# Maximin Share Guarantees via Limited Cost-Sensitive Sharing

AAAI Track

Hana Salavcova
Charles University
Prague, Czech Republic
salavcovah@gmail.com

Martin Černý
Charles University
Prague, Czech Republic
cerny@kam.mff.cuni.cz

Arpita Biswas
Rutgers University
New Jersey, United States
a.biswas@rutgers.edu

## ABSTRACT

We study the problem of fairly allocating indivisible goods when limited sharing is allowed, that is, each good may be allocated to up to $k$ agents, while incurring a cost for sharing. While classic maximin share (MMS) allocations may not exist in many instances, we demonstrate that allowing controlled sharing can restore fairness guarantees that are otherwise unattainable in certain scenarios. (1) Our first contribution shows that exact maximin share (MMS) allocations are guaranteed to exist whenever goods are allowed to be cost-sensitively shared among at least half of the agents and the number of agents is even; for odd numbers of agents, we obtain a slightly weaker MMS guarantee. (2) We further design a Shared Bag-Filling Algorithm that guarantees a $(1-C)(k-1)$-approximate MMS allocation, where $C$ is the maximum cost of sharing a good. Notably, when $(1-C)(k-1) \geq 1$, our algorithm recovers an exact MMS allocation. (3) We additionally introduce the *Sharing Maximin Share* (SMMS) fairness notion, a natural extension of MMS to the $k$-sharing setting. (4) We show that SMMS allocations always exist under identical utilities and for instances with two agents. (5) We construct a counterexample to show the impossibility of the universal existence of an SMMS allocation. (6) Finally, we establish a connection between SMMS and constrained MMS (CMMS), yielding approximation guarantees for SMMS via existing CMMS results. These contributions provide deep theoretical insights for the problem of fair resource allocation when a limited sharing of resources are allowed in multi-agent environments.

## KEYWORDS

fair allocation, maximin share, shared resources, cost of sharing, fair division

**ACM Reference Format:**
Hana Salavcova, Martin Černý, and Arpita Biswas. 2026. Maximin Share Guarantees via Limited Cost-Sensitive Sharing: AAAI Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. https://doi.org/10.65109/LGSS5881

## 1 INTRODUCTION

The theory of fair allocation of indivisible items has been receiving significant attention, driven both by theoretical interest and by a wide range of practical applications [4, 5, 30]. A substantial body of work has explored various generalizations, relaxations, and constraints that arise in real-world allocation problems. For instance, several studies have addressed settings where not all items must be allocated, leading to models where certain fairness guarantees exist, which is otherwise unattainable [8, 14, 16, 17]. These extensions reflect a growing recognition of the need for fair allocation settings and methods that are both flexible and applicable to complex, structured domains.

One of the most prominent fairness notions in the area of fair division of indivisible goods is *maximin share* (MMS) fairness, where each agent is guaranteed a bundle at least as much as their maximin share value. Informally, an agent's maximin share value represents the maximum utility they can guarantee themselves by partitioning the goods into as many bundles as there are agents, under the assumption that they receive the least valuable bundle. This concept captures a natural and compelling benchmark for fairness, as it guarantees what an agent could ensure for themselves in a worst-case, divide-and-choose scenario. MMS was formally introduced in [13] in the context of course allocation.

Despite its intuitive appeal, MMS allocations are not guaranteed to exist in general. In fact, counterexamples with as few as three agents and a small number of goods have shown that no MMS exists [18, 25]. This inherent impossibility has led researchers to focus on approximate MMS allocations, where each agent receives at least a fraction of their MMS value. The study of such approximations has become a central direction in the literature, yielding progressively stronger guarantees over time [6, 18, 20, 21] with the best known approximation factor of $\left(\frac{3}{4} + \frac{3}{3836}\right)$ under additive valuations established by Akrami and Garg [2].

Much of the prior work on fair allocation assumes that the set of items must be partitioned disjointly among the set of agents, meaning that no good can be shared among multiple agents. However, many practical scenarios, such as allocation of shared computing access [23] and community energy storage systems [15], may allow the sharing of resources in a structured way. For instance, consider allocating access to high-demand laboratory equipment at a university, such as a high-resolution electron microscope, a specialized DNA sequencer, and a powerful computing server. These devices are expensive, scarce, and cannot be physically subdivided, but each can be shared among a limited number of users via scheduled access time slots or usage quotas. However, sharing may come at a cost, and the effective value of a resource to each user can diminish as more users share it. A purely non-sharing allocation may leave some researchers without access to any resource, failing to meet basic fairness expectations. In contrast, allowing resources to be

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus.* © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/LGSS5881

---

The full version of the paper, which includes complete proofs and additional results, is available here: https://arxiv.org/abs/2602.20541.

shared in a structured, bounded manner can help achieve fairness even when sharing incurs a cost, motivating the study of fair allocation under bounded, cost-sensitive sharing. This leads to the central question:

*Can a fair allocation be obtained by relaxing the classical no-sharing constraint in structured ways, particularly in settings where fairness is otherwise unattainable?*

In this work, we allow each good to be allocated to at most $k$ agents, while explicitly accounting for the costs incurred due to sharing. This framework significantly broadens the applicability of fair allocation, making it more aligned with real-world scenarios where limited sharing is often inevitable or desirable. A consequence of the result by Akrami et al. [1] and Barman et al. [7] implies that sharing among two agents is sufficient to ensure MMS fairness, assuming no cost is incurred when sharing goods. However, including the cost of sharing makes the model more realistic, as sharing typically reduces the benefit that a good provides to each user. Therefore, we formalize the $k$-sharing fair allocation problem and introduce models to capture the cost of sharing (in Section 4). Our main theoretical contributions are as follows:

1. In Section 5, we show that exact maximin share (MMS) allocations can be achieved when goods are allowed to be cost-sensitively shared among at least half of the agents and when there is an even number of agents, under certain cost models. When agents are odd in number, we establish a slightly weaker MMS guarantee (Theorem 5.1). Further, for any cost-sharing model, we provide Shared Bag-Filling Algorithm for computing approximate MMS allocations in the shared setting (Theorem 5.4).
2. In Section 6, we introduce a stronger fairness notion for $k$-sharing allocations, *Sharing Maximin Share (SMMS)*. We show that SMMS always exists when there are only two agents (Proposition 6.2) and under identical utilities (Proposition 6.3). Interestingly, the existing MMS counterexamples [18, 25] admit feasible SMMS allocations, prompting us to construct a new counterexample to disprove universal SMMS existence (Theorem 6.5).
3. Additionally, we establish a connection between SMMS for shared settings and constrained MMS (CMMS), yielding approximation guarantees for SMMS (Proposition 7.3).

## 2 RELATED WORK

The literature on discrete fair allocation is extensive and rapidly evolving, typically operating under two key assumptions: *non-shareability*, where an agent's allocation is disjoint from others, and *completeness*, where all items must be allocated among all the agents. Recent survey articles, such as those by Walsh [30], Aziz et al. [5], and Amanatidis et al. [4], offer comprehensive insights into the algorithmic and complexity dimensions of the area of fair allocation. In addition to these broad overviews, more focused surveys have explored specific subdomains: for instance, Aleksandrov and Walsh [3] examined fairness in dynamic resource allocation, while Suksompong [29] and Biswas et al. [12] investigate fairness under various structural set constraints.

Despite significant advances, the well-studied maximin share (MMS) [13] fairness guarantee does not hold under the standard assumptions of non-shareability and completeness [18, 24, 25]. Recent research has made progress on achieving various fairness guarantees by relaxing these classical assumptions through innovative approaches: (1) allowing some items to remain unallocated (as *charity*) [11, 17], (2) treating items expandable such as school seats or course seats, and allowing their capacities to be increased [26, 28], (3) considering some items to be divisible [9, 27] and (4) creating duplicate copies of items to enable multi-allocation [1, 7]. Barman et al. [7] show that MMS fairness can be achieved under additive valuations by duplicating each item at most once, and Akrami et al. [1] show that MMS can be achieved with one duplicate copy of at most $\lfloor n/2 \rfloor$ items. While these results imply that allowing sharing among two agents helps achieving MMS fairness, the results fall short when the cost of sharing is considered. To the best of our knowledge, no prior work has considered cost-sensitive shared fair allocation setting.

## 3 PRELIMINARIES

An instance of a discrete fair allocation problem is typically represented by a triple $I = (N, M, v)$ where $N$ is a set of $n$ agents, $M$ is a set of $m$ indivisible items, and $v = (v_1, \ldots, v_n)$ is a valuation profile, where each $v_i : 2^M \to \mathbb{R}$ is a valuation function for agent $i$ assigning a value to every subset of items $S \subseteq M$. In this work, we assume non-negative and additive valuations, that is, $v_i(S) \in \mathbb{R}_0^+$ for all $S \subseteq M$ and for each agent $i \in N$, and $v_i(S) = \sum_{g \in S} v_i(\{g\})$. Since we assume non-negative valuations, we will refer to items as *goods*. For brevity, we will write $v_i(g)$ instead of $v_i(\{g\})$ throughout the paper.

In the classical setting, an allocation is a tuple $A = (A_1, A_2, \ldots, A_n)$ where each $A_i \subseteq M$ is the bundle allocated to an agent $i$ and satisfy two properties, namely (1) *non-shareability*: $A_i \cap A_j = \emptyset$ for all $i \neq j$ and (2) *completeness*: $\bigcup_{i \in N} A_i = M$. In other words, an allocation corresponds to a partition of the set of goods $M$ into $n$ disjoint subsets, with no good assigned to more than one agent.

In this work, we focus on *maximin share* (MMS) fairness [13], which is one of the most well-studied notions in discrete fair allocation literature. It guarantees each agent at least their *maximin share value* ($\text{MMS}_i^n(M)$), that is, $v_i(A_i) \geq \text{MMS}_i^n(M)$ where

$$\text{MMS}_i^n(M) := \max_{(A_1,\ldots,A_n) \in \mathcal{A}^n(M)} \min_{j \in [n]} v_i(A_j), \tag{1}$$

and $\mathcal{A}^n(M)$ denotes the set of all possible $n$-partitions of the set of goods $M$. In other words, the *maximin share value* $\text{MMS}_i$ for an agent $i$ is the maximum value she can guarantee for herself by partitioning the goods into $n$ bundles and receiving the least valued bundle.

While MMS often fails to exist in discrete fair allocation problems under the standard assumptions of non-shareability, we reveal new pathways to fairness by introducing cost-sensitive restricted sharing. We now formalize the problem setting in the subsequent section.

## 4 PROBLEM FORMULATION

We investigate fair allocation settings where each good is allowed to be shared among a fixed number of agents $k$. We introduce *$k$-sharing allocations.*

*Definition 4.1.* A *$k$-sharing allocation* is a tuple $A = (A_1, \dots, A_n)$ where the bundles $A_i \subseteq M$ satisfy two properties, namely (1) *$k$-limited shareability*: $\bigcap_{i \in S} A_i = \emptyset$, for every group $S \subseteq N$ of size $|S| > k$ and (2) *completeness*: $\bigcup_{i \in N} A_i = M$. Further, $\mathcal{A}_k^n(M)$ denotes the set of all possible $k$-sharing allocations of goods $M$.

The *$k$-limited shareability* condition ensures that each good can be shared among at most $k$ agents. To denote the set of agents sharing a good $g$ under allocation $A$, we use

$$N_g(A) := \{i \in N : g \in A_i\}.$$

According to the *$k$-limited shareability* condition, we are required to satisfy $|N_g(A)| \leq k$. Additionally, we say that a $k$-sharing allocation is *fully-shared* if every good is shared by exactly $k$ agents, that is, $|N_g(A)| = k$ for all $g \in M$.

In the classical (1-sharing) model, goods are allocated exclusively, and each agent's utility is determined solely by the goods in her bundle. However, in the $k$-sharing setting, the utility an agent derives from a good may depend not only on having access to it but also on how many share it. Thus, rather than defining utility based only on the agent's individual bundle, we consider utility functions $u_i$ that take the full allocation $A$ as input.

*Definition 4.2.* The *utility* of an agent $i$ in a k-sharing setting, $u_i \colon \mathcal{A}_k^n(M) \to \mathbb{R}_0^+$, is defined as

$$u_i(A) = \sum_{g \in A_i} \left[1 - c_{i,g}(N_g(A))\right] \cdot v_i(g) \tag{2}$$

with $c_{i,g}(N_g(A)) \in [0, 1]$ denoting the cost incurred by agent $i$ for sharing good $g$ with agents $N_g(A)$. Additionally,

$$c_{i,g}(N_g(A)) = 0 \text{ if } |N_g(A)| = 1. \tag{3}$$

In this work, we consider *goods-based* cost models, where for every agent $i \in N$ and good $g \in M$, the cost depends only on the number of agents sharing $g$, and is given by $c_{i,g}(N_g(A)) = c_g(|N_g(A)|) \in [0, 1]$.

1. **Cost-free sharing:** $c_g(|N_g(A)|) = 0$ for every $g \in M$ (no cost of sharing). Here, the utility of agent $i$ equals the valuation of their allocated bundle (including their shared goods), $u_i(A) = v_i(A_i)$.
2. **Equal-share cost-sharing:** $c_g(|N_g(A)|) = 1 - \frac{1}{|N_g(A)|}$ for all $g \in M$. Therefore, each agent receives a utility of $\frac{1}{|N_g(A)|}$ fraction of her valuation for the good $g$ allocated to her, $u_i(A) = \sum_{g \in A_i} \frac{1}{|N_g(A)|} \cdot v_i(g)$.
3. **Generous cost-sharing:** $c_g(|N_g(A)|) \in [0, 1 - \frac{1}{|N_g(A)|}]$ for every $g \in M$. Further, we assume costs are non-decreasing in the number of sharers: $c_g(\ell') \leq c_g(\ell)$ for $\ell' < \ell$. When $c_g(|N_g(A)|) = 0$ for all $g \in M$, the model becomes *cost-free.* When $c_g(|N_g(A)|) = 1 - \frac{1}{|N_g(A)|}$ for all $g \in M$, it represents *equal-share.*

Using all the components, we now denote a cost-sensitive $k$-sharing fair allocation instance as a tuple $(N, M, k, v, c)$ where $N$ is the set of $n$ agents, $M$ is the set of $m$ goods, $k$ is the sharing constraint, $v = (v_i, \dots, v_n)$ is the valuation profile, and the sharing-cost function is denoted by $c = (c_g(\ell)))_{g \in M, \ell \in \{1, \dots, k\}}$.

## 5 MMS UNDER COST-SENSITIVE SHARING

A $k$-sharing allocation $A \in \mathcal{A}_k^n(M)$ satisfies MMS if $u_i(A) \geq \mathrm{MMS}_i^n(M)$ for all $i \in N$ (as defined in Equation 1). While sharing increases the flexibility to achieve fairer outcomes—by allowing agents to receive more (potentially shared) goods—it also introduces costs that complicate the structure of the utility function. For any given good, the utility an agent derives can be, due to the costs, lower when the good is shared than when it is received exclusively.

We now highlight the generalizability of the equal-share cost-sharing model in terms of satisfying MMS, or any other threshold-based fairness notions.

Observation 1. *In a k-sharing setting, if an allocation satisfies MMS under equal-share cost-sharing model, then the allocation also satisfies MMS under any* generous *cost-sharing model, including cost-free sharing model.*

Observation 1 follows directly from the definitions of the cost-sharing models, implying that an agent's utility for a given $k$-sharing allocation under the *equal-share* model $c_g(|N_g(A)|) = 1 - \frac{1}{|N_g(A)|}$ increases when we consider lesser cost of sharing, $0 \leq c_g(|N_g(A)|) < 1 - \frac{1}{|N_g(A)|}$.

Therefore, we focus on the equal-share cost model to investigate the existence of MMS in the cost-sensitive $k$-sharing setting. Intuitively, allowing goods to be shared among more agents increases the likelihood of achieving a fair allocation. Indeed, when $k = n$, each item can be shared among all agents, trivially satisfying MMS.

### Sharing among at Least Half of the Agents

Assuming $k \geq n/2$ under equal-share cost-sharing model, we establish the existence of MMS when $n$ is even, with a weaker MMS guarantee when $n$ is odd.

Theorem 5.1. *In every $k$-sharing instance $(N, M, k, v, c)$ under equal-share cost-sharing model, with $k \geq n/2$ agents, there exists an allocation $A$ that for every agent i satisfies:*

- $u_i(A) \geq \mathrm{MMS}_i^n(M)$ *if $n$ is even,*
- $u_i(A) \geq \mathrm{MMS}_i^{n+1}(M)$ *if $n$ is odd.*

Proof. We provide a constructive proof for both the cases. First, we consider $n$ to be even. We arrange all the agents into $\ell = n/2$ disjoint pairs. For each pair of agents, $\{i, j\}$, we define a classical (no-sharing) fair allocation instance, $I_{i,j} = (N = \{i, j\}, M, v = (v_i, v_j))$. Since MMS allocation can always be obtained when there are 2 players, we find one such allocation and denote it as $A_{i,j} = (A_i, A_j)$. Note that the allocated bundles between each pair are disjoint, $A_i \cap A_j = \emptyset$. Now, let us consider all such allocated bundles from all $\ell$ pairs of agents $A = (A_1, \dots, A_n)$. Since the allocated bundles between each of the $\ell$ pairs were disjoint, each good $g$ appears in exactly $\ell$ different bundles in $A$, and the utility is reduced by a factor of $1/\ell$. Therefore, for any agent $i$, the utility of $A_i$ upon sharing each good with $\ell$ agents, under equal-share cost-sharing model is:

$$\begin{aligned} u_i(A) &= \sum_{g\in A_i} (1 - c_g(\ell)) \cdot v_i(g) \\ &= \sum_{g\in A_i} \frac{1}{\ell} \cdot v_i(g) = \frac{1}{\ell} \cdot v_i(A_i) \\ &\ge \frac{1}{\ell} \max_{B\in\mathcal{A}_1^2(M)} \min_{h\in\{i,j\}} v_i(B_h) = \frac{1}{\ell} \cdot \mathrm{MMS}_i^2(M) \end{aligned} \tag{4}$$

Now, let $B^* \in \mathcal{A}_1^{2\ell}(M)$ be the $\mathrm{MMS}_i^n(M)$ maximizer. Let $C$ be the union of the $\ell$ least-valued bundles in $B^*$ according to agent $i$, and let $C' = M \setminus C$. Both $C$ and $C'$ have total value at least $\ell \cdot \min_{h\in N} v_i(B_h^*)$. Therefore, the bipartition $(C, C') \in \mathcal{A}_1^2(M)$ satisfies:

$$\min\{v_i(C), v_i(C')\} \ge \ell \cdot \mathrm{MMS}_i^n(M). \tag{5}$$

Inequality (5) along with the fact that $\mathrm{MMS}_i^2(M) \ge \min\{v_i(D), v_i(D')\}$ for any $(D, D') \in \mathcal{A}_1^2(M)$, and combining with Inequality (4), we obtain $u_i(A) \ge \mathrm{MMS}_i^n(M)$ when $n$ is even.

Next, we consider $n$ to be odd. In this case, we add a *dummy* agent with zero valuation for all goods. With $n + 1$ (even) agents, we apply the construction for even agents, obtaining $k$-sharing allocation of $n + 1$ agents, $A$, which satisfies: $u_i(A) \ge \mathrm{MMS}_i^{n+1}(M)$ for all agents $i \in N$. Finally, the *dummy* agent is removed and its goods are arbitrarily distributed among all the agents. Since this can only increase the utility of each agent in $N$, the result follows. □

Next, we show how to compute an $\alpha$-approximate MMS allocation in polynomial time for any $k \ge 2$, possibly yielding a better approximation than algorithms without sharing.

## Approximate MMS via Cost-Sensitive k-Sharing

Our polynomial-time algorithm, Shared Bag-Filling, (Algorithm 1) works .for *any* cost-sharing model, not only for generous goods-based ones. This formally means that the costs satisfy $c_{i,g}(N_g(A)) \in [0, 1]$ without any additional restrictions. The algorithm computes an $\alpha$-approximate MMS allocation, with the guarantee depending on the maximum sharing cost:

$$C := \max_{i\in N, g\in M, S\subseteq N, |S|\le k} c_{i,g}(S).$$

Notably, Algorithm 1 requires neither the knowledge nor the computation of $C$.

Algorithm 1 builds on the classic bag-filling technique [19–21], adapted to the $k$-sharing setting. It runs in two phases.

Phase 1 allocates *large goods*, i.e., those with $v_i(g) \ge v_i(\tilde{M})/|\tilde{N}|$ for some $i \in \tilde{N}$. Since $g$ already meets agent $i$'s MMS guarantee (in the instance with $\tilde{M}$, $\tilde{N}$), we can assign it directly to her and exclude her from further allocation, as by the monotonicity of MMS (see Appendix A of the full version), this does not reduce the MMS values of the remaining agents on the remaining goods.

Phase 2 begins by normalizing the valuations of the remaining agents $N'$ by setting $v_i^{\mathrm{norm}}(M') = |N'|$, where $M'$ are the remaining goods. This guarantees $\mathrm{MMS}_i^{\mathrm{norm}} \le 1$. In the subsequent steps, $u_i^{\mathrm{norm}}(A) \ge \mathrm{MMS}_i^{\mathrm{norm}}$ is ensured for each remaining agent, which implies the original MMS guarantee. Details are deferred to Appendix A of the full version.

Each good $g \in M'$ is then replaced by $k$ virtual copies (shares), and valuations are scaled so that each share has value $\tilde{v}_i(g) = \frac{1}{k} \cdot v_i^{\mathrm{norm}}(g)$. This transformation preserves total value and ensures that the final outcome corresponds to a fully-shared allocation over $M'$, while reducing the problem to allocating small items.

**Algorithm 1** Shared Bag-Filling Algorithm

```
Require: $I = (N, M, k, v, c)$
Ensure: The output allocation $A$ is $\alpha$-MMS where $\alpha = \min\{1, (1 - C)(k - 1)\}$.
$\tilde{N} \leftarrow N, \tilde{M} \leftarrow M$                // remaining agents and goods

Phase 1
while $\exists i \in N, \exists g \in M$ s.t. $v_i(g) \ge \frac{v_i(\tilde{M})}{|\tilde{N}|}$ do
    $A_i \leftarrow \{g\}$;   $\tilde{M} \leftarrow \tilde{M} \setminus \{g\}$;   $\tilde{N} \leftarrow \tilde{N} \setminus \{i\}$
end while

Phase 2
$v_i^{\mathrm{norm}}(g) \leftarrow \frac{|\tilde{N}|}{v_i(\tilde{M})} \cdot v_i(g)$ for all $i \in \tilde{N}$ and all $g \in \tilde{M}$
Let $\mathcal{M}$ be the multiset of $k$ copies of each good from $\tilde{M}$    // denote the multiplicity of $g$ in $\mathcal{M}$ by $\xi_{\mathcal{M}}(g)$
$\tilde{v}_i(g) = \frac{1}{k} \cdot v_i^{\mathrm{norm}}(g)$ for all $i \in \tilde{N}, g \in \tilde{M}$
while $|\tilde{N}| > 1$ do
    Initialize empty bag $B \leftarrow \emptyset$
    while exists $g \in \mathcal{M}$ s.t. $\xi_{\mathcal{M}}(g) = |\tilde{N}|$ do
        $B \leftarrow B \cup \{g\}$
        $\mathcal{M} \leftarrow \mathcal{M} \setminus \{g\}$            // ensures $\xi_{\mathcal{M}}(g) < |\tilde{N}|$
    end while
    while $\tilde{v}_i(B) < \frac{k-1}{k}$ for all $i \in N$ do
        Select arbitrary $g \in \mathcal{M}$ such that $g \notin B$
        $B \leftarrow B \cup \{g\}$ and $\mathcal{M} \leftarrow \mathcal{M} \setminus \{g\}$
    end while
    Choose any agent $i^* \in \tilde{N}$ such that $\tilde{v}_{i^*}(B) \ge \frac{k-1}{k}$
    $A_{i^*} \leftarrow B$
    $\tilde{N} \leftarrow \tilde{N} \setminus \{i^*\}$         // invariant $\xi_{\mathcal{M}}(g) \le |\tilde{N}|$ holds
end while
$A_j \leftarrow \mathcal{M}$, where $j$ is the remaining agent in $\tilde{N}$
return $A = (A_1, \ldots, A_n)$
```

Phase 2 then proceeds similarly to classic bag-filling, with the constraint that no bag contains multiple shares of the same good. It starts with an empty bag $B$, adds one share of each good whose number of remaining shares equals the number of remaining agents, then continues adding shares arbitrarily until some agent $i$ satisfies $\tilde{v}_i(B) \ge \frac{k-1}{k}$. The bag is allocated to agent $i$, and the process repeats. The last agent receives all remaining goods. For first $|N'| - 1$ agents,

$$\begin{aligned} u_i^{\mathrm{norm}}(A) &\ge (1 - C) \cdot v_i^{\mathrm{norm}}(A_i) = (1 - C) \cdot k \cdot \tilde{v}_i(A_i) \\ &\ge (1 - C) \cdot k \cdot \frac{k-1}{k} = (1 - C) \cdot (k - 1). \end{aligned} \tag{6}$$

Since the valuations are normalized, we have $\mathrm{MMS}_i^{\mathrm{norm}} \le 1$. Therefore, it holds $u_i^{\mathrm{norm}}(A) \ge (1 - C)(k - 1) \cdot \mathrm{MMS}_i^{\mathrm{norm}}$.

There are two issues to address in the analysis of Phase 2. First, it is not immediately clear that the algorithm cannot get stuck—i.e., that at any point, there is always an item that can be added to the current bag when it is not acceptable to every agent. Second, it's not obvious that the last agent, who receives the remaining goods, values them at least at their $(1-C)(k-1)$-MMS value.

The key to understanding why these concerns do not arise lies in the following lemma, which shows that at every iteration of the algorithm there remain shares of sufficient amount and value to satisfy the remaining agents.

Lemma 5.2. *At the beginning and the end of every iteration of Phase 2 (Lines 8-21) of Algorithm 1, for every $i \in \tilde{N}$, it holds*

$$\tilde{v}_i(\mathcal{M}) \geq |\tilde{N}|,$$

*where $\mathcal{M}$ is the multiset of $k$ copies of each good from $\tilde{M}$.*

Proof. Initially, $\tilde{v}_i(\mathcal{M}) = k \cdot \tilde{v}_i(M') = k \cdot \frac{1}{k} \cdot v_i^{\text{norm}}(M') = |N'|$. Assume that the invariant holds at the beginning of some iteration, when the remaining shares are $\mathcal{M}$.

If the while loop at Line 14 does not execute, that is, already the initial bag $B$ (containing exactly one share of each good whose number of remaining shares equals the number of remaining agents) satisfies $\tilde{v}_{i^*}(B) \geq \frac{k-1}{k}$ for some agent $i^*$, then the bag is immediately assigned to agent $i^*$. In this case, the total value of $B$ is at most $\frac{1}{|\tilde{N}|}\tilde{v}_i(\mathcal{M})$ for every agent. This means

$$\tilde{v}_i(\mathcal{M} \setminus B) \geq \tilde{v}_i(\mathcal{M}) - \frac{1}{|\tilde{N}|}\tilde{v}_i(\mathcal{M}) = \frac{|\tilde{N}|-1}{|\tilde{N}|}\tilde{v}_i(\mathcal{M})$$

which is by the assumption larger or equal to $|\tilde{N}| - 1$.

Otherwise, if the loop starting at Line 14 executes at least once, additional goods are added to the bag $B$ according to Lines 15-18. During this process, before the last good was added, the value of $B$ was at most $(k-1)/k$ for every agent. As all shares are valued by all agents at most $1/k$, the value of the bag is at most 1. In both cases, the invariant is preserved. □

Combining Lemma 5.2 with the Phase 2 invariant—that the number of shares of every remaining item is at most the number of remaining agents—we conclude that a dead-end state cannot occur because if there is no share left to add to the bag, then every remaining agent must already prefer the bag.

Proposition 5.3. *Phase 2 does not enter a dead-end state.*

Proof. Suppose that at some point, no share can be added to the current bag $B$ without violating the "distinct goods" condition, while, at the same time, no agent values the bag at least $(k-1)/k$. This means $B$ already contains at least one share of every remaining good in $\mathcal{M}$ and since no good appears more than $|\tilde{N}|$ times in $\mathcal{M}$, it follows for every $i \in \tilde{N}$ that

$$\tilde{v}_i(B) \geq \frac{1}{|\tilde{N}|} \cdot \tilde{v}_i(\mathcal{M}) \geq 1 > \frac{k-1}{k}.$$

This contradicts the assumption of a dead-end state. □

The second issue is handled directly by Lemma 5.2. The lemma implies for the last agent $i$ and the remaining shares $\mathcal{M}$ of unique goods, $\tilde{v}_i(\mathcal{M}) \geq 1$, which means by receiving all of these shares, his utility is at least his MMS value in the instance with agents and goods from the start of Phase 2. By monotonicity of MMS (Appendix A of the full version), this completes the argument and yields the following result.

**Table 1: MMS Approximation factor obtained using Algorithm 1 for various values of $C$ and $k$**

| $k \backslash C$ | 0.0 | 0.1 | 0.2 | 0.3 | 0.5 | 0.7 | 0.8 | 0.9 | 0.99 |
|---|---|---|---|---|---|---|---|---|---|
| 2 | 1.0 | 0.9 | 0.8 | 0.7 | 0.5 | 0.3 | 0.2 | 0.1 | 0.01 |
| 3 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 0.6 | 0.4 | 0.2 | 0.02 |
| 4 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 0.9 | 0.6 | 0.3 | 0.03 |
| 5 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 0.8 | 0.4 | 0.04 |
| 6 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 0.5 | 0.05 |
| 8 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 0.7 | 0.07 |
| 10 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 0.9 | 0.09 |
| 15 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 0.14 |
| 20 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 0.19 |
| 25 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 0.24 |

Theorem 5.4. *Algorithm 1 computes in polynomial time the exact MMS $k$-sharing allocation, if $(1-C)(k-1) \geq 1$, or $(1-C)(k-1)$-MMS $k$-sharing allocation otherwise.*

Proof. All players addressed in Phase 1 receive at least their MMS (guaranteed by monotonicity; Appendix A of the full version).

Phase 2 operates on a normalized instance (guaranteed by scale-invariance; Appendix A of the full version) where $\text{MMS}_i^n(M)_i^{\text{norm}} \leq 1$ for all remaining agents. By Proposition 5.3, no dead-end occurs, therefore each agent except the last receives a bag $B$ satisfying $u_i^{\text{norm}}(B) \geq (1-C)(k-1)$ (Equation (6)), and the last agent receives all remaining goods, ensuring utility at least 1.

Thus, the algorithm guarantees an exact MMS for all agents when $(1-C)(k-1) \geq 1$, and $(1-C)(k-1)$-MMS otherwise. Clearly, the algorithm runs in polynomial time in the number of agents and goods. □

We note that for equal-share cost, where $C = (k-1)/k$, the algorithm produces $\frac{k-1}{k}$-MMS using this algorithm. However, when $C$ is a fixed constant, exact MMS can always be achieved for sufficiently large $k$—specifically, whenever $k \geq 1 + \frac{1}{1-C}$. Equivalently, for a fixed $k$, exact MMS is attainable as long as $C \leq \frac{k-2}{k-1}$. Table 1 illustrates this trade-off between the maximal cost of sharing and sharing degree.

## 6 DEFINING STRONGER MMS FOR K-SHARING

A natural generalization of the maximin share (MMS) to the $k$-sharing setting is to guarantee, for each agent $i$, the maximum over all $k$-sharing allocations of the minimum utility they could receive under some assignment of bundles. Formally, we introduce *Sharing Maximin Share (SMMS).*

*Definition 6.1 (SMMS).* A $k$-sharing allocation $A$ is said to satisfy *$k$-sharing maximin share (SMMS)* if for each agent $i \in N$: $u_i(A) \geq \text{SMMS}_i^{n,k}(M)$, where

$$\text{SMMS}_i^{n,k}(M) := \max_{B \in \mathcal{A}_k^n(M)} \min_{j \in [n]} u_i(B_{i\leftrightarrow j}), \tag{7}$$

$\mathcal{A}_k^n(M)$ is the set of all possible $k$-sharing allocations, and $B_{i\leftrightarrow j}$ denotes the modified allocation obtained from $B$ by swapping the bundles assigned to agents $i$ and $j$. Note that although in this work we assume the sharing-cost and utility of any good depends only on the number of agents sharing the good, the representation of modified allocation as $B_{i\leftrightarrow j}$ allows the sharing-cost and utility to be modeled as functions of the specific agents involved, rather than merely the size of the sharing group.

In Definition 6.1, $\min_{j\in[n]} u_i(B_{i\leftrightarrow j})$, captures the worst-case bundle in $B$ for agent $i$, taking into account that in the $k$-sharing model the utility depends on the entire allocation, not just the agent's own bundle. The utility $u_i(B_{i\leftrightarrow j})$ can be expressed as

$$u_i(B_{i\leftrightarrow j}) = \sum_{g\in B_j} \left[1 - c_g(|N_g(B)|)\right] \cdot v_i(g).$$

First, we investigate the existence of SMMS assuming some restricted cases: (1) only two agents and (2) identical valuations.

Proposition 6.2. *For any $k$-sharing instance $(N, M, k, v, c)$ with generous cost-sharing model and $|N| = k = 2$, an SMMS allocation always exists.*

Proposition 6.3. *For any $k$-sharing instance $(N, M, k, v, c)$ with goods-based cost-sharing instance and identical valuations $v_i(g) = v_j(g)$ for $i \neq j$ and for all $g \in M$, an SMMS allocation always exists.*

The proof of Propositions 6.2 and 6.3 are provided in the full version of the paper.

For three agents, computing the SMMS becomes significantly more challenging because full-sharing (each good shared by exactly $k$ agents) may not be the best option.

*Example 6.4 (Full-sharing is not optimal).* Consider an equal-share 2-sharing instance with three agents and two goods, $g_1$ and $g_2$. Agents have identical valuations: $v(g_1) = 1$ and $v(g_2) = 2$. An optimal allocation assigns $g_1$ to one agent, while $g_2$ is shared between the two remaining agents.

In contrast to the equal-share model, the cost-free model always favors full sharing as the optimal choice. We note that the SMMS value of an agent under the generous cost-sharing model is within a factor of $(1 - C)$ of their value in the cost-free model, where $C$ is the maximal cost per good. The details of this comparison are provided in Appendix D of the full version.

Surprisingly, we found that the existing MMS counterexamples [18, 25], in fact, admit SMMS allocations, which we explore next.

## Incompatibility of MMS and SMMS

In this section, we show that SMMS allocations can exist even in instances where MMS allocations do not. Conversely, we also provide a counterexample demonstrating that an SMMS allocation may fail to exist even when an MMS allocation does exist. Feige et al. [18] considers $n = 3$ agents and $m = 9$ goods in which no MMS allocation exists under 1-sharing. The instance is defined by agent-specific valuations over the goods, represented as $3 \times 3$ matrices:

$$V_1 = \begin{bmatrix} 1 & 16 & 23 \\ 26 & 4 & 10 \\ 12 & 19 & 9 \end{bmatrix}, V_2 = \begin{bmatrix} 1 & 16 & 22 \\ 26 & 4 & 9 \\ 13 & 20 & 9 \end{bmatrix}, V_3 = \begin{bmatrix} 1 & 15 & 23 \\ 25 & 4 & 10 \\ 13 & 20 & 9 \end{bmatrix}.$$

The goods are indexed from 1 to 9 in row-major order. Although no MMS allocation exists under 1-sharing, we observe that SMMS allocations do exist for both the *equal-share* and *cost-free* cost-sharing models even under 2-sharing.

*Equal-Share SMMS Allocation.* In the equal-share model, the sum of utilities of all bundles, given an allocation $A$ is equal to $v_i(M)$ for every agent $i$. As $v_i(M) = 120$, $\text{SMMS}_i^{n,k}(M) \leq 40$. At the same time, we have the following allocation with the corresponding utility, proving $\text{SMMS}_i^{n,k}(M) = 40$ for every agent $i$:

- Agent 1 receives goods $\{1, 2, 3\}$, utility: **40.00**,
- Agent 2 receives goods $\{4, 5, 8\}$, utility: **40.00**,
- Agent 3 receives goods $\{6, 7, 8, 9\}$, utility: **42.00**,

This is thus SMMS allocation. Notice that it was enough to share only one of the goods to achieve SMMS.

*Cost-Free SMMS Allocation.* As the utility of the least bundle in equal-share is bounded by 40, in cost-free sharing, the value of the least bundles is valued at most twice this price, i.e., $\text{SMMS}_i^{n,k}(M) \leq 80$ for every $i \in N$. The following allocation:

- Agent 1 receives goods $\{1, 2, 3, 4, 5, 6\}$, utility: **80.00**,
- Agent 2 receives goods $\{2, 3, 7, 8, 9\}$, utility: **80.00**,
- Agent 3 receives goods $\{4, 5, 6, 7, 8, 9\}$, utility: **81.00**,

is thus an SMMS allocation.

A similar behaviour can be observed for the example from Kurokawa et al. [25], where they constructed an instance with $n = 3$ and $m = 12$ goods, in which no 1-sharing MMS allocation exists. The valuation of agent $i \in N$ for a good $(k, \ell) \in M$ is given by

$$v_i(k, \ell) = 10^6 \cdot S_{k,\ell} + 10^3 \cdot T_{k,\ell} + E^i_{k,\ell},$$

where

$$S = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}, \quad T = \begin{bmatrix} 17 & 25 & 12 & 1 \\ 2 & 22 & 3 & 28 \\ 11 & 0 & 21 & 23 \end{bmatrix}$$

and

$$E^1 = \begin{bmatrix} 3 & -1 & -1 & -1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, E^2 = \begin{bmatrix} 3 & -1 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{bmatrix}, E^3 = \begin{bmatrix} 3 & 0 & -1 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}.$$

*Equal-Share SMMS Allocation.* For each agent, $v_i(M) = 12,165,000$ for every $i \in N$, thus $\text{SMMS}_i^{n,k}(M) \leq \frac{v_i(M)}{3} = 4,055,000$. The following allocation:

- Agent 1 - $\{1, 2, 3, 5, 10\}$, utility: **4,055,001**,
- Agent 2 - $\{4, 7, 8, 12\}$, utility: **4,055,000**,
- Agent 3 - $\{5, 6, 9, 10, 11\}$, utility: **4,055,000**.

is thus an SMMS allocation.

*Cost-Free SMMS Allocation.* Once again, assuming the SMMS value in the cost-free model is at most twice the SMMS value in equal share, we get $\text{SMMS}_i^{n,k}(M) \leq 8,110,000$. The following allocation is thus an SMMS allocation:

- Agent 1 - $\{1, 2, 3, 5, 7, 8, 10, 12\}$, utility: **8,110,001**,
- Agent 2 - $\{1, 3, 4, 6, 7, 9, 11, 12\}$, utility: **8,110,001**,
- Agent 3 - $\{2, 4, 5, 6, 8, 9, 10, 11\}$, utility: **8,110,000**.

Given these examples, it may seem that achieving SMMS is easier than achieving MMS; however, we provide a counterexample for SMMS by using the structure of the instance proposed by Kurokawa et al. [25], where SMMS does not exist, even though MMS does.

Theorem 6.5. *There exists an instance of cost-free* 2*-sharing problem with* $n = 3$ *and* $m = 12$*, which does not admit an SMMS allocation, but admits MMS allocation.*

Proof. Let $M = \{(k, \ell) \mid k \in [3], \ell \in [4]\}$. The valuation of agent $i \in N$ for a good $(k, \ell) \in M$ is given by

$$v_i(k, \ell) = 10^7 \cdot S_{k,\ell} + 10^3 \cdot T_{k,\ell} + E^i_{k,\ell},$$

$$S = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}, \quad T = \begin{bmatrix} 17 & 25 & 12 & 1 \\ 2 & 22 & 3 & 28 \\ 11 & 0 & 21 & 23 \end{bmatrix},$$

$$E^1 = \left[\begin{smallmatrix} -2 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 1 & 0 \end{smallmatrix}\right], E^2 = \left[\begin{smallmatrix} 0 & 0 & -2 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{smallmatrix}\right], E^3 = \left[\begin{smallmatrix} -1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{smallmatrix}\right].$$

The choice of matrix $S$ enforces that for each agent $i$, his $\mathrm{SMMS}_i^{3,2}(M)$ is achieved for an allocation which assigns each good to exactly two agents, i.e. each agent receives 8 goods. Under such 2-sharing allocations, $\mathrm{SMMS}_i^{3,2}(M) \sim 8 \cdot 10^7$ for every agent, which cannot be achieved for allocations where any share contains less than 8 goods (in such case, such agent receives strictly less than $8 \cdot 10^7$).

Matrix $T$ was chosen in such a way that if we label it using three different types of labels (numeric: 1, 2, 3, Greek: $\alpha, \beta, \gamma$, and symbols: $+, -, *$), entries corresponding to one specific label sum to exactly 55:

$$T = \begin{bmatrix} {}^{\alpha}17^{1}_{+} & {}^{\alpha}25^{1}_{-} & {}^{\beta}12^{1}_{+} & {}^{\gamma}1^{1}_{*} \\ {}^{\alpha}2^{2}_{-} & {}^{\beta}22^{2}_{*} & {}^{\gamma}3^{2}_{+} & {}^{\gamma}28^{2}_{-} \\ {}^{\alpha}11^{3}_{+} & {}^{\beta}0^{3}_{-} & {}^{\beta}21^{3}_{*} & {}^{\gamma}23^{3}_{+} \end{bmatrix}.$$

Since the total sum of entries in matrix $T$ is $3 \times 55 = 165$, any other partitioning of the goods into three bundles of four must include one whose sum in $T$ is strictly less than 55.

Now, consider an allocation that divides $M$ into three bundles, each containing 8 goods. Each bundle is constructed by selecting goods associated with two labels of the same type—for instance, one bundle includes goods labeled with $\alpha$ and $\beta$, another with $\beta, \gamma$, and the third with $\alpha, \gamma$. In this allocation, the sum of entries in $T$ for each bundle is exactly 110. By the argument above, any allocation that does not follow this label-based structure yields at least one bundle for which the sum of entries in $T$ is strictly less than 110.

This means, if there is an SMMS allocation, it must be one of those described above; assigning pairs of labels to agents

1. according to 1,2,3,
2. according to $\alpha, \beta, \gamma$,
3. according to $+, -, *$.

Each of the three ways guarantees $\mathrm{SMMS}_i^{3,2}(M)$ to one of the agents $i$, which is the same for all of the agents and equal to 8, 110, 000. However, none of the three ways assign $\mathrm{SMMS}_i^{3,2}(M)$ to all of the agents at the same time. This is enforced by the structure of matrices $E^1$, $E^2$, and $E^3$.

In division according to 1,2,3, both agent 2 and agent 3 need to be assigned elements with labels 2 and 3, otherwise their value is 8, 109, 999. Similarly, in division according to $\alpha, \beta, \gamma$, both agents 1 and 3 need to be assigned elements with $\beta$ and $\gamma$, or, once again, their value is less than their SMMS value. Finally, in division according to $+, -, *$, both agents 1 and 2 need to be assigned elements with $-$ and $*$.

Finally, assigning each agent their corresponding row (i.e., the first row to the first agent and so on) yields a valid 1-sharing MMS allocation. □

Despite the lack of general existence, we show that an $\alpha(1 - C)$-SMMS allocation of a goods-based cost-sharing model with $C$ being the maximal cost can be computed in polynomial time, via a reduction to $\alpha$-CMMS.

## 7 APPROXIMATE SMMS

We provide a connection of SMMS with the cardinality-constrained maximin share (CMMS). An instance of the *fair allocation problem under cardinality constraints* [10, 22] is denoted as $(N, M, v, b)$ where $N$ is the set of agents, $M = \{C_1, \ldots, C_\ell\}$ is a set of goods partitioned into sets $C_i$ according to $\ell$ types and $b = (b_1, \ldots, b_\ell)$ is the budget profile of types. In this setting, only a subset of allocations $\mathcal{F} \subseteq \mathcal{A}_1^n(M)$ is feasible; under feasible allocation, no player can receive more than $b_i$ goods of type $i$. Formally, $A \in \mathcal{F}$ satisfies for every $i \in [n]$, $j \in [\ell]$ that $|A_i \cap C_j| \leq b_j$. Here, the *cardinality-constrained maximin share* (CMMS) value is

$$\mathrm{CMMS}_i^n(M) := \max_{A \in \mathcal{F}} \min_{j \in [n]} v_i(A_j). \tag{8}$$

In any goods-based cost-sharing setting with $c_g(|N_g(A)|) \in [0, 1]$, fully-shared SMMS allocations correspond to allocations of CMMS of a special case of the fair allocation model under cardinality constraints, as stated in Proposition 7.1 (see Appendix E, full version).

Proposition 7.1. *For every instance* $I = (N, M, u)$ *of the goods-based* $k$*-sharing problem, there exists an instance* $\tilde{I} = (N, \tilde{M}, \tilde{v}, b)$ *of a fair allocation under cardinality constraints such that fully-shared allocations* $A \in \mathcal{A}_k^n$ *correspond bijectively to feasible allocations* $\tilde{A} \in \mathcal{F}$*, and the bijection preserves agents' utilities:* $\forall i \in N, u_i(A) = \tilde{v}_i(\tilde{A}_i)$.

If we define the *full-sharing maximin share value* $\overline{\mathrm{SMMS}}_i$ by restricting the SMMS definition to full-$k$-sharing allocations, then it follows from Proposition 7.1 that $\mathrm{CMMS}_i = \overline{\mathrm{SMMS}}_i$ for every $i \in N$ and also any $\alpha$-CMMS allocation $\tilde{A}$ corresponds to an $\alpha$-$\overline{\mathrm{SMMS}}$ allocation $A$. As $\overline{\mathrm{SMMS}}_i$ is defined over less allocations than $\mathrm{SMMS}_i$, we have $\mathrm{SMMS}_i \geq \overline{\mathrm{SMMS}}_i$. More importantly, we can bound $\mathrm{SMMS}_i$ and $\mathrm{MMS}_i$ by $\overline{\mathrm{SMMS}}_i$ from above, allowing us to construct approximations of these by $\alpha$-CMMS allocations.

Lemma 7.2. *For every goods-based* $k$*-sharing model with maximal possible sharing cost* $C$*, it holds*

1. $\overline{\mathrm{SMMS}}_i \geq (1 - C) \cdot \mathrm{SMMS}_i$,
2. $\overline{\mathrm{SMMS}}_i \geq k \cdot (1 - C) \cdot \mathrm{MMS}_i$.

Proof. (1.) Let a $k$-sharing allocation $A$ be an $\mathrm{SMMS}_i$ maximizer , and let $B$ be a fully-shared allocation obtained by taking $A$ and assigning the remaining shares of each good to arbitrary agents. Consequently, $A_j \subseteq B_j$ for every $j \in N$. Since $A$ is $\mathrm{SMMS}_i$ maximizer, $\mathrm{SMMS}_i$ is equal to

$$\begin{aligned} \min_{j \in N} u_i(A_{i \leftrightarrow j}) &= \min_{j \in N} \sum_{g \in A_j} \left(1 - c_g\left(|N_g(A)|\right)\right) \cdot v_i(g) \\ &\leq \min_{j \in N} \sum_{g \in A_j} v_i(g) \end{aligned} \tag{9}$$

where the last inequality follows from the fact that removing costs increases utility. Further, since $B$ is a fully-shared allocation, we have that $\overline{\text{SMMS}}_i$ is at least

$$\begin{aligned}\min_{j\in N} u_i(B_{i\leftrightarrow j}) &= \min_{j\in N} \sum_{g\in B_j} (1 - c_g(|N_g(B)|)) \cdot v_i(g) \\ &\geq (1 - C) \cdot \min_{j\in N} \sum_{g\in B_j} v_i(g) \qquad (10)\end{aligned}$$

where the final inequality uses that $C$ is the maximum cost. Finally, since $A_j \subseteq B_j$, we have

$$(1 - C) \cdot \min_{j\in N} \sum_{g\in B_j} v_i(g) \geq (1 - C) \cdot \min_{j\in N} \sum_{g\in A_j} v_i(g)$$

which in combination with (9) and (10) concludes the proof.

(2.) Let a 1-sharing allocation $A \in \mathcal{A}_1^n$ be an $\text{MMS}_i$ maximizer. We construct a different fully-shared allocation $B$ by giving each agent their own bundle from $A$ and the next $k - 1$ consecutive bundles, looping around to the beginning if we run out. Without the costs, each agent $i$ values each bundle in $B$ more than $k$-times the least valuable bundle in $A$, i.e., $\forall j \in N$, we have

$$v_i(B_j) \geq k \cdot \min_{\ell\in N} v_i(A_\ell) \geq k \cdot \text{MMS}_i\,. \qquad (11)$$

Note that (10) still holds for this allocation $B$, since the inequality relied only on $B$ being full-sharing. Combining (10) with (11) therefore gives us:

$$\overline{\text{SMMS}}_i \geq (1 - C) \cdot \min_{j\in N} v_i(B_j) \geq k \cdot (1 - C) \cdot \text{MMS}_i\,.$$

□

By combining the equivalence between CMMS and $\overline{\text{SMMS}}$ with the upper bounds from Lemma 7.2, we obtain approximation results for both MMS and SMMS.

Proposition 7.3. *Let $\alpha \in (0, 1]$, and suppose that for every fair division instance with cardinality constraints, there exists an $\alpha$-*CMMS *allocation. Then, for any goods-based $k$-sharing model with maximum cost $C$, there exist:*

- *an $\alpha \cdot (1 - C)$-*SMMS *allocation, and*
- *an $\alpha \cdot k \cdot (1 - C)$-*MMS *allocation.*

The best known $\alpha = 1/2$ [22] gives a $\frac{1}{2}$-SMMS guarantee in the cost-free model and $\frac{1}{2k}$-SMMS under equal-share. Also, while a larger $k$ leads to a worse approximation factor, the SMMS value may increase due to the greater flexibility in sharing.

Regarding MMS approximations, we note that $\alpha = 1/2$ does not improve upon the bag-filling algorithm (Proposition 5.4), which guarantees $(k-1)(1-C)$-MMS. However, for $\alpha > \frac{k-1}{k}$, the approach via $\alpha$-CMMS yields better guarantees—meaning that improvements occur even for $\alpha = 1/2 + \varepsilon$ with arbitrarily small $\varepsilon > 0$.

## 8 CONCLUSION

In this paper, we studied how allowing limited sharing of indivisible goods affects fairness, specifically focusing on Maximin Share (MMS) fairness. By permitting goods to be shared among up to $k$ agents, we found that previously unattainable fairness guarantees become achievable. We introduced Sharing Maximin Share (SMMS) fairness and for both notions explored when fair solutions exist, revealing important trade-offs between fairness, sharing costs, and the maximum number of agents allowed to share each good.

Our findings show that allowing bounded cost-sensitive sharing greatly improves the practicality of fair allocations, provided that the costs of sharing are not excessive, making our proposed mechanisms suitable for scenarios where sharing is possible. Table 1 could be viewed as a baseline for future work refining the boundary where sharing can or cannot lead to improved MMS guarantees depending on the parameters $k$ and $C$. A key open question is determining the lower bound on the number of agents $k$ needed to achieve fairness under *generous cost-sharing* model.

The theoretical results established in this paper open up new research directions for fair resource allocation when a limited sharing of resources are allowed in multi-agent environments. It opens the door to exploring various cost-sharing models that reflect structured, real-world constraints. One promising direction is to consider models where the cost of sharing is determined not by the goods themselves, but by the group of other agents they are sharing with—for example, agents may differ in how effectively they can make use of shared resources, with some incurring lower costs when participating in sharing, and costs being divided equally when similarly capable agents share the same good.

## ACKNOWLEDGMENTS

This work was carried out while H.S. and M.C. were participants in the 2025 DIMACS REU program at Rutgers University, supported by the NSF grant CCF-2447342. H.S. and M.C. were partly supported by the RSJ Foundation and by the Department of Applied Mathematics and the Computer Science Institute of Charles University. H.S. was supported by the Horizon Europe Programme under Grant Agreement No. 101183743 (AGATE).

## REFERENCES

[1] Hannaneh Akrami, Alon Eden, Michal Feldman, Amos Fiat, and Yoav Gal-Tzur. 2025. Fair Division via Resource Augmentation. arXiv:2502.09377 [cs.GT]

[2] Hannaneh Akrami and Jugal Garg. 2024. Breaking the 3/4 barrier for approximate maximin share. In *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 74–91.

[3] Martin Aleksandrov and Toby Walsh. 2020. Online fair division: A survey. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 13557–13562.

[4] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. 2023. Fair Division of Indivisible Goods: Recent Progress and Open Questions. *Artificial Intelligence* 319 (2023). https://doi.org/10.1016/j.artint.2023.103965

[5] Haris Aziz, Bo Li, Hervé Moulin, and Xiaowei Wu. 2022. Algorithmic Fair Allocation of Indivisible Items: A Survey and New Questions. *ACM SIGecom Exchanges* 20, 1 (July 2022), 24–40. https://doi.org/10.1145/3614593.3614597

[6] Siddharth Barman and Sanath Kumar Krishnamurthy. 2020. Approximation Algorithms for Maximin Fair Division. *ACM Transactions on Economics and Computation* 8, 1 (2020), 5:1–5:28. https://doi.org/10.1145/3381525

[7] Siddharth Barman, Satyanand Rammohan, and Aditi Sethia. 2025. Exact Maximin Share Fairness via Adjusted Supply. arXiv:2502.03789 [cs.GT]

[8] Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. 2022. Almost Full EFX Exists for Four Agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 4826–4833. https://doi.org/10.1609/aaai.v36i5.20410

[9] Samuel Bismuth, Ivan Bliznets, and Erel Segal-Halevi. 2024. Fair Division with Bounded Sharing: Binary and Non-degenerate Valuations. In *Algorithmic Game Theory: 17th International Symposium, SAGT 2024, Amsterdam, The Netherlands, September 3–6, 2024, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 89–107. https://doi.org/10.1007/978-3-031-71033-9_6

[10] Arpita Biswas and Siddharth Barman. 2018. Fair Division Under Cardinality Constraints. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)*. International Joint Conferences on Artificial Intelligence Organization, Stockholm, Sweden, 91–97. https://doi.org/10.24963/ijcai.2018/13

[11] Arpita Biswas, Yiduo Ke, Samir Khuller, and Quanquan C Liu. 2023. An Algorithmic Approach to Address Course Enrollment Challenges. In *4th Symposium on Foundations of Responsible Computing*.

[12] Arpita Biswas, Justin Payan, Rik Sengupta, and Vignesh Viswanathan. 2023. The Theory of Fair Allocation Under Structured Set Constraints. In *Ethics in Artificial Intelligence: Bias, Fairness and Beyond*, Animesh Mukherjee, Juhi Kulshrestha, Abhijnan Chakraborty, and Srijan Kumar (Eds.). Studies in Computational Intelligence, Vol. 1123. Springer, 115–129. https://doi.org/10.1007/978-981-99-7184-8_7

[13] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103. https://doi.org/10.1086/664613

[14] Ioannis Caragiannis, Nick Gravin, and Xin Huang. 2019. Envy-Freeness Up to Any Item with High Nash Welfare: The Virtue of Donating Items. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC '19)*. ACM, Phoenix, AZ, USA, 527–545. https://doi.org/10.1145/3328526.3329574

[15] Hsiu-Chuan Chang, Bissan Ghaddar, and Jatin Nathwani. 2022. Shared community energy storage allocation and optimization. *Applied Energy* 318 (2022), 119160.

[16] Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta, and Pranabendu Misra. 2021. Improving EFX Guarantees through Rainbow Cycle Number. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC '21)*. 310–311. https://doi.org/10.1145/3465456.3467605

[17] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. 2021. A Little Charity Guarantees Almost Envy-Freeness. *SIAM J. Comput.* 50, 4 (2021), 1244–1275. https://doi.org/10.1137/20M1359134

[18] Uriel Feige, Ariel Sapir, and Laliv Tauber. 2021. A Tight Negative Example for MMS Fair Allocations. In *Proceedings of the 17th International Conference on Web and Internet Economics (WINE)* (Potsdam, Germany) *(Lecture Notes in Computer Science, Vol. 13112)*. Springer, 355–372.

[19] Jugal Garg, Peter McGlaughlin, and Setareh Taki. 2019. Approximating maximin share allocations. In *2nd Symposium on Simplicity in Algorithms (SOSA 2019)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 20–1.

[20] Jugal Garg and Setareh Taki. 2021. An Improved Approximation Algorithm for Maximin Shares. *Artificial Intelligence* 300 (2021), 103547. https://doi.org/10.1016/j.artint.2021.103547

[21] Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2021. Fair Allocation of Indivisible Goods: Improvement. *Mathematics of Operations Research* 46, 3 (2021), 1038–1053. https://doi.org/10.1287/moor.2020.1096

[22] Halvard Hummel and Magnus Li Hetland. 2022. Maximin Shares Under Cardinality Constraints. *Multi-Agent Systems* 13210 (2022), 188–206. https://doi.org/10.1007/978-3-031-20614-6_11

[23] Pierre Jacquet, Thomas Ledoux, and Romain Rouvoy. 2024. Sweetspotvm: Oversubscribing CPU without sacrificing VM performance. In *2024 IEEE 24th International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*. IEEE, 148–157.

[24] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2016. When Can the Maximin Share Guarantee Be Guaranteed?. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 523–529.

[25] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2018. Fair Enough: Guaranteeing Approximate Maximin Shares. *J. ACM* 65, 2 (2018), 8:1–8:27. https://doi.org/10.1145/3140756

[26] Ariel Procaccia, Isaac Robinson, and Jamie Tucker-Foltz. 2024. School redistricting: Wiping unfairness off the map. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2704–2724.

[27] Fedor Sandomirskiy and Erel Segal-Halevi. 2022. Efficient Fair Division with Minimal Sharing. *Oper. Res.* 70, 3 (May 2022), 1762–1782. https://doi.org/10.1287/opre.2022.2279

[28] K. A. Santhini, Kamesh Munagala, Meghana Nasre, and Govind S. Sankar. 2024. Group Fairness and Multi-criteria Optimization in School Assignment. arXiv:2403.15623 [cs.DS] https://arxiv.org/abs/2403.15623 To appear in FORC 2025.

[29] Warut Suksompong. 2021. Constraints in Fair Division. *ACM SIGecom Exchanges* 19, 2 (November 2021), 46–61. https://dl.acm.org/doi/10.1145/3508352.3508357

[30] Toby Walsh. 2020. Fair Division: The Computer Scientist's Perspective. *International Joint Conference on Artificial Intelligence* (2020), 4966–4972.