\frac{2}{3}, 1)$ -MMS for $n < 9$. Yet, computing such allocations for any n , if possible, will require further techniques to circumvent computing exact MMS bounds. Another interesting avenue for future research is exploring PTAS algorithms that guarantee optimal-MMS [Heinen *et al.*, 2018] for a fraction of agents to achieve $(\alpha, 1)$ -MMS for $\alpha \in [\frac{2}{3}, \frac{n-1}{n}]$.

Acknowledgements

HH acknowledges the support from NSF grant #1850076. We are grateful to Erel Segal-Halevi for his valuable input that improved the proof of Theorem 3.

References

- [Aigner-Horev and Segal-Halevi, 2019] Elad Aigner-Horev and Erel Segal-Halevi. Envy-free matchings in bipartite graphs and their applications to fair division. *arXiv preprint arXiv:1901.09527*, 2019.
- [Amanatidis et al., 2017] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms (TALG)*, 13(4):52, 2017.
- [Barman and Krishna Murthy, 2017] Siddharth Barman and Sanath Kumar Krishna Murthy. Approximation algorithms for maximin fair division. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 647–664. ACM, 2017.
- [Bouveret and Lemaître, 2016] Sylvain Bouveret and Michel Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290, Mar 2016.
- [Brams and Taylor, 1996] Steven J Brams and Alan D Taylor. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press, 1996.
- [Budish, 2011] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis et al., 2016] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 305–322. ACM, 2016.
- [Chen and Shah, 2018] Yiling Chen and Nisarg Shah. Ignorance is often bliss: Envy with incomplete information. Technical report, 2018.
- [Chevalerey et al., 2007] Yann Chevaleyre, Ulle Endriss, Sylvia Estivie, and Nicolas Maudet. Reaching envy-free states in distributed negotiation settings. In *Proceedings of the 20th international joint conference on Artificial intelligence*, pages 1239–1244. Morgan Kaufmann Publishers Inc., 2007.
- [Feldman and Kirman, 1974] Allan Feldman and Alan Kirman. Fairness and envy. *The American Economic Review*, 64(6):995–1005, 1974.
- [Foley, 1967] Duncan K. Foley. Resource allocation and the public sector. *Yale Economic Essays*, 7:45–98, 1967.
- [Freeman et al., 2019] Rupert Freeman, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. Equitable allocations of indivisible goods. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 280–286, 7 2019.
- [Garg and Taki, 2020] Jugal Garg and Setareh Taki. An improved approximation algorithm for maximin shares. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC ’20, page 379–380, New York, NY, USA, 2020. Association for Computing Machinery.
- [Garg et al., 2018] Jugal Garg, Peter McGlaughlin, and Setareh Taki. Approximating maximin share allocations. In *2nd Symposium on Simplicity in Algorithms (SOSA 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [Ghods et al., 2018] Mohammad Ghods, MohammadTaghi Haji-Aghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 539–556. ACM, 2018.
- [Gourvès and Monnot, 2019] Laurent Gourvès and Jérôme Monnot. On maximin share allocations in matroids. *Theoretical Computer Science*, 754:50–64, 2019.
- [Heinen et al., 2018] Tobias Heinen, Nhan-Tam Nguyen, Trung Thanh Nguyen, and Jörg Rothe. Approximation and complexity of the optimization and existence problems for maximin share, proportional share, and minimax share allocation of indivisible goods. *Autonomous Agents and Multi-Agent Systems*, 32(6):741–778, 2018.
- [Hosseini and Searns, 2021] Hadi Hosseini and Andrew Searns. Guaranteeing maximin shares: Some agents left behind. *arXiv preprint arXiv:2105.09383*, 2021.
- [Hosseini et al., 2020] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, Jun Wang, and Lirong Xia. Fair Division through Information Withholding. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*, pages 2014–2021, 2020.
- [Kurokawa et al., 2016] David Kurokawa, Ariel D Procaccia, and Junxing Wang. When can the maximin share guarantee be guaranteed? In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 523–529. AAAI Press, 2016.
- [Kurokawa et al., 2018] David Kurokawa, Ariel D Procaccia, and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM (JACM)*, 65(2):8, 2018.
- [Lipton et al., 2004] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 125–131. ACM, 2004.
- [Nguyen and Rothe, 2014] Trung Thanh Nguyen and Jörg Rothe. Minimizing envy and maximizing average nash social welfare in the allocation of indivisible goods. *Discrete Applied Mathematics*, 179:54–68, 2014.
- [Nguyen et al., 2017] Nhan-Tam Nguyen, Trung Thanh Nguyen, and Jörg Rothe. Approximate solutions to max-min fair and proportionally fair allocations of indivisible goods. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 262–271, 2017.
- [Nyman et al., 2020] Kathryn Nyman, Francis Edward Su, and Shira Zerbib. Fair division with multiple pieces. *Discrete Applied Mathematics*, 283:115 – 122, 2020.
- [Ortega, 2018] Josué Ortega. Social integration in two-sided matching markets. *Journal of Mathematical Economics*, 78:119–126, 2018.
- [Procaccia and Wang, 2014] Ariel D Procaccia and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 675–692. ACM, 2014.
- [Robertson and Webb, 1998] Jack Robertson and William Webb. *Cake-Cutting Algorithms: Be fair if You Can*. AK Peters/CRC Press, 1998.
- [Schneckenburger et al., 2017] Sebastian Schneckenburger, Britta Dorn, and Ulle Endriss. The atkinson inequality index in multi-agent resource allocation. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 272–280, 2017.
- [Woeginger, 1997] Gerhard J Woeginger. A polynomial-time approximation scheme for maximizing the minimum machine completion time. *Operations Research Letters*, 20(4):149–154, 1997.