Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI-23)
Journal Track

Ordinal Maximin Share Approximation for Goods (Extended Abstract)*

Hadi Hosseini[1] , **Andrew Searns**[2] and **Erel Segal-Halevi**[3]
[1]Pennsylvania State University
[2]Johns Hopkins University Applied Physics Laboratory
[3]Ariel University
hadi@psu.edu, andrew.searns@jhuapl.edu, 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 *in*divisible. 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. 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 and Wang, 2014; Kurokawa *et al.*, 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. The best currently known fraction is $\frac{3}{4} + \frac{1}{12n}$ [Garg and 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 sensitive to agents' precise cardinal valuations. 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 $\frac{3}{4}$ 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).

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 that guarantee 1-out-of-d MMS for some integer $d > n$.

The aforementioned guarantee can be naturally generalized to ℓ-out-of-d MMS [Babaioff *et al.*, 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?*

*Full paper appeared in the Journal of Artificial Intelligence Research [Hosseini *et al.*, 2022b].

1.1 Our Contributions

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

In **Section 3**, 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 and 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 4**, we present polynomial-time algorithms that find an ℓ-out-of-$(d + o(n))$ MMS-fair allocation, where $d = (\ell + \frac{1}{2})n$. 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-$\lfloor 3n/2\rfloor$ MMS up to an additive gap of at most 1. For $\ell > 1$, we present a different polynomial-time algorithm for finding an ℓ-out-of-$\lceil(\ell + \frac{1}{2})n + O(n^{2/3})\rceil$ MMS allocation (Theorem 3).

In the full version of the paper [Hosseini *et al.*, 2022b], 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.[1] 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 3 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-$\lfloor 3n/2\rfloor$ 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 the full version ([Hosseini *et al.*, 2022b]), 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 and 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 3).

Bin Covering. To develop a polynomial-time algorithm when $\ell = 1$, we extend an algorithm of Csirik *et al.* [1999] for the *bin covering* problem—a dual of 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.

1.3 Related Work

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 and 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 *et al.*; Babaioff *et al.* [2019; 2021], and further studied by Segal-Halevi [2020].

The multiplicative approximation to MMS originated in the computer science literature [Procaccia and Wang, 2014]. The non-existence of MMS allocations [Kurokawa *et al.*, 2018] and its intractability [Bouveret and Lemaître, 2016; Woeginger, 1997] have given rise to a number of approximation techniques. The currently known algorithms guarantee $\beta \geq 2/3$ [Kurokawa *et al.*, 2018; Amanatidis *et al.*, 2017; Garg *et al.*, 2018] and $\beta \geq 3/4$ [Ghodsi *et al.*, 2018; Garg and Taki, 2020] in general, and $\beta \geq 7/8$ [Amanatidis *et al.*, 2017] as well as $\beta \geq 8/9$ [Gourvès and Monnot, 2019] when there are only three agents. Amanatidis *et al.*; Halpern and Shah [2016; 2021] show that the highest attainable multiplicative approximation of MMS is $\Theta(1/\log n)$ when agents report only a ranking over the *goods*.

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 receives at least a fraction β of its MMS. However, Searns and Hosseini [2020] 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.

[1] 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 [Hosseini *et al.*, 2022b]).

2 Preliminaries

2.1 Agents and Goods

Let $N = [n] := \{1, \ldots, n\}$ 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 $v_i(\emptyset) = 0$.[2]

An *instance* of the problem is denoted by $I = \langle N, M, V \rangle$, where $V = (v_1, \ldots, v_n)$ 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* $A = (A_1, \ldots, A_n)$ is an n-partition of M that allocates the bundle of goods in A_i to each agent $i \in N$.

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.

2.2 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

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

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.

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.

2.3 The Lone Divider Algorithm

The Lone Divider algorithm [Kuhn, 1967], as described in Aigner-Horev and Segal-Halevi [2022], accepts as input a set M of items and a threshold value t_i for each agent i. The threshold value t_i should satisfy the condition that, if some k bundles with values below the threshold are removed from M, then there exists a partition of the remaining items into $n - k$ bundles with value at least t_i. This guarantees that, in any iteration, any agent either receives a bundle with value at least t_i, or become a divider in a later iteration and guarantee a bundle with value at least t_i.

[2] In the full version of the paper [Hosseini *et al.*, 2022b] we complement our results with a non-existence result for the more general class of *responsive* preferences.

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

Require: An instance $\langle N, M, V \rangle$ and an integer $\ell \geq 1$.
Ensure: An ℓ-out-of-$\lfloor(\ell + \frac{1}{2})n\rfloor$ MMS allocation.
1: Scale the valuations of all agents such that the ℓ-out-of-$\lfloor(\ell + \frac{1}{2})n\rfloor$ MMS equals ℓ (see full version for details).
2: Order the instance as in [Bouveret *et al.*, 2016; Barman and Krishna Murthy, 2017].
3: Run the Lone Divider algorithm 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.

3 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 1) for achieving the above MMS bound. Algorithm 1 starts with two normalization steps, some of which appeared in previous works and some are specific to our algorithm. The algorithm applies to the normalized instance an adaptation of the Lone Divider algorithm, in which the divider in each step must construct an *ℓ-balanced* partition.

An ℓ-balanced partition is a partition in which each part contains exactly one good from among the n most-valuable goods; exactly one good from among the n second-most-valuable goods; and so on up to exactly one good from the n ℓth-most-valuable goods (yielding exactly ℓ goods from the ℓn most valuable goods). We prove that, when the partition in each iteration is an ℓ-balanced partition, the threshold of ℓ satisfies the condition required for the Lone Divider algorithm to succeed.

4 Ordinal Approximation for Goods in Polynomial Time

Algorithm 1 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 1 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 *et al.* [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

Algorithm 2 Bidirectional bag-filling

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

```
Order the instance in descending order of value, so that
  for each agent i, v_i(g_1) ≥ ··· ≥ v_i(g_m).
for k = 1, 2, ...: do
    Initialize a bag with the good g_k.
    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.
    Give the goods in the bag to an arbitrary agent i who values it at least t_i.
    If every remaining agent i values the remaining goods at less than t_i, stop.
end for
```

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 $b \geq 0$ 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.

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 2 generalizes the aforementioned algorithm to MMS allocation of goods. Thus, the algorithm of Csirik *et al.* [1999] corresponds to a special case of Algorithm 2 wherein (i) All agents have the same v_i (describing the item sizes); and (ii) All agents have the same t_i (describing the bin size).

To compute a threshold for agent i, we simulate Algorithm 2 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 1. 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 2 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.

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.

We show that a BBFS-fair allocation always exists, and can be found in time polynomial in the length of the binary representation of the problem. Moreover, a BBFS-fair allocation is also 1-out-of-$\lceil 3n/2 \rceil$ MMS-fair, although the BBFS may be larger than 1-out-of-$\lceil 3n/2 \rceil$ MMS.

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.*

Example 1 (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 2 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 2 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 1. 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 experiments on instances generated uniformly at random, Algorithm 2 significantly outperforms the theoretical guarantee of 1-out-of-$\lceil 3n/2 \rceil$ MMS ([Hosseini *et al.*, 2022b]).

The case when $\ell > 1$. So far, we could not adapt Algorithm 2 to finding an ℓ-out-of-$\lfloor(\ell + 1/2)n\rfloor$ MMS allocation for $\ell \geq 2$. Using a weaker approximation to MMS, along with other approximations of bin-covering, we prove 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 $\mathrm{MMS}_i^{\ell\text{-out-of-}\lceil d+15d^{2/3}+\ell\rceil}(M)$, in time $\widetilde{O}\left(n \cdot m^4\right)$.*

5 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.

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 *et al.*, 2022a].

Acknowledgments

Hadi Hosseini acknowledges support from NSF IIS grants #2144413, #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.

References

[Aigner-Horev and Segal-Halevi, 2022] Elad Aigner-Horev and Erel Segal-Halevi. Envy-free matchings in bipartite graphs and their applications to fair division. *Information Sciences*, 587:164–187, 2022.

[Amanatidis *et al.*, 2016] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. On truthful mechanisms for maximin share allocations. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 31–37, 2016.

[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.

[Assmann *et al.*, 1984] Susan F Assmann, David S. Johnson, Daniel J. Kleitman, and JY-T Leung. On a dual version of the one-dimensional bin packing problem. *Journal of algorithms*, 5(4):502–525, 1984.

[Babaioff *et al.*, 2019] Moshe Babaioff, Noam Nisan, and Inbal Talgam-Cohen. Fair allocation through competitive equilibrium from generic incomes. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 180–180, 2019.

[Babaioff *et al.*, 2021] Moshe Babaioff, Noam Nisan, and Inbal Talgam-Cohen. Competitive equilibrium with indivisible goods and generic budgets. *Mathematics of Operations Research*, 46(1):382–403, 2021.

[Barman and Krishna Murthy, 2017] Siddharth Barman and Sanath Kumar Krishna Murthy. Approximation algorithms for maximin fair division. In *Proceedings of the ACM Conference on Economics and Computation*, pages 647–664, 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.

[Bouveret *et al.*, 2016] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible items. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jerome Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13, pages 284–310. Cambridge University Press, 2016.

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

[Csirik *et al.*, 1999] János Csirik, J. B. G. Frenk, M. Labbè, and S. Zhang. Two simple algorithms for bin covering. *Acta Cybernetica*, 14(1):13–25, 1999.

[Garg and Taki, 2020] Jugal Garg and Setareh Taki. An improved approximation algorithm for maximin shares. In *Proceedings of the ACM Conference on Economics and Computation*, pages 379–380, 2020.

[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.

[Ghodsi *et al.*, 2018] Mohammad Ghodsi, MohammadTaghi HajiAghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 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.

[Halpern and Shah, 2021] Daniel Halpern and Nisarg Shah. Fair and efficient resource allocation with partial information. In *Proceedings of the International Joint Conference on Artificial Intelligence IJCAI*, pages 224–230, 2021.

[Hosseini and Searns, 2021] Hadi Hosseini and Andrew Searns. Guaranteeing maximin shares: Some agents left behind. In Zhi-Hua Zhou, editor, *Proceedings of the International Joint Conference on Artificial Intelligence IJCAI*, pages 238–244. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Main Track.

[Hosseini *et al.*, 2022a] Hadi Hosseini, Andrew Searns, and Erel Segal-Halevi. Ordinal maximin share approximation for chores. In *Proceedings of the International Conference on Autonomous Agents and MultiAgent Systems*, page forthcoming, 2022.

[Hosseini *et al.*, 2022b] Hadi Hosseini, Andrew Searns, and Erel Segal-Halevi. Ordinal maximin share approximation for goods. *Journal of Artificial Intelligence Research*, 74:353–391, 2022.

[Johnson, 1973] David S Johnson. *Near-optimal bin packing algorithms*. PhD thesis, Massachusetts Institute of Technology, 1973.

[Kuhn, 1967] Harold W. Kuhn. *On games of fair division*, pages 29–37. Princeton University Press, 1967.

[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.

[Nguyen *et al.*, 2017] Nhan-Tam Nguyen, Trung Thanh Nguyen, and Jörg Rothe. Approximate solutions to maxmin fair and proportionally fair allocations of indivisible goods. In *Proceedings of the Conference on Autonomous*

Agents and MultiAgent Systems, pages 262–271. International Foundation for Autonomous Agents and Multiagent Systems, 2017.

[Procaccia and Wang, 2014] Ariel D Procaccia and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the ACM conference on Economics and computation*, pages 675–692. ACM, 2014.

[Searns and Hosseini, 2020] Andrew Searns and Hadi Hosseini. Fairness does not imply satisfaction (student abstract). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(10):13911–13912, Apr. 2020.

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

[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.