



# Maximin Fair Allocation of Indivisible Items Under Cost Utilities

Sirin Botan, Angus Ritossa, Mashbat Suzuki<sup>(✉)</sup>, and Toby Walsh

UNSW Sydney, Kensington, Australia  
{s.botan,mashbat.suzuki,t.walsh}@unsw.edu.au,  
a.ritossa@student.unsw.edu.au

**Abstract.** We study the problem of fairly allocating indivisible goods among a set of agents. Our focus is on the existence of allocations that give each agent their *maximin fair share*—the value they are guaranteed if they divide the goods into as many bundles as there are agents, and receive their lowest valued bundle. An *MMS allocation* is one where every agent receives at least their maximin fair share. We examine the existence of such allocations when agents have *cost utilities*. In this setting, each item has an associated *cost*, and an agent’s valuation for an item is the cost of the item if it is useful to them, and zero otherwise.

Our main results indicate that cost utilities are a promising restriction for achieving MMS. We show that for the case of three agents with cost utilities, an MMS allocation always exists. We also show that when preferences are restricted slightly further—to what we call *laminar set approvals*—we can guarantee MMS allocations for any number of agents. Finally, we explore if it is possible to guarantee each agent their maximin fair share while using a strategyproof mechanism.

**Keywords:** Fair Division · Maximin Fair Share · Resource Allocation

## 1 Introduction

How to fairly divide a set of indivisible resources is a problem that has been studied by computer scientists, economists, and mathematicians [10, 11, 25]. Because of the fundamental nature of the problem, there is a large number of applications ranging from course allocations [26], to division of assets [19], and air traffic management [27].

Among the fairness notions studied, two of the most commonly studied are those of envy-freeness—how to ensure no agent envies another, and maximin fair share—our focus in this paper. The notion of the maximin fair share was introduced by Budish [12], and generalises the well known cut-and-choose protocol. Conceptually, an agent’s *maximin fair share* is the value they can achieve by partitioning the items into as many bundles as there are agents, and receiving their least preferred bundle. The ideal outcome is of course an *MMS allocation*, where every agent receives at least their maximin fair share.

There has been a significant amount of work on MMS in the general additive valuations setting. Unfortunately, results are often quite negative. In general, MMS allocations cannot be guaranteed to exist, even in the case of three agents [16, 24]. Furthermore, for instances where MMS allocations do exist (for example, when agents have identical valuations), computing an MMS allocation is NP-hard. As a result, a large body of work has been focused on establishing the existence of MMS allocations in more restricted settings [4, 8, 9, 15]. In this paper, we study the problem under a natural class of valuation functions—what we call *cost utilities*—that allow us to provide fairness guarantees that are not achievable for general additive valuations. Cost utilities describe the setting where each item has an associated cost. An agent’s value for any item is the cost of the item if it is useful to them, and zero otherwise. Our focus in this work is on the existence of MMS allocations under cost utilities.

We are not the first to study this restriction in the context of fair division. Bansal and Sviridenko [7] provided an approximation of egalitarian welfare maximisation under cost utilities, that was then improved upon by Asadpour et al. [5] and Cheng and Mao [14]. Camacho et al. [13] and Akrami et al. [1] focus on envy-freeness, and show that an EFX allocation always exists under cost utilities<sup>1</sup>. There are clear practical advantages to studying this particular class of valuations. In many real-life settings, the price of items are known, and elicitation of preferences boils down to asking an agent whether they want the item or not—a task that can be accomplished easily.

*Related Work.* Given that MMS allocations cannot be guaranteed for general additive valuations, the work done on MMS in fair division has focused on two main approaches to circumvent this impossibility. The first—which is the route we employ in this paper—is to consider a restriction on the valuations of the agents. Examples of such restrictions under which MMS allocations always exist include binary valuations [9], and ternary valuations [4]—where item values belong to  $\{0, 1, 2\}$ , and Borda utilities [21]. Existence of MMS allocations also holds for personalised bivalued valuations—where for each agent  $i$ , the value of an item belongs to  $\{1, p_i\}$  for  $p_i \in \mathbb{N}$ , and weakly lexicographical valuations—where each agent values each good more than the combined value of all items that are strictly less preferred [15].

The second approach is to examine *how close* we can get to MMS, meaning how far each agent is from receiving their maximin fair share. An allocation is said to be  $\rho$ -MMS, if each agent receives a  $\rho$  fraction of their MMS value. Garg and Taki [17] show that for instances with more than five agents a  $(\frac{3}{4} + \frac{1}{12n})$ -MMS allocation always exists. On the more negative side, [16] show that there exist instances such that no allocation is  $39/40$ -MMS. For valuations that are beyond additive, the picture is arguably gloomier. [18] show an existence of  $\frac{1}{3}$ -MMS allocations and a PTAS for computing such allocations. They also show

<sup>1</sup> Bansal and Sviridenko [7] call them *restricted assignment valuations*, while Camacho et al. [13] call them *generalised binary valuations*. Akrami et al. [1] study them under the name *restricted additive valuations*. We use the term “cost utilities” as we find it conceptually the most appealing and descriptive.

that for submodular valuations, there exist instances that do not admit any  $\frac{3}{4}$ -MMS allocation.

There have been several works focused on achieving both fairness along with strategyproofness. Amanatidis et al. [2] show that when there are two agents and  $m$  items there is no truthful mechanism that outputs an  $\frac{1}{\lfloor m/2 \rfloor}$ -MMS allocation. On the positive side Halpern et al. [20] and Babaioff et al. [6] show that when agents have binary valuations there is a polynomial time computable mechanism that is strategyproof and outputs an MMS allocation along with several other desirable properties.

*Our Contribution.* We know that for some restricted settings—bivalued and ternary valuations—MMS allocations can always be found. Amanatidis et al. [3] highlight an open problem regarding the existence of other classes of structured valuations for which an MMS allocation is guaranteed to exist. Our paper answers this in the affirmative for a new class of valuation functions. We first show that MMS allocations exist for three agents under cost utilities, in contrast to the case of general additive utilities. We also show that when valuations are restricted slightly further to laminar set approvals, MMS allocations are guaranteed to exist for any number of agents. Additionally, for the case of  $n$  agents and  $n + 2$  items, we show there is a strategyproof polynomial time algorithm for computing Pareto optimal MMS allocations.

Interestingly, to the best of our knowledge, our results on cost utilities are first of its kind for which (other than identical valuations) the computation of the maximin fair share value is NP-hard, while existence of MMS allocation is still guaranteed. For previously known classes where an MMS allocation is guaranteed, the computation of the maximin fair share value can be done in polynomial time.

*Paper Outline.* In Sect. 2 we introduce the framework of fair division of indivisible items, and present the central preference and fairness notions of the paper. Section 3 is focused on when we can achieve MMS allocations for cost utilities. Section 4 looks at a strategyproof mechanism for finding MMS allocations. Section 5 concludes.

## 2 Preliminaries

Let  $N$  be a set of  $n$  agents, and  $M$  a set of  $m$  indivisible goods (or items). Our goal is to divide  $M$  among the agents in  $N$  according to their preferences over the items.

*Preferences.* Each agent  $i \in N$  has a valuation function  $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$  that determines how much they value any bundle of items. For all agents  $i$ , we assume that  $v_i$  is additive, so  $v_i(S) = \sum_{g \in S} v_i(g)$ . For singleton bundles, we write  $v_i(g)$  in place of  $v_i(\{g\})$  for simplicity. We write  $\mathbf{v} = (v_1, \dots, v_n)$  to denote the vector of all valuation functions for agents in  $N$ .

Our focus in this paper is on a restricted domain of preferences—cost utilities. For these preferences, it is easy to think of each agent as submitting an approval set. Let  $A_i$  be the *approval set* of agent  $i$ . More formally, we say  $A_i = \{g \in M \mid v_i(g) > 0\}$ . We say agents have *cost utilities* if there exists a *cost function*  $c$  such that  $v_i(S) = c(S \cap A_i)$  for all  $S \subseteq M$  and all agents  $i \in N$ . We require that the cost function is additive, as well as non-negative.

*Allocations and Mechanism.* An *allocation*  $\mathbf{B} = (B_1, \dots, B_n)$  is an  $n$ -partition of the set of items  $M$ , where  $B_i \subseteq M$  is the *bundle* assigned to agent  $i$  under the allocation  $\mathbf{B}$ . We write  $\mathbf{B}|_{N'}$  to denote the restriction of the allocation  $\mathbf{B}$  to only the bundles assigned to agents in  $N' \subseteq N$ . For a set of goods  $M$ , we write  $\mathcal{B}_n(M)$  to mean all possible allocations of the goods in  $M$  to  $n$  agents. An *instance*  $\mathcal{I} = (N, M, \mathbf{v})$  of a fair division problem is defined by a set of agents, a set of goods, and the agents' valuations over those goods.

Given an instance  $\mathcal{I}$ , our goal is to find an allocation  $\mathbf{B}$  that satisfies certain normative properties. An *allocation mechanism* for  $n$  agents and  $m$  items is a function  $f : \mathbf{V}_n \rightarrow \mathcal{B}_n(M)$ , where  $\mathbf{V}_n$  is the set of possible valuation profiles—i.e. vectors of  $n$  valuation functions.

*Fairness and Efficiency.* For an agent  $i \in N$ , their *maximin fair share* in an instance  $\mathcal{I} = (N, M, \mathbf{v})$  is defined as

$$\text{MMS}_i^n(\mathcal{I}) = \max_{\mathbf{B} \in \mathcal{B}_n(M)} \min_{j \in N} v_i(B_j).$$

We sometimes write  $\text{MMS}_i^n(M)$  when the instance is clear from context. When the set of goods and the value of  $n$  is fixed, we will also sometimes write  $\text{MMS}_i$ .

An *MMS allocation*  $\mathbf{B} \in \mathcal{B}_n(M)$  is an allocation such that  $v_i(B_i) \geq \text{MMS}_i$  for all agents  $i \in N$ .

We say an allocation  $\mathbf{B} \in \mathcal{B}_n(M)$  is *Pareto efficient* if there is no allocation  $\mathbf{B}' \in \mathcal{B}_n(M)$  such that  $v_i(B'_i) \geq v_i(B_i)$  for all  $i \in N$  and  $v_{i^*}(B'_{i^*}) > v_{i^*}(B_{i^*})$  for some  $i^* \in N$ .

### 3 Maximin Fair Share Guarantees

In this section, we will look at two settings where cost utilities can aid in finding cases where MMS allocations can be guaranteed to exist. Section 3.1 focuses on cases with only three agents. Section 3.2 considers any number of agents but is limited to laminar approval sets. This is a restriction that captures the idea of items belonging to different categories.

#### 3.1 MMS Allocations for Three Agents

For the case of three agents, restricting our scope to considering only cost utilities yields positive results. As we have seen in the introduction, this is not the case

for the more general case of additive preferences. Theorem 1 is therefore a very welcome result.

In this section, we will sometimes speak about items approved exclusively by two agents. We denote by  $A_{ij} = (A_i \cap A_j) \setminus A_{i^*}$ —where  $i^* \in N$  and  $i^* \neq i, j$ —the set of items approved by agents  $i$  and  $j$ , and no third agent.

Before we state our main result in this section, we present the following two lemmas that we need in order to prove Theorem 1. Our first lemma simply tells us that adding items approved only by a single agent does not affect the existence of an MMS allocation.

**Lemma 1.** *If an MMS allocation exists for instance  $\mathcal{I} = (N, M, \mathbf{v})$ , then an MMS allocation also exists for the instance  $\mathcal{I}' = (N, M \cup S, \mathbf{v})$ , where  $S$  is a set of items approved by a single agent  $i \in N$ , and  $S \cap M = \emptyset$ .*

*Proof.* Suppose we have an instance  $\mathcal{I} = (N, M, \mathbf{v})$  where  $\mathbf{B}$  is an MMS allocation. Suppose further that  $\mathcal{I}' = (N, M \cup S, \mathbf{v})$  is an instance where  $S$  is a set of items approved by a single agent  $i \in N$ , and  $S \cap M = \emptyset$ . We show that  $\mathbf{B}'$  where  $B'_j = B_j$  for all  $j \neq i$  and  $B'_i = B_i \cup S$  is an MMS allocation. Since for any  $j \neq i$  we have  $v_j(B_j) \geq \text{MMS}_j^n$ , we only need to show that agent  $i$  gets her MMS fair share.

Suppose for contradiction that we have  $v_i(S) + \text{MMS}_i^n(M) < \text{MMS}_i^n(M \cup S)$ . Let  $\mathbf{W} = (W_1, \dots, W_n)$  be an  $n$ -partition of  $(M \cup S)$  such that  $v_i(W_k) \geq \text{MMS}_i^n(M \cup S)$  for  $1 \leq k \leq n$ . Note that for any  $W_k$  in the partition we have that  $W_k = (W_k \cap M) \cup (W_k \cap S)$ . Thus we have the following:

$$\begin{aligned} \text{MMS}_i^n(M \cup S) &\leq v_i(W_k) \\ &= v_i(W_k \cap M) + v_i(W_k \cap S) \\ &\leq v_i(W_k \cap M) + v_i(S) \\ &< v_i(W_k \cap M) + \text{MMS}_i^n(M \cup S) - \text{MMS}_i^n(M) \end{aligned}$$

where the last inequality follows from our assumption that  $v_i(S) < \text{MMS}_i^n(M \cup S) - \text{MMS}_i^n(M)$ . It follows that  $v_i(W_k \cap M) > \text{MMS}_i^n(M)$ . As  $k$  was chosen arbitrarily, this implies existence of a partition of  $M$  into  $n$  sets  $(W_k \cap M)_{k \in [n]}$  such that each set has value strictly larger than  $\text{MMS}_i^n(M)$ , a contradiction.  $\square$

Our second lemma is a more technical one. In yet another simplification of notation, we write  $\mu_{ij} = \text{MMS}_i^2(A_{ij}) = \text{MMS}_j^2(A_{ij})$  to mean the maximin fair share of agents  $i$  and  $j$  when dividing exactly the goods only the two of them approve among themselves.

**Lemma 2.** *Let  $N = \{1, 2, 3\}$ , and let  $\mathbf{S} = (S_1, S_2, S_3)$  be a 3-partition of  $A_1$  such that  $v_1(S_r) \geq \text{MMS}_1$  for all  $r \in \{1, 2, 3\}$ . Then there exist distinct  $k, \ell \in \{1, 2, 3\}$  such that*

$$\begin{aligned} c(S_k \cap A_{12}) &\leq \mu_{12}, \text{ and} \\ c(S_\ell \cap A_{13}) &\leq \mu_{13}. \end{aligned}$$

*Proof.* Note that, by the definition of maximin fair share, there cannot be two elements  $k_1, k_2 \in \{1, 2, 3\}$  such that  $c(S_{k_1} \cap A_{12}) > \mu_{12}$  and  $c(S_{k_2} \cap A_{12}) > \mu_{12}$ —this would imply that we could divide  $A_{12}$  into two bundles such that both agents 1 and 2 are guaranteed strictly more than their maximin fair share.

Therefore, there must exist at least two distinct  $k, k' \in \{1, 2, 3\}$  such that both  $c(S_k \cap A_{12}) \leq \mu_{12}$  and  $c(S_{k'} \cap A_{12}) \leq \mu_{12}$ . The same argument tells us there are distinct  $\ell, \ell' \in \{1, 2, 3\}$  such that  $c(S_\ell \cap A_{13}) \leq \mu_{13}$  and  $c(S_{\ell'} \cap A_{13}) \leq \mu_{13}$ . Applying a pigeonhole argument, we conclude there must be distinct  $k, \ell \in \{1, 2, 3\}$  such that  $c(S_k \cap A_{12}) \leq \mu_{12}$  and  $c(S_\ell \cap A_{13}) \leq \mu_{13}$ , as desired.  $\square$

We are now ready to state the main result of this section.

**Theorem 1.** *For three agents with cost utilities, there always exists a Pareto efficient MMS allocation.*

*Proof.* Given a set of agents  $N = \{1, 2, 3\}$ , let  $MMS_i = MMS_i^3(M)$ —the maximin fair share of agent  $i$  when dividing the items in  $M$  among the three agents. We assume that for any item  $g \in M$ , we have that  $g$  is approved by at least two agents. By Lemma 1, we know the claim will also hold for the remaining cases where there are additional goods approved by a single agent.

Finally, we define the following three values:

$$q_1 = MMS_1 + \mu_{23}$$

$$q_2 = MMS_2 + \mu_{13}$$

$$q_3 = MMS_3 + \mu_{12}$$

Without loss of generality, we assume that  $q_1 \geq q_2$  and  $q_1 \geq q_3$ . We can rewrite this, and express it as follows:

$$MMS_1 + \mu_{23} - \mu_{13} \geq MMS_2 \tag{1}$$

$$MMS_1 + \mu_{23} - \mu_{12} \geq MMS_3 \tag{2}$$

Our method for finding an allocation that satisfies the maximin property and is Pareto efficient, takes as basis a partition of the goods where each bundle reaches the maximin fair share of agent 1. Let  $\mathbf{S} = (S_1, S_2, S_3)$  be a 3-partition of  $A_1$  such that  $v_1(S_r) \geq MMS_1$  for all  $r \in \{1, 2, 3\}$ . Note that such a partition always exists by definition of  $MMS_1$ . By Lemma 2 we know there exist distinct  $k, \ell \in \{1, 2, 3\}$  such that

$$c(S_k \cap A_{12}) \leq \mu_{12}, \tag{3}$$

$$c(S_\ell \cap A_{13}) \leq \mu_{13}. \tag{4}$$

We can now describe the allocation  $\mathbf{B}$ , which we claim is a Pareto efficient MMS allocation.

We divide  $A_{23}$  into two disjoint sets  $T_1$  and  $T_2$  such that  $c(T_1) \geq \mu_{23}$  and  $c(T_2) \geq \mu_{23}$ . Note that such a partition exists by the definition of  $\mu_{23}$ . Let  $S_x$  be the third bundle in  $\mathbf{S}$ —i.e.  $x \in \{1, 2, 3\} \setminus \{k, \ell\}$ . We then allocate the goods in  $M$  as follows:

$$\begin{aligned}
 B_1 &= (S_\ell \setminus A_2) \cup (S_k \setminus A_3) \cup S_x \\
 B_2 &= (S_\ell \cap A_2) \cup T_1 \\
 B_3 &= (S_k \cap A_3) \cup T_2
 \end{aligned}$$

In words, agent 2 receives  $T_1$  and everything in  $S_\ell$  that she wants, agent 3 receives  $T_2$  and everything in  $S_k$  that she wants, and agent 1 receives the remaining items in  $S_k$  and  $S_\ell$  as well as the entire bundle  $S_x$ . Note that all items have been allocated as  $A_1 \cup A_{23} = M$ , and no item is allocated to more than one agent as  $S_x$  and  $A_{23}$  are disjoint. By definition, we have that  $v_1(B_1) \geq \text{MMS}_1$ —agent 1 clearly receives their maximin fair share as she receives one of the original bundles,  $S_x$ , and then some. We now show that the same must hold for the other two agents.

For agent 2, we need to show that  $v_2(B_2) \geq \text{MMS}_2$ . Note that we can express the value of agent 2’s bundle using the cost function  $c$  as follows (where  $S_\ell \cap A_{13}$  is the portion of  $S_\ell$  that agent 2 values at 0).<sup>2</sup>

$$\begin{aligned}
 v_2(B_2) &= v_2(S_\ell \cap A_2) + v_2(T_1) \\
 &= c(S_\ell \cap A_2) + c(T_1) \\
 &= c(S_\ell) - c(S_\ell \cap A_{13}) + c(T_1)
 \end{aligned}$$

Because of the way we’ve defined the partition  $\mathbf{S}$  and  $A_{13}$ , we know that  $c(S_\ell) \geq \text{MMS}_1$  and  $c(T_1) \geq \mu_{23}$ . Additionally, by Eq. 4, we know that  $c(S_\ell \cap A_{13}) \leq \mu_{13}$ . From this, we can conclude the following, where the last inequality follows from Eq. 1.

$$\begin{aligned}
 v_2(B_2) &= c(S_\ell) - c(S_\ell \cap A_{13}) + c(T_1) \\
 &\geq \text{MMS}_1 - \mu_{13} + \mu_{23} \\
 &\geq \text{MMS}_2
 \end{aligned}$$

Putting this all together, we have shown that  $v_2(B_2) \geq \text{MMS}_2$ , as desired. The proof for agent 3 proceeds analogously, using Eqs. 2 and 3. Thus, we have shown that  $\mathbf{B}$  is an MMS allocation.

Finally, we see that no item has been allocated to an agent who values it at 0, meaning the allocation is indeed Pareto efficient.  $\square$

Theorem 1 establishes a clear improvement when dealing with cost utilities over general additive valuations.

### 3.2 MMS Allocations for Laminar Set Approvals

In this section we present our results for agents with laminar set approvals. This restriction on the agents’ preferences has a very natural interpretation, in that it

<sup>2</sup> This is possible because we know that any good in the set is either approved by all three agents, or a subset of two. Agent 2 is a member of any subset of size two except  $A_{13}$ .


describes the notion of items falling into categories and subcategories quite well. We can think of agents as approving categories as a whole. For example, one agent might want all vegetarian dishes, while another wants only the seafood. A third agent might want the pasta-based vegetarian dishes, which would constitute a subcategory of vegetarian.

We say agents with cost utilities have *laminar set approvals* if for a vector  $\mathbf{A} = (A_1, \dots, A_n)$  of approval sets, we have that for any  $i, j \in N$ , either  $A_i \cap A_j = A_j$ ,  $A_i \cap A_j = \emptyset$ , or  $A_i \cap A_j = A_i$ . In words, for any two agents, one approval set is either a subset of the other, or the sets are disjoint. Note that in this paper, we only examine laminar set approvals within the context of cost utilities.

We first present a technical lemma that we will apply inductively in the proof of Theorem 2. Lemma 3 allows us to carry the existence of an MMS allocation from cases where all agents submit the whole set  $M$  of goods as their approval, to cases where fewer and fewer agents do so, until we reach a single agent approving all goods.

**Lemma 3.** *For  $n$  agents with cost utilities and laminar set approvals, and  $k \geq 1$ , if an MMS allocation exists for all instances where  $k+1$  agents approve all items in  $M$ , then an MMS allocation exists for any instance where  $k$  agents approve all items.*

*Proof.* Consider an instance  $\mathcal{I} = (N, M, \mathbf{v})$  where there are  $k \geq 1$  agents whose approval set equals  $M$ . We call this set of agents  $N'$ . Let  $i \in N \setminus N'$  be an agent such that  $A_i \not\subseteq A_j$  for all  $j \in N \setminus N'$  in the instance  $\mathcal{I}$ . Note that such an agent must exist, as agents have laminar set approvals. See Fig. 1 for a visual representation. We will continue to use this figure throughout this proof.



**Fig. 1.** An illustration of the sets involved in the proof of Lemma 3. The largest set is  $M$ —the set of goods. Note how the approval sets  $A_1, \dots, A_k$  are equivalent to the whole set of goods, and the same holds for  $A_{i^*}$  and  $A'_i$ . The approval set  $A_i$  is “one level below” the sets approving all items. The bundle  $B'_{i^*}$ —represented in blue—is the bundle in the allocation  $\mathbf{B}|_{N' \cup \{i\}}$  that is highest valued according to  $v_i$ .

Our aim is to show that there exists an MMS allocation for the instance  $\mathcal{I}$ . To this end, we define a second instance  $\mathcal{I}' = (N, M, \mathbf{v}')$  such that  $A'_i = M$ , and



$A'_j = A_j$  for all agents  $j \neq i$ —i.e. the instance  $\mathcal{I}'$  only differs from  $\mathcal{I}$  in that agent  $i$  now approves all items. Thus, we have  $k + 1$  agents whose approval set is  $M$  in the instance  $\mathcal{I}'$ .


Suppose  $\mathbf{B}'$  is an MMS allocation for  $\mathcal{I}'$ , such an allocation is guaranteed to exist by the assumption of the lemma. We construct an MMS allocation  $\mathbf{B}$  for our initial instance by building on  $\mathbf{B}'$ . We first define  $i^* \in \operatorname{argmax}_{j \in N' \cup \{i\}} v_i(B'_j)$ . This is an agent who gets the highest value bundle in  $\mathbf{B}'_{|N' \cup \{i\}}$  according to  $v_i$ —agent  $i$ 's valuation in the initial instance. Because the value  $n$  is fixed, we will write  $\operatorname{MMS}_i(\mathcal{I})$  to mean  $\operatorname{MMS}_i^n(\mathcal{I})$ . We consider two cases.

*Case 1:* Suppose  $v_i(B'_{i^*}) \geq \operatorname{MMS}_i(\mathcal{I})$ . Then agent  $i$  values agent  $i^*$ 's bundle at least as much as their maximin fair share in the initial instance. We define an allocation  $\mathbf{B}$  and claim that it is an MMS allocation for the instance  $\mathcal{I}$ .

$$B_j = \begin{cases} B'_{i^*} & \text{if } j = i \\ B'_i & \text{if } j = i^* \\ B'_j & \text{otherwise} \end{cases}$$

First note that for any agent  $j \notin \{i, i^*\}$ , their maximin fair share is the same across both instances, and they receive the same bundle under  $\mathbf{B}$  and  $\mathbf{B}'$ . Thus, they receive at least their maximin fair share in the allocation  $\mathbf{B}$ .

We now show the same holds for  $i$  and  $i^*$ . For agent  $i$ , this follows by assumption since  $v_i(B_i) = v_i(B'_{i^*}) \geq \operatorname{MMS}_i(\mathcal{I})$ . For agent  $i^*$  then, we only need to consider when  $i^* \neq i$ . In that case, as  $i^* \in N'$ , we have that  $A_{i^*} = A'_i = M$ . Then agent  $i^*$  must also receive their maximin fair share in the allocation  $\mathbf{B}$ , because  $v_{i^*}(B_{i^*}) = v'_i(B'_i) \geq \operatorname{MMS}_i(\mathcal{I}') = \operatorname{MMS}_{i^*}(\mathcal{I})$ . Note that this holds because the agents have cost utilities, and both  $v_{i^*}$  and  $v'_i$  are equivalent to the cost function  $c$  since  $A_{i^*} = A'_i = M$ . As  $\mathbf{B}$  guarantees everyone at least their maximin fair share, it is an MMS allocation for  $\mathcal{I}$ .



**Fig. 2.** An illustration of the sets involved in Case 2 of the proof of Lemma 3. The possible approval set of agent  $j$  in each case is represented in green.

*Case 2:* Suppose instead that  $v_i(B'_{i^*}) < \operatorname{MMS}_i(\mathcal{I})$ . In this case, agent  $i$  values agent  $i^*$ 's bundle strictly less than their maximin fair share in the initial instance. Recall that  $v_i(B'_j) \leq v_i(B'_{i^*})$  for all  $j \in N' \cup \{i\}$ —agent  $i^*$ 's bundle is still the

“best” one among those in  $\mathbf{B}|_{N' \cup \{i\}}$ . Given our initial assumption, we then have that

$$v_i(B'_j) < \text{MMS}_i(\mathcal{I}) \text{ for all } j \in N' \cup \{i\}. \tag{5}$$

Before we proceed, we will need to define a third instance over only the goods in  $A_i$ . Let  $\mathcal{I}^* = (N, A_i, \mathbf{v})$  be a restriction of the instance  $\mathcal{I}$  to only the items in  $A_i$ —meaning  $A_j^* = A_j \cap A_i$  for all  $j \in N$ . Note that in  $\mathcal{I}^*$ , there are at least  $k+1$  agents whose approval set is  $A_i$ —the initial  $k$  agents who approved all items in  $\mathcal{I}$ , and agent  $i$ . Let  $\mathbf{B}''$  be an MMS allocation for  $\mathcal{I}^*$ . We now proceed with defining an allocation  $\mathbf{B}$  by using both allocations  $\mathbf{B}'$  and  $\mathbf{B}''$ . In particular, we define

$$B_j = (B'_j \setminus A_i) \cup B''_j \text{ for all } j \in N.$$

Note that no item is allocated more than once because  $B''_j \subseteq A_i$  for all  $j \in N$ . We claim that  $\mathbf{B}$  is an MMS allocation for the instance  $\mathcal{I}$ . Because agents have laminar set approvals, there are three possible cases for any agent  $j$ : either *i*)  $A_j \subseteq A_i$ , or *ii*)  $A_j \cap A_i = \emptyset$ , or *iii*)  $A_i \subset A_j$ . See Fig. 2 for a visual representation.

- i) Suppose  $A_j \subseteq A_i$ . Then agent  $j$  was only approving items in  $A_i$  and their approval set remains the same in the restriction  $\mathcal{I}^*$ , implying that their maximin fair share also remains the same in both instances. Additionally, we have that  $v_j(B'_j \setminus A_i) = 0$  given that  $A_j \subseteq A_i$ , and so  $v_j(B_j) = v_j(B''_j)$ . Since  $j$  receives their maximin fair share in  $\mathbf{B}''$ , they also do so in  $\mathbf{B}$ .
- ii) Suppose instead  $A_j \cap A_i = \emptyset$ . Because agent  $j$  does not approve any items in  $A_i$ , we have that  $v_j(B'_j) = v_j(B'_j \setminus A_i)$  and  $v_j(B''_j) = 0$ . Then  $v_j(B_j) = v_j(B'_j)$ , and because  $A'_j = A_j$  their maximin fair share is the same in  $\mathcal{I}$  and  $\mathcal{I}'$ . Thus  $j$  receives their maximin fair share in  $\mathbf{B}$ .
- iii) Finally, suppose  $A_i \subset A_j$ . This is only possible if  $j \in N'$ , meaning  $j$  is one of the agents approving all items. We know that

$$\begin{aligned} v_j(B_j) &= v_j(B'_j \setminus A_i) + v_j(B''_j) \\ &= (B'_j) - v_j(B'_j \cap A_i) + v_j(B''_j) \\ &= (B'_j) - v_i(B'_j) + v_j(B''_j) \end{aligned} \tag{6}$$

Where the last line follows from the fact that agents have cost utilities, meaning  $v_j(B'_j \cap A_i) = v_i(B'_j)$ .

Recall that Eq. 5 tells us  $v_i(B'_j) < \text{MMS}_i(\mathcal{I})$ . This fact, combined with Eq. 6 (and some reshuffling of the terms), tells us it must be the case that

$$v_j(B_j) > v_j(B'_j) - \text{MMS}_i(\mathcal{I}) + v_j(B''_j). \tag{7}$$

Since  $\mathbf{B}''$  is an MMS allocation for  $\mathcal{I}^*$ , it follows that  $v_j(B''_j) \geq \text{MMS}_j(\mathcal{I}^*)$ . Further, since  $A_i \subset A_j$ , and  $\mathcal{I}^*$  is an instance over only  $A_i$ , we have that  $\text{MMS}_j(\mathcal{I}^*) = \text{MMS}_i(\mathcal{I})$ . Thus,  $v_j(B''_j) \geq \text{MMS}_i(\mathcal{I})$ .

We can then transform Eq. 7 as follows:

$$v_j(B_j) > v_j(B'_j) - \text{MMS}_i(\mathcal{I}) + \text{MMS}_i(\mathcal{I}),$$

meaning it must be the case that  $v_j(B_j) > v_j(B'_j)$ . Because we know agent  $j$  has identical valuations in  $\mathcal{I}$  and  $\mathcal{I}'$ , and  $\mathbf{B}'$  is an MMS allocation, we can conclude that agent  $j$  receives at least their maximin fair share in  $\mathbf{B}$ .

Thus we have shown for any agent  $j \in N$  that they receive their maximin fair share in the allocation  $\mathbf{B}$ , meaning it must be an MMS allocation. Since  $\mathcal{I}$  was an arbitrary instance where exactly  $k$  agents submit the approval set  $M$ , this concludes the proof.  $\square$

We can now (finally) present the main result of this section.

**Theorem 2.** *For  $n$  agents with cost utilities and laminar set approvals, there always exists an MMS allocation.*

*Proof.* First, note that given agents with laminar set approvals, if no agent has  $M$  as their approval set, then we can find a  $k$ -partition  $(N_1, \dots, N_k)$  of agents and pairwise disjoint subsets  $M_1, \dots, M_k$  of items such that agents in  $N_\ell$  do not approve any items in  $M \setminus M_\ell$ , and there is an agent  $i \in N_\ell$  such that  $A_i = M_\ell$ . It is clear—because the agents are partitioned such that each partition considers a distinct set of items from  $M$ —that if we find an MMS allocation for each of the  $k$  sub-cases, this gives us an MMS allocation in the global case. Therefore, without loss of generality, we assume for any instance that at least one agent submits  $M$  as their approval set.

Now suppose there are  $n$  agents with cost utilities who all submit  $M$  as their approval set. Then an MMS allocation trivially exists. Applying Lemma 3 inductively, we see that for agents with cost utilities and laminar set approvals, an MMS allocation always exists given that at least one agent submits  $M$  as their approval set.  $\square$

*Remark 1.* If an MMS allocation exists, then an MMS and PO allocation always exists since after each Pareto improvement agent's utility is weakly increasing. Thus, Theorem 2 implies that under cost utilities and laminar set approvals, MMS and PO allocation always exist.

## 4 Strategyproof MMS Allocations

In this section, we study the strategic guarantees possible under cost utilities. We first show that for cost utilities, the *Sequential Allocation* mechanism is strategyproof. Let us first define what we mean by strategyproofness.

An allocation mechanism  $f$  is *manipulable* if there is some agent  $i \in N$  such that  $v_i(f(\mathbf{v}_{-i}, v'_i)_i) > v_i(f(\mathbf{v})_i)$ , where  $(\mathbf{v}_{-i}, v'_i)$  is the valuation that results when  $v_i$  is replaced by  $v'_i$ . In other words, agent  $i$  can misrepresent their preferences by submitting an untruthful valuation  $v'_i$ , thereby getting a more preferred outcome. We say  $f$  is *strategyproof* if it is not manipulable by any agent.

We now define the *Sequential Allocation mechanism* from previous studies [22, 23]. We first define a *picking sequence* as a sequence of agents in  $N$ . Note that

the sequence of agents can be of any length, and any agent might appear multiple times in the sequence. We can think of Sequential Allocation as proceeding sequentially (as the name indicates), through the ordering of agents. At each step, the agent whose turn it is chooses the item with the highest cost that a) is still available and b) is in their approval set. Note that we “force” agents to pick their most wanted item, as reported in their approvals. If there are no remaining items that an agent finds useful then we skip this agent and continue with the next. The mechanism allows some items to remain unallocated only if they are not approved by any agent.

In fact, Sequential Allocation is a family of mechanisms, each defined by the picking sequence. As we will see, the properties of the mechanism also heavily depend on the picking sequence in question. For example, it is well known that Sequential Allocation is not strategyproof in general unless an agent’s picks are all consecutive [23].

In the rest of this section, we will assume that the goods in  $M = \{g_1, \dots, g_m\}$  are ordered from lowest cost to highest cost—i.e.  $c(g_k) \leq c(g_\ell)$  for all  $k < \ell$ .

**Proposition 1.** *For agents with cost utilities, there exists a picking sequence such that Sequential Allocation is strategyproof and results in a Pareto efficient allocation.*<sup>3</sup>

*Proof.* We define a sequence  $S$  of agents of length  $n + 2$ , and a sequence  $T$  of agents where every agent appears exactly once. Let  $S = 1, 2, \dots, n - 1, n, n, n$ , and  $T = n, n - 1, \dots, 2, 1$ . Our picking sequence is  $S$ , followed by  $m$  copies of each element in the sequence  $T$ . We can think of this as running through  $S$ , then letting each agent in  $T$  choose all the items they want when it is their turn in  $T$ . We now show that this gives us a strategyproof mechanism.

It is immediately clear that agent  $n$  has no incentive to manipulate. They cannot move themselves up in the picking sequence, and once it is their turn, they can essentially grab all the items they want.

For any other agent  $i \in N$ , let  $X_i$  be the items remaining immediately before agent  $i$  received their first item, and let  $x$  be the item with highest cost in  $X_i \cap A_i$ . Then, agent  $i$  receives  $x$ . After this, all items in the approval sets of agents  $i + 1, i + 2, \dots, n$  are allocated before agent  $i$  receives all remaining items in  $A_i$ . Thus, agent  $i$  receives the bundle  $x \cup (A_i \cap (X_i \setminus (A_{i+1} \cup A_{i+2} \cup \dots \cup A_n)))$ . Note that the preferences of agent  $i$  do not decide the set  $X_i$ . Hence, by misreporting, agent  $i$  is unable to gain any additional items that they approve.

Note that the final allocation is Pareto optimal because items are only allocated to agents that want them. As agents have cost utilities, all agents who want an item will value it the same. This concludes the proof.  $\square$

We now consider whether there are picking sequences that can give us an MMS allocation along with truthfulness for a restricted number of items. Such

<sup>3</sup> We prove Proposition 1 for a picking sequence used in the proof of Proposition 2, but note that there are simpler picking sequences for which it holds.

a restriction is needed because computation of an agent’s MMS value is NP-hard for an arbitrary number of items, which implies that no picking sequence is guaranteed to output an MMS allocation.<sup>4</sup> We start with a lemma that will be used to prove Proposition 2.

**Lemma 4.** *For  $n$  agents and  $n+2$  goods, let  $|A_i| \geq n+k$  where  $k \in \{0, 1, 2\}$ . The  $(n - k)$ -th most valuable item in  $A_i$  is guaranteed to give agent  $i$  their maximin fair share.*

*Proof.* Note that for any  $n$ -partition of the items in  $A_i$ , there is at most  $k$  bundles that are not singletons, meaning at least  $(n - k)$  of the bundles have just a single item. Any of these bundles will give agent  $i$  their maximin fair share. Of these  $(n - k)$  singleton bundles, the highest possible value for the lowest valued bundle is the cost of the  $(n - k)$ -th most valuable item in the agent’s approval set.

**Proposition 2.** *For  $n$  agents with cost utilities, and  $n + 2$  goods, there exists a picking sequence such that Sequential Allocation is strategyproof, and returns a Pareto efficient MMS allocation.*

*Proof.* We first show that there is a picking sequence such that Sequential Allocation returns an MMS allocation. If an agent approves fewer than  $n$  items, they still receive their maximin fair share even when no items are allocated to them. We therefore focus on agents who approve at least  $n$  items. We define the picking sequence based on the cost of the items in  $M$ .

- ▶ If  $c(g_4) > c(g_2) + c(g_3)$ , our picking sequence is  $1, 2, \dots, n - 1, n, n, n$ .
- ▶ Otherwise, our picking sequence is  $1, 2, \dots, n - 1, n, n, n - 1$ .

Note that these differ only in who gets to pick the last item. The fact that agents 1 through  $n - 2$  are guaranteed their maximin fair share for both picking sequences follows from Lemma 4. It remains to show that the same holds for agent  $n - 1$  and agent  $n$ . If agent  $n - 1$  or agent  $n$  approve at most  $n$  items, then we already know they are guaranteed their maximin fair share. If agent  $n - 1$  approves  $n + 1$  items, their  $(n - 1)$ -th most valuable item is still up for grabs, and by Lemma 4 this will guarantee them their maximin fair share.

We now consider what happens when agent  $n$  approves  $n + k$  items—for  $k \in \{1, 2\}$ , and when agent  $(n - 1)$  approves  $n + 2$  items. We look at each potential picking sequence separately.

*Case 1:* Suppose  $c(g_4) > c(g_2) + c(g_3)$ . If agent  $n$  approves  $n + k$  items, they will receive at least  $k + 1$  items, as they pick last and can pick up to three items if they want, given the picking sequence  $1, 2, \dots, n - 1, n, n, n$ . Clearly a bundle of size  $k + 1$  guarantees them their maximin fair share.

What remains is to check what happens when agent  $(n - 1)$  approves all items in  $M$ , so suppose this to be the case. We first show that the maximin fair share of agent  $n - 1$  is  $\min(c(\{g_1, g_2, g_3\}), c(g_4))$ . Consider a partition  $\mathbf{B}$  of  $M$  into  $n$  bundles, where  $c(B_i) \geq \text{MMS}_{n-1}^n(M)$  for each  $i \in N$ . At least

---

<sup>4</sup> This is under the assumption  $P \neq \text{NP}$ .

$n - 2$  of these bundles must contain a single item, and so we know that either *i*)  $n - 2$  bundles contain one item and two bundles contain two items, or *ii*)  $n - 1$  bundles contain one item and one bundle contains three items. We know that  $c(g_4) > c(g_2) + c(g_3)$  by assumption, and the non-singleton bundles will be made up of the four lowest value items— $g_1, \dots, g_4$ . Then the best we can do is one 3-item bundle  $B = \{g_1, g_2, g_3\}$  and all the remaining items in singleton bundles. It follows that the maximin fair share of agent  $n - 1$  is  $\min(c(B), c(g_4))$ . When it is agent  $n - 1$ 's turn to pick, in the worst case, the only remaining goods will be  $g_1, \dots, g_4$ , in which case agent  $n - 1$  can pick item  $g_4$  to guarantee their maximin fair share.

*Case 2:* Suppose instead that  $c(g_4) \leq c(g_2) + c(g_3)$ . If agent  $n$  approves  $n + 1$  items, they will receive two items, guaranteeing them their maximin fair share. If agent  $n$  approves all items in  $M$ , their maximin fair share in this case is determined by the lowest value bundle between the two bundles of size two, and the cheapest singleton. In particular, agent  $n$ 's maximin fair share is  $\min(c(\{c_1, c_4\}), c(\{c_2, c_3\}), c(\{g_5\}))$ . With this picking sequence, agent  $n$  receives two items and in the worst case, this will be the bundle  $B = \{g_2, g_3\}$ . Clearly this guarantees agent  $n$  their maximin fair share.

Finally, we look at when agent  $(n - 1)$  approves  $n + 2$  items. In this case, we know that their maximin fair share is determined by  $\min(c(\{g_1, g_4\}), c(\{g_2, g_3\}), c(\{g_5\}))$ , as was the case for agent  $n$ . As we did for agent  $n$  we know that agent  $(n - 1)$  will receive two items, and in the worst case this will be the bundle  $B = \{g_4, g_1\}$ , which gives the agent their maximin fair share.

Strategyproofness and Pareto efficiency for the first case follows directly from Proposition 1. We now prove strategyproofness and Pareto efficiency for the second case, where  $c(g_4) \leq c(g_2) + c(g_3)$ . In this case, our picking sequence is  $1, 2, \dots, n - 1, n, n, n - 1$ .

For any agent  $i \in N$ , if  $i < n - 1$ , it is clear that there is no way for the agent to manipulate as they only get one pick. For agent  $n$ , because their picks are right after each other, they also have no incentive to manipulate. Thus, we need only consider agent  $n - 1$ . Let  $X$  be the items remaining immediately before agent  $n - 1$  received their first item, and let  $x$  be the item with highest cost in  $X \cap A_{n-1}$ . Agent  $n - 1$  will pick  $x$  by definition of the mechanism. Agent  $n$  then receives their two highest valued remaining items if they exist (call these items  $y$  and  $y'$ ), and then finally agent  $n - 1$  potentially receives the last item they approve (call this item  $z$ ).

First, consider the case where agent  $n - 1$  misreports that they approve some item  $x'$ , and they receive  $x'$  instead of  $x$ . Then, the bundle of agent  $n - 1$  will consist of  $x'$  (which they value at 0), and potentially some other item  $z'$  with  $v_{n-1}(z') \leq v_{n-1}(x)$ . Thus, agent  $n - 1$  is not better off in this case. Otherwise, if agent  $n - 1$  instead misreports that they do not approve item  $x$ , then they will pick some other item  $x''$  instead, where  $c(x'') \leq c(x)$ . If  $x'' \neq y$  and  $x'' \neq y'$ , then we must have  $v_{n-1}(x'') \leq v_{n-1}(z)$ , and so agent  $n - 1$  is not better off. Otherwise, if  $x'' = y$  or  $x'' = y'$ , then agent  $n - 1$  will have strictly fewer options

for their final pick (compared to the case where they do not misreport), and so they are still not any better off.

Therefore, the mechanism is strategyproof. It is clear that no agent is assigned an item they do not want, and all items that are wanted by at least one agent are assigned to someone. Thus the allocation is Pareto efficient.  $\square$

We remark that Proposition 2 is tight in the sense that it no longer holds when there are  $n$  agents and  $n + 3$  items.

**Proposition 3.** *For agents with cost utilities, there exists an instance with  $n = 2$  agents and  $m = 5$  goods such that no strategyproof mechanism can guarantee a Pareto efficient MMS allocation.*

*Proof.* Let  $n = 2$ , and  $M = \{g_2, g_3, g_4, g_5, g_6\}$  such that  $c(g_i) = i$ . We will show that no allocation mechanism can satisfy strategyproofness while also guaranteeing a Pareto Efficient MMS allocation. Our aim is to start from an instance  $\mathcal{I}_1$  and—by repeatedly applying the three axioms—reach a contradiction.

First, consider the instance  $\mathcal{I}_1$ , where both agents approve all items—this corresponds to the top row of Table 1. Then, their maximin fair share is 10, and the only way to reach an MMS allocation is to give  $g_4$  and  $g_6$  to one agent, and  $g_2, g_3$  and  $g_5$  to another. Suppose without loss of generality that  $\{g_2, g_3, g_5\}$  is allocated to agent 1, and  $\{g_4, g_6\}$  is allocated to agent 2. We will consider 5 further instances.

$\mathcal{I}_2$  differs only on agent 2’s approval set—they now only approve items  $g_4, g_5$ , and  $g_6$ . By strategyproofness, agent 2 must still receive a bundle she values at 10. If this were a higher value the agent could manipulate from  $\mathcal{I}_1$ , and if it were lower, they could manipulate from  $\mathcal{I}_2$  to  $\mathcal{I}_1$ .

Instance  $\mathcal{I}_3$  differs from instance  $\mathcal{I}_2$  only on agent 1’s approval set—they now only approve items  $g_3, g_4, g_5$ , and  $g_6$ . As agent 1 is the only one approving item  $g_3$ , they must be allocated this item by Pareto efficiency. The maximin value of agent 1 in this instance is 9, so they must receive one of the following bundles:  $\{(g_3, g_6)\}$ ,  $\{g_3, g_4, g_5\}$ ,  $\{g_3, g_4, g_6\}$ ,  $\{g_3, g_5, g_6\}$ , or  $\{g_3, g_4, g_5, g_6\}$ . All but  $\{g_3, g_6\}$  break strategyproofness, as agent 1 would have an incentive to manipulate from  $\mathcal{I}_2$  to  $\mathcal{I}_3$ .

Instance  $\mathcal{I}_4$  differs from instance  $\mathcal{I}_3$  only on agent 1’s approval set—they now only approve item  $g_6$ . Agent 1 must be allocated  $g_6$ . If this were not the case, they would have an incentive to manipulate from  $\mathcal{I}_4$  to  $\mathcal{I}_3$  as they do receive item 6 in that instance.

Instance  $\mathcal{I}_5$  differs from instance  $\mathcal{I}_4$  only on agent 1’s approval set—they now approve items  $g_2, g_3$  and  $g_6$ . As agent 1 is the only one approving items  $g_2$  and  $g_3$ , they must be allocated these items by Pareto efficiency. If agent 1 is not also given  $g_6$ , they would have an incentive to manipulate from  $\mathcal{I}_4$  to  $\mathcal{I}_3$  as their bundle in that instance is valued at 6 (which is greater than  $2 + 3$ , the value of the bundle  $\{g_2, g_3\}$ ). Note that this gives them a bundle valued at 11.

Finally, instance  $\mathcal{I}_6$  differs from instance  $\mathcal{I}_5$  only on agent 1’s approval set—they now approve all items. If agent 1 is given a bundle valued lower than 11, they would have an incentive to manipulate from  $\mathcal{I}_6$  to  $\mathcal{I}_5$ . Note however that

$\mathcal{I}_6 = \mathcal{I}_2$ , and our axioms dictated in that instance that agent 1 must receive utility of 10. This gives us our contradiction.  $\square$

**Table 1.** Table showing the approval sets corresponding to each instance in the proof of Proposition 3. For example, (23456)(456) denotes the instance where agent 1 approves all items, and agent 2 approves items  $g_4$ ,  $g_5$ , and  $g_6$ . The second column describes outcomes consistent with MMS, Pareto efficiency, and strategyproofness. Note that we omit items not approved by either agents, as they can be allocated to anyone without affecting any of the three axioms.

Instance	Approval Sets	Allocation
$\mathcal{I}_1$	(23456)(23456)	(235)(46)
$\mathcal{I}_2$	(23456)(456)	(235)(46)
$\mathcal{I}_3$	(3456)(456)	(36)(45)
$\mathcal{I}_4$	(6)(456)	(6)(45)
$\mathcal{I}_5$	(236)(456)	(236)(45)
$\mathcal{I}_6$	(23456)(456)	(at least 11) (at most 9)

## 5 Conclusion

Fair division of indivisible resources is a challenging yet important problem with wide-ranging applications. In this paper, we have established that cost utilities are a useful restriction to study, especially in the context of MMS allocations. We have shown that there are several classes of instances where MMS allocations always exist under cost utilities. We also show that cost utilities are helpful in circumventing problems of strategic manipulation.

The topic of MMS allocations in general, and for cost utilities in particular, poses many challenging questions. One might consider various fair division problems with constraints under cost utilities. A prime example is cardinality constraints—or more generally, budget constraints—which are quite natural in this setting.

Our work serves as a further indication that fair division under cost utilities is a fruitful research direction.

**Acknowledgements.** This project was partially supported by the ARC Laureate Project FL200100204 on “Trustworthy AI”.

## References

1. Akrami, H., Rezvan, R., Seddighin, M.: An EF2X allocation protocol for restricted additive valuations. In: Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI) (2022)



2. Amanatidis, G., Birmpas, G., Christodoulou, G., Markakis, E.: Truthful allocation mechanisms without payments: characterization and implications on fairness. In: Proceedings of the 18th ACM Conference on Economics and Computation (EC) (2017)
3. Amanatidis, G., Birmpas, G., Filos-Ratsikas, A., Voudouris, A.A.: Fair division of indivisible goods: a survey. In: Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI) (2022)
4. Amanatidis, G., Markakis, E., Nikzad, A., Saberi, A.: Approximation algorithms for computing maximin share allocations. *ACM Trans. Algorithms* **13**(4), 1–28 (2017b)
5. Asadpour, A., Feige, U., Saberi, A.: Santa Claus meets hypergraph matchings. *ACM Trans. Algorithms* **8**(3), 1–9 (2012)
6. Babaioff, M., Ezra, T., Feige, U.: Fair and truthful mechanisms for dichotomous valuations. In: Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI) (2021)
7. Bansal, N., Sviridenko, M.: The Santa Claus problem. In: Proceedings of the 38th Annual ACM Symposium on Theory of Computing, p. 31–40 (2006)
8. Bouveret, S., Cechlárová, K., Elkind, E., Igarashi, A., Peters, D.: Fair division of a graph. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI) (2017)
9. Bouveret, S., Lemaître, M.: Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Auton. Agents Multi-Agent Syst.* **30**(2), 259–290 (2016)
10. Brams, S.J., Taylor, A.D.: Fair Division: From cake-cutting to dispute resolution. Cambridge University Press (1996)
11. Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D.: Handbook of Computational Social Choice. Cambridge University Press (2016)
12. Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. *J. Polit. Econ.* **119**(6), 1061–1103 (2011)
13. Camacho, F., Fonseca-Delgado, R., Pino Pérez, R., Tapia, G.: Generalized binary utility functions and fair allocations. *Math. Soc. Sci.* **121**, 50–60 (2022)
14. Cheng, S., Mao, Y.: Integrality gap of the configuration LP for the restricted maximin fair allocation (2018), <http://arxiv.org/abs/1807.04152>
15. Ebadian, S., Peters, D., Shah, N.: How to fairly allocate easy and difficult chores. In: Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2022)
16. Feige, U., Sapir, A., Tauber, L.: A tight negative example for MMS fair allocations. In: Proceedings of the 17th International Conference on Web and Internet Economics (WINE) (2021)
17. Garg, J., Taki, S.: An improved approximation algorithm for maximin shares. In: Proceedings of the 21st ACM Conference on Economics and Computation (EC) (2020)
18. Ghodsi, M., Hajiaghayi, M., Seddighin, M., Seddighin, S., Yami, H.: Fair allocation of indivisible goods: Improvements and generalizations. In: Proceedings of the 19th ACM Conference on Economics and Computation (EC) (2018)
19. Goldman, J., Procaccia, A.D.: Spliddit: unleashing fair division algorithms. *SIGecom Exchanges* **13**(2), 41–46 (2015)
20. Halpern, D., Procaccia, A.D., Psomas, A., Shah, N.: Fair division with binary valuations: one rule to rule them all. In: Proceedings of the 16th International Conference on Web and Internet Economics (WINE), Springer (2020)

21. Heinen, T., Nguyen, N.-T., Nguyen, T.T., Rothe, J.: Approximation and complexity of the optimization and existence problems for maximin share, proportional share, and minimax share allocation of indivisible goods. *Auton. Agents Multi-Agent Syst.* **32**(6), 741–778 (2018). <https://doi.org/10.1007/s10458-018-9393-0>
22. Kalinowski, T., Narodytska, N., Walsh, T.: A social welfare optimal sequential allocation procedure. In: *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)* (2013)
23. Kalinowski, T., Narodytska, N., Walsh, T., Xia, L.: Strategic behavior when allocating indivisible goods sequentially. In: *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)* (2013)
24. Kurokawa, D., Procaccia, A.D., Wang, J.: Fair enough: guaranteeing approximate maximin shares. *J. ACM* **65**(2), 1–27 (2018)
25. Moulin, H.: *Fair division and collective welfare*. MIT press (2004)
26. Othman, A., Sandholm, T., Budish, E.: Finding approximate competitive equilibria: efficient and fair course allocation. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2010)
27. Vossen, T.: *Fair allocation concepts in air traffic management*. PhD thesis (2002)