Fairness and Stability for Shared Resource Allocation Problems

Jiazhu Fang[1], **Qizhi Fang**[1], **Minming Li**[2*], **Wenjing Liu**[1]

[1]School of Mathematical Sciences, Ocean University of China, Qingdao, Shandong, China.
[2]Department of Computer Science, City University of Hong Kong, Kowloon Tong, Hong Kong, China.
fjz@stu.ouc.edu.cn, qfang@ouc.edu.cn, minming.li@cityu.edu.hk, liuwj@ouc.edu.cn

Abstract

This paper investigates the problem of shared resource allocation, where a set of agents must be assigned to heterogeneous resources, with each agent allocated exactly one resource and each resource potentially shared by multiple agents. An agent's utility for a given resource is jointly determined by the resource's type and the number of agents sharing it. We focus on two fundamental classes of monotone valuations: monotone nondecreasing and monotone nonincreasing, where an agent's utility respectively increases or decreases with the number of agents sharing the resource. Within this shared resource framework, we examine classical notions of fairness and stability, including maximin-share fairness, envy-freeness, Nash stability, and two epistemic relaxations—epistemic envy-freeness and epistemic Nash stability—as well as swap stability. We propose formal definitions adapted to this setting and systematically analyze the relationships among these concepts. The primary contributions of this work consist of establishing existence and computational complexity results for each notion under both monotonicity assumptions and developing polynomial-time algorithms in cases where fair or stable allocations are guaranteed to exist.

1 Introduction

Fair allocation of a set of indivisible resources among self-interested agents has long been a central concern, and remains a core problem in algorithmic game theory. Over the past decades, economists, mathematicians, and computer scientists have systematically studied this problem with the goal of providing *feasible fairness* guarantees (Steinhaus 1948; Thomson 1983; Moulin 2004; Aziz et al. 2015). In the classical fair division framework, each resource can be allocated to at most one agent, and strong fairness notions such as *envy-freeness* (Foley 1966) or *proportionality* (Steinhaus 1948) are often unattainable. This has spurred a rich literature on relaxed fairness concepts: for instance, *maximin-share fairness* (Budish 2011; Kurokawa, Procaccia, and Wang 2018), *envy-freeness up to one item* (Lipton et al. 2004), and *envy-freeness up to any item* (Caragiannis et al. 2019). However, these traditional frameworks neglect scenarios where resources can serve multiple agents simultaneously.

Many modern applications involve resources that, while indivisible, are inherently *shareable*: a single resource can serve multiple agents simultaneously, and an agent's utility depends on the number of co-users. For example, in ridesharing among travelers, a vehicle can carry multiple passengers, and a passenger's utility for a given vehicle may vary with both the type of vehicle and the number of fellow passengers. Some individuals enjoy social interaction and derive higher utility from a larger group of co-riders, while others dislike crowding and experience decreasing utility as group size grows. As another example, in cloud computing or data center scheduling, virtual machines or computational tasks are assigned to servers that can host multiple tasks concurrently; a server under heavy load may degrade each task's performance, whereas multiple tasks sharing a server's cache can improve execution efficiency.

Despite the prevalence of such scenarios, to the best of our knowledge, no general framework for the fair allocation of shared resources has been systematically developed. A recent study by (Gan, Li, and Li 2023) considers a related college dormitory assignment problem, where each student is assigned to a room with multiple roommates. The authors define separate valuation functions for room quality and roommate group, demonstrating that a strictly envy-free assignment (where no student prefers another's room and roommate combination) is generally impossible. By introducing a relaxed concept of *Pareto envy-freeness* (where students compare either room quality or roommate group, but not both simultaneously), they show that an envy-free assignment exists when room capacity is at most two. However, their results focus on specific envy notions in this constrained setting. In contrast, our objective is to develop a comprehensive fairness and stability framework for shared resources under *general* valuations, without imposing restrictions on room capacity or separating preferences for room quality and roommates.

In this paper, we formally introduce and study the *shared resource allocation* problem, where the allocation perspective is reversed: instead of assigning resources to agents, we assign each agent to a resource, with each resource potentially shared by multiple agents. An agent's utility for a given resource depends on both the resource's type and the number of agents sharing it. We focus on two natural classes of valuation functions: *monotone nondecreasing* valuations, where utility increases with the number of co-users, and *monotone*

*Corresponding Authors.

nonincreasing valuations, where utility decreases with more co-users. Notably, our model differs from those involving divisible resources, where agents can take fractional shares of each resource. In contrast, resources in our model are indivisible. Each agent either uses a whole resource (possibly shared with others) or not at all, which leads to discrete externalities rather than linear utility scaling. Furthermore, while our problem resembles an "anonymous" roommate assignment with arbitrary room sizes, prior work either assumes a maximum of two agents per resource or imposes separable preferences. Our framework generalizes these models by allowing arbitrary monotone valuations and group sizes.

1.1 Our Contributions

We introduce the shared resource allocation problem and extend the classical fairness and stability concepts to this model, analyzing the existence and computational complexity of solutions under two classes of monotone valuations. Our main results are summarized in Table 1.

Maximin Share. In fair division, an agent's *maximin share* (MMS) value is the maximum utility she can guarantee for herself by partitioning the items into n bundles and receiving the least valued bundle. An allocation is MMS if every agent receives utility at least equal to her MMS value. This definition does not directly apply to shared resources, since we partition agents among resources rather than partitioning resources themselves. We therefore generalize the MMS concept to our setting: an agent's MMS value is the maximum utility she can guarantee by proposing an *allocation vector* (specifying how many agents would share each resource) and then receiving the resource with minimum utility under that vector. Under monotone nonincreasing valuations, we show that each agent's MMS value and an MMS allocation can be computed in polynomial time. This stands in stark contrast to classical fair division with indivisible goods, where determining an agent's MMS or finding an MMS allocation is NP-hard. In contrast, under monotone nondecreasing valuations, an agent's MMS value remains polynomial-time computable, but an MMS allocation may not exist, and deciding the existence of an MMS allocation is NP-complete.

Envy-Freeness and Nash Stability. We extend *envy-freeness* (EF) to the shared setting by requiring that no agent prefers another agent's assigned resource (considering the number of sharers) and, additionally, that no agent would prefer an *empty* resource (i.e., one assigned to no agent) over their own. The empty resource condition prevents the degenerate allocation where all agents share a single resource, which would trivially eliminate envy but is highly inefficient. We also adapt *Nash stability* (NS) from hedonic coalition formation: an allocation is NS if no agent can unilaterally deviate to another resource and become strictly better off. Interestingly, EF and NS imply each other under opposite monotonicity regimes. When valuations are monotone nonincreasing, any EF allocation is also NS; conversely, under monotone nondecreasing valuations, any NS allocation is also EF. We establish existence and complexity results for these concepts in our model. For monotone nonincreasing valuations, an NS allocation always exists and can be found in polynomial time, whereas an EF allocation may not exist. For monotone nondecreasing valuations, an NS allocation may not exist when there are more than two resources, and deciding its existence is NP-complete. On the positive side, if there are only two resources or if all agents have identical valuations, then an NS (and hence EF) allocation is guaranteed to exist. The existence of an EF allocation for general instances with more than two resources under monotone nondecreasing valuations remains an open question.

Epistemic Relaxations: Epistemic Envy-Freeness and Epistemic Nash Stability. Given the negative results above, we explore their *epistemic* relaxations, as introduced by (Aziz et al. 2018). An allocation is *epistemically envy-free* (EEF) if any envy perceived by an agent can be eliminated by hypothetically reassigning agents sharing other resources among those resources, while maintaining the sharing configuration of their own resource. Similarly, an allocation is *epistemically Nash stable* (ENS) if an agent's deviation incentive can be removed by a suitable reassignment of agents sharing other resources. These relaxations capture scenarios where agents have limited information about the sharing of other resources. Clearly, any EF (resp. NS) allocation is also EEF (resp. ENS), and the relationship between EEF and ENS mirrors that between EF and NS. We find that epistemic relaxations often restore existence. Under monotone nondecreasing valuations, an EEF and an ENS allocation are guaranteed to exist and can be computed in polynomial time. Under monotone nonincreasing valuations, ENS allocations always exist and are efficiently computable, while EEF allocations may not exist, similar to exact EF.

Swap Stability. Finally, we investigate *swap stability* (SS), a concept from matching theory requiring that no pair of agents can mutually improve their utilities by exchanging their assigned resources. We prove that, even without any monotonicity assumptions, an SS allocation always exists in our model and can be found in polynomial time. In fact, any allocation maximizing social welfare (the sum of all agents' utilities) must be swap-stable. However, computing a welfare-maximizing allocation is NP-hard, even under monotone nondecreasing valuations. Due to space constraints, we present these results in the Appendix.

1.2 Related Work

Our study relates to several areas of research, with key distinctions from each. We begin with the literature on fair division, which can be categorized by resource type. The first category concerns *divisible* resources, also known as cake-cutting, where heterogeneous resources can be infinitely divided (Steinhaus 1948; Aziz and Mackenzie 2016; Dubins and Spanier 1961). The second category considers *indivisible* resources, in which each item must be allocated whole to a single agent (Budish 2011; Lipton et al. 2004). Two central fairness notions in these settings are *envy-freeness* (EF) (Foley 1966), which requires that no agent prefers another agent's allocation, and *maximin-share fairness* (MMS) (Budish 2011; Kurokawa, Procaccia, and Wang 2018), which guarantees that each agent receives at least the maximum value she could guarantee by partitioning the resources into n

Concept	Monotone Nondecreasing		Monotone Nonincreasing	
	Existence	Complexity	Existence	Complexity
MMS	× [Ex.1]	NP–c [Thm.3]	✓[Thm.2]	P [Thm.2]
NS	× [Ex.2]	NP–c [Thm.4]	✓[Thm.6]	P [Thm.6]
EF	Open	NP–c* [Thm.5]	× [Ex.3]	NP–c* [Thm.5]
EEF	✓[Cor.2]	P [Cor.2]	× [Cor.4]	Open
ENS	✓[Thm.7]	P [Thm.7]	✓[Cor.3]	P [Cor.3]
SS	✓[Appendix]	P [Appendix]	✓[Appendix]	P [Appendix]

Table 1: Our Main Results. Abbreviations: EF = envy-freeness, MMS = maximin share fairness, NS = Nash stability, SS = swap stability, EEF = epistemic envy-freeness, ENS = epistemic Nash stability, NP–c = NP-complete, NP–c* = NP-complete result under valuations with heterogeneous monotonicity, P = polynomial time computable, × = may not exist, ✓= always exists.

bundles and taking the least valuable bundle. For divisible resources, MMS coincides with proportionality, i.e., each agent obtains at least her proportional share of the total value (Steinhaus 1948), and both EF and MMS allocations are known to always exist (Aziz and Mackenzie 2016). In contrast, for indivisible resources, EF and MMS may not exist (Aziz et al. 2015; Bhaskar, Sricharan, and Vaish 2021; Aziz et al. 2017; Kurokawa, Procaccia, and Wang 2018), motivating relaxed notions such as envy-freeness up to one item (Lipton et al. 2004), envy-freeness up to any item (Caragiannis et al. 2019), and epistemic envy-freeness (Aziz et al. 2018). Similarly, approximate MMS allocations have been extensively studied (Huang and Lu 2021; Amanatidis et al. 2017; Garg and Taki 2020; Feige, Sapir, and Tauber 2022). Note that while divisible resources can be shared via fractional allocations, with utility derived from those fractions, in our model resources remain indivisible yet shareable, and utility depends on the number of sharers. Another line of research on shareable resources is *group fairness* in fair division, where agents are prepartitioned into groups that share the resources allocated to their group (Manurangsi and Suksompong 2017; Segal-Halevi and Nitzan 2019). Fairness notions studied include group envy-freeness (Conitzer et al. 2019) and group maximin-share fairness (Barman et al. 2018; Suksompong 2018). These models assume predetermined group membership and therefore cannot capture the more complex scenario in which groups form endogenously through the allocation and each agent's utility depends on the resulting group size.

Furthermore, our work is connected to coalition formation games and congestion games. *Hedonic games* model coalition formation where each agent's utility depends solely on the set of coalition members (Aziz et al. 2019; Bilò et al. 2018; Bogomolnaia and Jackson 2002; Drèze and Greenberg 1980). Standard solution concepts in hedonic games include *core stability* and *Nash stability*. In particular, the class of *anonymous* hedonic games assumes that an agent's utility depends only on coalition size, not on the identities of the other members (Jang, Shin, and Tsourdos 2018). These models focus on the stability of coalitions rather than fairness across agents, and do not involve distinct resource types. One notable exception is the *group activity selection* games (Darmann et al. 2012; Darmann 2015; Igarashi, Peters, and Elkind 2017), where agents choose activities, and the utility depends on both the activity type and the group size. However, that model allows agents to opt out and emphasizes individual rationality, whereas in our setting every agent must be assigned to some resource.

Additionally, *Congestion games* study scenarios where multiple agents choose resources and incur costs that increase or decrease with the number of users (Holzman and Law-Yone 1997; Milchtaich 1996; Rosenthal 1973). It is well known that every congestion game admits a pure Nash equilibrium (Rosenthal 1973). In fact, congestion games correspond to a special case of our shared resource model when all agents' valuations are identical and determined solely by the resource's cost function. However, analyses of congestion games focus on selfish cost-minimization and do not consider fairness criteria such as EF or MMS.

Finally, matching theory offers further parallels, including the *college admissions* problem and the *stable roommates* problem (Gale and Shapley 1962; Chung 2000). The college admissions problem seeks a many-to-one matching between students and schools based on mutual preferences, and research has primarily focused on stability and strategyproofness (Gale and Shapley 1962). In that model, an agent's preference does not depend on how many others share the same assignment, so there are no externalities from group size. The classic stable roommates problem initially assumes a set of $2n$ students to be paired into n identical rooms, considering only preferences over roommates. Subsequent work (Huzhang et al. 2017; Chan et al. 2016; Gan, Li, and Li 2023) extended this setting to heterogeneous rooms, and defined envy-based fairness notions or their relaxations for this context. However, these studies typically assume a room capacity of two, which is not applicable to our framework.

2 Preliminaries

2.1 Shared Resource Allocation Problem

We consider the shared resource allocation problem, defined as follows. Let N be a set of n agents and M be a set of m indivisible resources, each of which may be shared by multiple agents. The valuation of each agent $i \in N$ is determined by the resource to which she is assigned, and also influenced by the number of agents sharing that resource. Formally, agent i's valuation is given by a function $v_i : M \times [n] \to \mathbb{R}$, where $v_i(j, k)$ denotes the value agent i obtains when assigned to resource j, and exactly k agents in total share that resource.

We denote the *valuation profile* by $\mathbf{v} = (v_1, \ldots, v_n)$ and an *instance* by $I = (N, M, \mathbf{v})$ or simply by I.

An *allocation* is a partition $\mathbf{X} = (X_1, \ldots, X_m)$ of the agent set N into m (possibly empty) pairwise disjoint subsets such that $\bigcup_{j=1}^{m} X_j = N$. Here, X_j represents the set of agents assigned to (i.e., sharing) resource j, and each agent appears in exactly one subset. The corresponding *allocation vector* is the m-tuple $\mathbf{x} = (x_1, \ldots, x_m)$, where $x_j = |X_j|$, for each $j \in [m]$. For each agent $i \in N$, let $r_i(\mathbf{X})$ denote the resource to which i is assigned in $\mathbf{X}$. Note that each allocation $\mathbf{X}$ induces a unique allocation vector $\mathbf{x}$, although the converse does not necessarily hold.

In this work, we investigate the existence of allocations that satisfy certain fairness or stability criteria under particular classes of valuation functions. A natural and widely adopted assumption is that all agents' valuation functions share a common monotonicity property.

Definition 1. *A valuation function v_i is called* monotone nonincreasing *(resp.* monotone nondecreasing*) if for every resource $j \in M$ and all integers $1 \leq k_1 \leq k_2 \leq n$, we have $v_i(j, k_1) \geq v_i(j, k_2)$ (resp. $v_i(j, k_1) \leq v_i(j, k_2)$).*

2.2 Solution Concepts

We next introduce several desirable properties of allocations. Some of these properties extend classical fairness notions in fair division to our shared resource setting.

We begin with the notion of *maximin share* (MMS) — one of the most extensively studied threshold-fairness notions in fair division. An allocation satisfies MMS fairness if each agent i receives a utility at least equal to her maximin share value MMS_i. Intuitively, MMS_i represents the maximum utility agent i can guarantee for herself by partitioning the resources into n bundles (one per agent) and then receiving the worst of those bundles. However, in our shared model, agent i instead reports an allocation vector $\mathbf{x} = (x_1, \ldots, x_m)$, after which a scheduler will assign x_j agents to each resource j; agent i aims to maximize her worst-case utility under the allocation vector she chooses. Formally, for any instance I, let $\mathcal{F}(n, m)$ denote the set of all possible allocation vectors for n agents and m resources. Then the maximin share of agent i is defined as

$$\mathrm{MMS}_i(n, m) = \max_{\mathbf{x} \in \mathcal{F}(n,m)} \min_{j \in [m]: x_j > 0} v_i(j, x_j).$$

When n and m are clear from context, we abbreviate this value as MMS_i. An allocation vector $\mathbf{x} = (x_1, \ldots, x_m)$ is called an MMS_i *allocation vector for agent i* if $\min_{j \in [m]: x_j > 0} v_i(j, x_j) \geq \mathrm{MMS}_i$.

Definition 2. *An allocation $\mathbf{X} = (X_1, \ldots, X_m)$ is called an* MMS allocation *if for every agent $i \in N$,*

$$v_i(r_i(\mathbf{X}), |X_{r_i(\mathbf{X})}|) \geq \mathrm{MMS}_i.$$

In classical fair division, *envy-freeness* requires that no agent prefers another's bundle. In our model, the analogous requirement is that no agent prefers the resource assigned to another, accounting for the number of sharers. To avoid the degenerate allocation where all agents share a single resource, which would trivially satisfy envy-freeness but be highly inefficient, we additionally require that no agent envies an empty resource, i.e., one that is assigned to no agent.

Definition 3. *An allocation $\mathbf{X} = (X_1, \ldots, X_m)$ is called* envy-free (EF) *if for every agent $i \in N$, the following hold:*

1. $v_i(r_i(\mathbf{X}), |X_{r_i(\mathbf{X})}|) \geq v_i(j, 1)$, $\forall j \in M$ with $X_j = \emptyset$.
2. $v_i(r_i(\mathbf{X}), |X_{r_i(\mathbf{X})}|) \geq v_i(r_{i'}(\mathbf{X}), |X_{r_{i'}(\mathbf{X})}|)$, $\forall i' \in N$.

An allocation is *Nash-stable* if no agent can increase her utility by unilaterally switching to another resource.

Definition 4. *An allocation $\mathbf{X} = (X_1, \ldots, X_m)$ is called* Nash-stable (NS) *if $v_i(r_i(\mathbf{X}), |X_{r_i(\mathbf{X})}|) \geq v_i(j, |X_j| + 1)$ for every agent $i \in N$ and every resource $j \in M \setminus \{r_i(\mathbf{X})\}$.*

However, for certain classes of valuations both EF and NS may be too strong. We thus consider epistemic relaxations of these notions, following the literature.

As introduced by (Aziz et al. 2018), *epistemic envy-freeness* is a relaxation of envy-freeness that naturally applies to our shared setting. Intuitively, an allocation is epistemic envy-free if any envy an agent may feel can be eliminated by suitably "reshuffling" agents sharing other resources, without altering the sharing configuration of the resource assigned to that agent.

Definition 5. *An allocation $\mathbf{X} = (X_1, \ldots, X_m)$ is called* epistemic envy-free (EEF) *if for every agent $i \in N$ there exists another allocation $\mathbf{X}^{(i)} = (X_1^{(i)}, \ldots, X_m^{(i)})$ such that $X^{(i)}_{r_i(\mathbf{X})} = X_{r_i(\mathbf{X})}$, and for all $j \in M$,*

$$v_i(r_i(\mathbf{X}), |X_{r_i(\mathbf{X})}|) \geq \begin{cases} v_i(j, 1), & \text{if } X_j^{(i)} = \emptyset, \\ v_i(j, |X_j^{(i)}|), & \text{if } X_j^{(i)} \neq \emptyset. \end{cases}$$

Similarly, we define *epistemic Nash-stability*, which requires that each agent's incentive to deviate be eliminated through a suitable reshuffling.

Definition 6. *An allocation $\mathbf{X} = (X_1, \ldots, X_m)$ is called* epistemic Nash-stable (ENS) *if for every agent $i \in N$ there exists an allocation $\mathbf{X}^{(i)} = (X_1^{(i)}, \ldots, X_m^{(i)})$ such that $X^{(i)}_{r_i(\mathbf{X}^{(i)})} = X_{r_i(\mathbf{X})}$, and for all $j \in M \setminus \{r_i(\mathbf{X})\}$,*

$$v_i(r_i(\mathbf{X}), |X_{r_i(\mathbf{X})}|) \geq v_i(j, |X_j^{(i)}| + 1).$$

Note that EEF and ENS are defined by restricting the information available to agents. Under EEF and ENS, each agent knows only the sharing configuration of their assigned resource and is unaware of the configurations of other resources. In contrast, under EF and NS, agents have full knowledge of the entire resource allocation.

2.3 Connections Between Concepts

We now consider the relationships among the solution concepts under two classes of monotone valuations, as illustrated in Figure 1 and Figure 2. By definition, it is straightforward to verify that EF implies both EEF and SS, and NS implies ENS, for any valuation profile.

Proposition 1. *Every EF allocation is also EEF and SS; every NS allocation is also ENS.*

$$\begin{array}{ccccccc} \text{MMS} & \Rightarrow & \text{NS} & \Rightarrow & \text{EF} & \Rightarrow & \text{SS} \\ & & \Downarrow & & \Downarrow & & \\ & & \text{ENS} & \Rightarrow & \text{EEF} & & \end{array}$$

Figure 1: Implications between concepts when all valuations are monotone nondecreasing.

$$\begin{array}{ccccc} \text{NS} & \Leftarrow & \text{EF} & \Rightarrow & \text{SS} \\ \Downarrow & & \Downarrow & & \\ \text{ENS} & \Leftarrow & \text{EEF} & \Rightarrow & \text{MMS} \end{array}$$

Figure 2: Implications between concepts when all valuations are monotone nonincreasing.

From the definitions and monotonicity, the mutual implications between EF and NS under different monotonicity assumptions also hold. The relationship between EF and NS is analogous to that between EF and EFX (EF1), with the distinction that in the former case, the difference arises from the impact of a single item, whereas in our case, it is the impact of a single agent. Naturally, the relationship between EEF and ENS mirrors that between EF and NS.

Proposition 2. *If all valuations are monotone nonincreasing, then every (E)EF allocation is also (E)NS; conversely, if all valuations are monotone nondecreasing, then every (E)NS allocation is also (E)EF.*

Under monotone nondecreasing valuations, note that each agent i's MMS_i value is simply the maximum utility she can obtain, i.e. $\mathrm{MMS}_i = \max_{j \in M} v_i(j, n)$. Therefore, any MMS allocation guarantees that every agent achieves her maximum possible value, and thus an MMS allocation satisfies all of the above solution concepts.

Proposition 3. *If all valuations are monotone nondecreasing, then every MMS allocation is also NS.*

Under monotone nonincreasing valuations, it can be further shown that EEF implies MMS.

Proposition 4. *If all valuations are monotone nonincreasing, then every EEF allocation is also an MMS allocation.*

3 MMS Allocation

In classical fair division, computing an MMS allocation or even a single agent's MMS value is NP-hard under monotone additive valuations. In contrast, in the shared resource model, these computational barriers can be overcome. When valuations are monotone, every agent's MMS value can be computed in polynomial time. Furthermore, when valuations are monotone nonincreasing, an MMS allocation always exists and can be computed efficiently.

Monotone Nonincreasing Valuations. By the definition of maximin share, computing i's MMS_i value reduces to determining an appropriate MMS_i allocation vector for i. The main challenge is that we do not know in advance which resource will be the least-valued in i's MMS_i allocation vector.

To circumvent this, we consider each resource $j \in M$ under the assumption that j is the least-valued resource in i's desired allocation vector. For each such j, we compute an allocation vector that satisfies this assumption and maximizes i's utility from j. Specifically, we start with $\mathbf{x}^{(j)} = (x_1^{(j)}, \ldots, x_m^{(j)})$ initialized by setting $x_j^{(j)} = n$ and $x_{j'}^{(j)} = 0$ for every $j' \neq j$. Then, while there exists another resource $j' \neq j$ such that $v_i(j', x_{j'}^{(j)} + 1) \geq v_i(j, x_j^{(j)} - 1)$, we move one agent from j to j', updating $x_j^{(j)} \leftarrow x_j^{(j)} - 1$ and $x_{j'}^{(j)} \leftarrow x_{j'}^{(j)} + 1$. This iterative process maintains j as the least-valued among all nonempty resources. By Theorem 1, upon termination, $\mathbf{x}^{(j)}$ is an allocation vector in which j is indeed the least-valued resource, and among all such vectors it maximizes agent i's utility for j.

Performing the above procedure for every $j \in M$ yields m candidate allocation vectors. We then select the vector that maximizes the minimum utility (i.e., the utility of the assumed least-valued resource), as detailed in Algorithm 1.

Algorithm 1: Computing an MMS_i allocation vector for agent i with monotone nonincreasing valuation v_i

---

Input: An instance I with nonincreasing valuations.
Output: An MMS_i allocation vector $\mathbf{x}$ for i.

```
Initialize: For each j ∈ [m], set x^(j) with x_j^(j) = n and x_{j'}^(j) = 0 for all j' ≠ j.
for all j ∈ [m] do
    while there exists j' ≠ j such that v_i(j, x_j^(j) − 1) ≤ v_i(j', x_{j'}^(j) + 1) do
        Set x_j^(j) = x_j^(j) − 1 and x_{j'}^(j) = x_{j'}^(j) + 1.
    end while
end for
Let j* ∈ arg max_{j∈[m]} v_i(j, x_j^(j)).
return x^(j*) = (x_1^(j*), ..., x_m^(j*)).
```

Theorem 1. *When agents have monotone nonincreasing valuations, an MMS_i allocation vector for any agent $i \in N$ can be computed in polynomial time.*

Proof Sketch. For each resource $j \in M$, the algorithm constructs an allocation vector $\mathbf{x}^{(j)}$ by iteratively transferring agents while ensuring that j remains the least-valued nonempty resource for agent i. This invariant follows directly from the transfer condition in the inner loop. We can then argue, by contradiction, that $\mathbf{x}^{(j)}$ maximizes i's utility from j among all allocations where j is the least-valued nonempty resource. Otherwise, an additional improving transfer would exist, contradicting the loop's termination. Finally, selecting $\mathbf{x}^{(j^*)}$ with least-valued resource j^* that maximizes $v_i(j^*, x_{j^*}^{(j^*)})$ over all m vectors gives the desired result. □

Using Algorithm 1, we can compute an MMS_i allocation vector $\mathbf{x}^{(i)} = (x_1^{(i)}, \ldots, x_m^{(i)})$ for each agent $i \in N$. For any such vector, if agent i is assigned to resource j and the number of agents on j does not exceed $x_j^{(i)}$, then i's valuation will be no less than MMS_i. Algorithm 2 describes

Algorithm 2: Computing an MMS allocation for monotone nonincreasing valuations

Input: An instance I with nonincreasing valuations; for each $i \in N$, an MMS_i allocation vector $\mathbf{x}^{(i)} = (x_1^{(i)}, \ldots, x_m^{(i)})$.
Output: An MMS allocation $\mathbf{X} = (X_1, \ldots, X_m)$.

Initialize: $\mathbf{X} = (\emptyset, \ldots, \emptyset)$, $N' \leftarrow N$.
for all $j \in [m]$ **do**
  while there exists $i^* = \arg\max_{i \in N'} x_j^{(i)}$ such that $|X_j \cup \{i^*\}| \leq x_j^{(i^*)}$ **do**
    Set $X_j = X_j \cup \{i^*\}$ and $N' = N' \setminus \{i^*\}$.
  end while
end for
return $\mathbf{X} = (X_1, \ldots, X_m)$.

how to construct an MMS allocation from these n vectors. Specifically, for each resource j, we iteratively assign to j the remaining unassigned agent who "most prefers" j (i.e., the agent $i^* \in N'$ maximizing $x_j^{(i^*)}$ among those $i \in N'$ still unassigned), as long as assigning i^* to j does not exceed her bound $x_j^{(i^*)}$. This process continues until no further agent can be added to j without exceeding her bound, at which point we proceed to the next resource.

We illustrate the execution of Algorithm 2 with an example. Consider two resources and five agents whose MMS allocation vectors are $(3, 2)$, $(2, 3)$, $(1, 4)$, $(4, 1)$, and $(2, 3)$. Initially, $(X_1, X_2) = (\emptyset, \emptyset)$. For resource 1, agent 4 has the maximum first component and is assigned to X_1 since $|X_1 \cup \{4\}| = 1 \leq 4$. Next, agent 1 with $(3, 2)$ is assigned as $|X_1 \cup \{1\}| = 2 \leq 3$. Agents 2 and 5, both with first component 2, cannot be added since $|X_1| + 1 = 3 > 2$. Moving to resource 2, agents 3, 2, and 5 are assigned sequentially based on their allocation vectors. The final allocation is $(X_1, X_2) = (\{1, 4\}, \{2, 3, 5\})$.

Theorem 2. *When agents have monotone nonincreasing valuations, an MMS allocation always exists and can be computed in polynomial time.*

Proof Sketch. First, we show that every agent is assigned to some resource. Suppose for contradiction that agent i is unassigned in the final allocation $\mathbf{X} = (X_1, \ldots, X_m)$. Then, for every resource j, we must have $x_j^{(i)} < |X_j| + 1$, where $\mathbf{x}^{(i)}$ is i's MMS_i allocation vector. Because the entries are integral, this gives $x_j^{(i)} \leq |X_j|$ for all j; summing over j yields $n = \sum_j x_j^{(i)} \leq \sum_j |X_j| < n$, a contradiction. Next, we show that each agent receives at least her maximin share under $\mathbf{X}$. Suppose that agent i is finally assigned to resource j. Let i' be the last agent added to X_j. The algorithm ensured $|X_j| \leq x_j^{(i')}$. Since i was chosen earlier, we have $x_j^{(i)} \geq x_j^{(i')} \geq |X_j|$. Thus, $v_i(j, |X_j|) \geq v_i(j, x_j^{(i)}) \geq \mathrm{MMS}_i$. □

Monotone Nondecreasing Valuations. If agents have monotone nondecreasing valuations, the MMS_i value of each agent i is simply the maximum utility she can achieve. Specifically, $\mathrm{MMS}_i = \max_{j \in M} v_i(j, n)$. Intuitively, if agent i could dictate an allocation vector, her best strategy would be to place all n agents on the single resource that gives her the highest value at full capacity. However, it is not always possible for every agent to achieve her maximum attainable valuation (see Example 1 in the Appendix).

Despite the possible nonexistence of MMS allocations under nondecreasing valuations, one may question whether determining their existence is computationally feasible. We show that this decision problem is, in fact, intractable.

Theorem 3. *When agents have monotone nondecreasing valuations, determining whether a given instance admits an MMS allocation is NP-complete.*

4 EF and NS Allocation

In this section, we investigate the existence of EF and NS allocations under two distinct classes of monotone valuations.

Monotone Nondecreasing Valuations. We first note that under monotone nondecreasing valuations, an NS allocation may not exist (see Example 2 in the Appendix). We then prove that deciding its existence is NP-complete.

Theorem 4. *When agents have monotone nondecreasing valuations, determining whether a given instance admits an NS allocation is NP-complete.*

Finally, we characterize the conditions that guarantee the existence of an NS allocation.

Observation 1. *If agents have monotone nondecreasing valuations and $|M| = 2$, then an NS allocation always exists and can be found in $O(n)$ time.*

When all agents have identical valuation functions (i.e. $v_1 = v_2 = \cdots = v_n$), our problem reduces to a class of congestion games. In this case, an NS allocation corresponds exactly to a pure Nash equilibrium of the game.

Observation 2. *If agents have identical valuations, an NS allocation always exists.*

It is well-known that computing a pure strategy Nash equilibrium in potential games is PLS (Polynomial Local Search) complete. Therefore, when agents have identical valuations, computing an NS allocation is also PLS-complete.

Corollary 1. *When agents have identical valuations, computing an NS allocation is PLS-complete.*

Monotone Nonincreasing Valuations. Next, we consider the case where agents have monotone nonincreasing valuations. We first observe that under monotone nonincreasing valuations, an EF allocation may not exist (see Example 3 in the Appendix). We aim to further explore whether determining the existence of EF allocations is computationally hard under monotone nonincreasing valuations. However, we are unable to establish such a result. Instead, we prove a complexity result for the EF existence problem under the more general assumption that agents' valuations exhibit different types of monotonicity.

Theorem 5. *Determining whether a given instance admits an EF allocation is NP-complete, even when all agents' valuations are monotone (either nonincreasing or nondecreasing).*

Algorithm 3: Computing an NS allocation for monotone nonincreasing valuations

Input: An instance I with nonincreasing valuations.
Output: An NS allocation $\mathbf{X} = (X_1, \ldots, X_m)$.

Initialize: $X_1 = \cdots = X_m = \emptyset$, $N' = N$, and $l = 0$.
while $N' \neq \emptyset$ **do**
  Let i' be an arbitrary agent in N'.
  Let j' be the resource that i' most prefers, i.e., $j' \in \arg\max_{j \in M} v_{i'}(j, |X_j| + 1)$.
  Set $X_{j'} = X_{j'} \cup \{i'\}$, $N' = N' \setminus \{i'\}$, and $l = j'$.
  while there exists an agent $i \in X_l$ and a resource $j \in M$ such that $v_i(l, |X_l|) < v_i(j, |X_j \cup \{i\}|)$ **do**
    Let j^* be the resource that i most wants to deviate to, i.e., $j^* \in \arg\max_{j \in M} v_i(j, |X_j \cup \{i\}|)$.
    Set $X_{j^*} = X_{j^*} \cup \{i\}$, $X_l = X_l \setminus \{i\}$, and $l = j^*$.
  end while
end while
return $\mathbf{X} = (X_1, \ldots, X_m)$.

Although EF allocations do not always exist, NS allocations are guaranteed to exist. We provide a polynomial-time algorithm (Algorithm 3) to compute an NS allocation. The high-level idea of the algorithm is to start with an empty allocation, which is also an NS allocation, and iteratively assign agents while maintaining Nash stability.

At each iteration of the outer `while` loop, an unassigned agent i' is assigned to the resource j' that maximizes i''s utility, i.e., $j' \in \arg\max_{j \in M} v_{i'}(j, |X_j \cup \{i'\}|)$. After assigning i' to j', we iteratively reassign agents to maintain Nash stability: as long as there is an agent i in the last updated resource $l = j'$ who now has an incentive to deviate, we move i to the resource j^* that i most prefers, i.e., $j^* \in \arg\max_{j \in M} v_i(j, |X_j \cup \{i\}|)$). This may in turn cause a different resource to be the last updated, and we continue the process until no agent has an incentive to deviate. It can be shown that in each outer `while` iteration, every already assigned agent moves at most once, so the inner loop terminates after at most $n - 1$ reassignments. We then proceed to the next unassigned agent, and so on, until all agents are assigned.

Theorem 6. *When agents have monotone nonincreasing valuations, an NS allocation always exists and can be computed in $O(n^2)$ time.*

5 Epistemic Relaxations

In this section, we focus on two epistemic relaxations: ENS and EEF. We show that for both classes of monotone valuations, an ENS allocation always exists and can be computed in polynomial time. In contrast, under monotone nonincreasing preferences, an EEF allocation may not exist.

Although an NS allocation does not always exist under monotone nondecreasing valuations, we observe that its relaxation, ENS, always exists and can be computed efficiently (see Algorithm 4): we initialize by assigning all agents to an arbitrary resource j (so $X_j = N$), then iteratively reassign any agent $i \in X_j$ who can improve its utility by moving to another resource $j^* \in \arg\max_{j' \in M \setminus \{j\}} v_i(j', |X_{j'}| + 1)$, and repeat until no agent has an incentive to deviate.

Algorithm 4: Computing an ENS allocation for monotone nondecreasing valuations

Input: An instance I with nondecreasing valuations.
Output: An ENS allocation $\mathbf{X} = (X_1, \ldots, X_m)$.

Initialize: $X_1 = N$; $X_j = \emptyset$ for all $j \neq 1$.
while there exist $i \in X_1$ and $j \in M \setminus \{1\}$ such that $v_i(1, |X_1|) < v_i(j, |X_j| + 1)$ **do**
  Let $j^* \in \arg\max_{j \in M \setminus \{1\}} v_i(j, |X_j| + 1)$.
  Set $X_{j^*} = X_{j^*} \cup \{i\}$ and $X_1 = X_1 \setminus \{i\}$.
end while
return $\mathbf{X} = (X_1, \ldots, X_m)$.

Theorem 7. *When agents have monotone nondecreasing valuations, an ENS allocation always exists and can be computed in polynomial time.*

Under monotone nondecreasing valuations, ENS implies EEF, yielding the following corollary by Theorem 7.

Corollary 2. *When agents have monotone nondecreasing valuations, an EEF allocation always exists and can be computed in polynomial time.*

Under monotone nonincreasing valuations, Theorem 6 guarantees the existence of an NS allocation. Since any NS allocation is trivially ENS, we have the following corollary.

Corollary 3. *When agents have monotone nonincreasing valuations, an ENS allocation always exists and can be computed in polynomial time.*

Furthermore, in the special case of two resources, EEF coincides with EF. Since each agent knows the number of sharers on her own resource, she can infer the number of sharers on the other resource, thus eliminating any information asymmetry. Therefore, Example 3 (the two-resource counterexample) demonstrates that EEF allocations may not exist under monotone nonincreasing valuations.

Corollary 4. *When agents have monotone nonincreasing valuations, an EEF allocation may not exist.*

6 Conclusion and Future Directions

In this work, we introduced the shared resource allocation problem and extended several classical fairness and stability concepts to this new setting. We analyzed the existence and computational complexity of fair allocations under two classes of monotone valuations (summarized in Table 1).

Several open questions and promising directions remain for future work. One natural problem is to determine the existence and computational complexity of EF allocations under monotone nondecreasing valuations. Another worthwhile direction is to explore combined fairness and efficiency criteria. For instance, one could seek an allocation that is both epistemic envy-free and Pareto optimal. More broadly, almost every research theme from classical fair division can be revisited in the context of shared resources. Finally, from an applied perspective, it would be valuable to study online variants of the model in which agents arrive dynamically.

Ethical Statement

There are no ethical issues.

Acknowledgments

This research was supported in part by the National Natural Science Foundation of China (Project No. 12201590, 12171444) and Natural Science Foundation of Shandong Province (Project No. ZR2024MA031). Minming Li is supported by a grant from Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. CityU 11216725).

References

Amanatidis, G.; Markakis, E.; Nikzad, A.; and Saberi, A. 2017. Approximation Algorithms for Computing Maximin Share Allocations. *ACM Trans. Algorithms*, 13(4).

Aziz, H.; Bouveret, S.; Caragiannis, I.; Giagkousi, I.; and Lang, J. 2018. Knowledge, Fairness, and Social Constraints. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).

Aziz, H.; Brandl, F.; Brandt, F.; Harrenstein, P.; Olsen, M.; and Peters, D. 2019. Fractional Hedonic Games. *ACM Trans. Econ. Comput.*, 7(2).

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

Aziz, H.; and Mackenzie, S. 2016. A Discrete and Bounded Envy-Free Cake Cutting Protocol for Any Number of Agents. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, 416–427.

Aziz, H.; Rauchecker, G.; Schryen, G.; and Walsh, T. 2017. Algorithms for Max-Min Share Fair Allocation of Indivisible Chores. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1).

Barman, S.; Biswas, A.; Krishnamurthy, S.; and Narahari, Y. 2018. Groupwise Maximin Fair Allocation of Indivisible Goods. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).

Bhaskar, U.; Sricharan, A. R.; and Vaish, R. 2021. On Approximate Envy-Freeness for Indivisible Chores and Mixed Resources. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX-/RANDOM 2021)*, 1:1–1:23. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

Bilò, V.; Fanelli, A.; Flammini, M.; Monaco, G.; and Moscardelli, L. 2018. Nash stable outcomes in fractional hedonic games: Existence, efficiency and computation. *Journal of Artificial Intelligence Research*, 62: 315–371.

Bogomolnaia, A.; and Jackson, M. O. 2002. The Stability of Hedonic Coalition Structures. *Games and Economic Behavior*, 38(2): 201–230.

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

Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Econ. Comput.*, 7(3).

Chan, P.; Huang, X.; Liu, Z.; Zhang, C.; and Zhang, S. 2016. Assignment and Pricing in Roommate Market. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1).

Chung, K.-S. 2000. On the Existence of Stable Roommate Matchings. *Games and Economic Behavior*, 33(2): 206–230.

Conitzer, V.; Freeman, R.; Shah, N.; and Vaughan, J. W. 2019. Group Fairness for the Allocation of Indivisible Goods. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 1853–1860.

Darmann, A. 2015. Group Activity Selection from Ordinal Preferences. In *Algorithmic Decision Theory*, 35–51. Cham: Springer International Publishing.

Darmann, A.; Elkind, E.; Kurz, S.; Lang, J.; Schauer, J.; and Woeginger, G. 2012. Group Activity Selection Problem. In Goldberg, P. W., ed., *Internet and Network Economics*, 156–169. Berlin, Heidelberg: Springer Berlin Heidelberg.

Drèze, J. H.; and Greenberg, J. 1980. Hedonic Coalitions: Optimality and Stability. *Econometrica*, 48(4): 987–1003.

Dubins, L. E.; and Spanier, E. H. 1961. How to Cut a Cake Fairly. *The American Mathematical Monthly*, 68(1): 1–17.

Feige, U.; Sapir, A.; and Tauber, L. 2022. A Tight Negative Example for MMS Fair Allocations. In Feldman, M.; Fu, H.; and Talgam-Cohen, I., eds., *Web and Internet Economics*, 355–372. Cham, Switzerland: Springer International Publishing.

Foley, D. K. 1966. *Resource allocation and the public sector*. Yale University.

Gale, D.; and Shapley, L. S. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, 69(1): 9–15.

Gan, J.; Li, B.; and Li, Y. 2023. Your college dorm and dormmates: Fair resource sharing with externalities. *Journal of Artificial Intelligence Research*, 77: 793–820.

Garg, J.; and Taki, S. 2020. An Improved Approximation Algorithm for Maximin Shares. In *Proceedings of the 21st ACM Conference on Economics and Computation*, 379–380. New York, NY, USA: Association for Computing Machinery.

Holzman, R.; and Law-Yone, N. 1997. Strong Equilibrium in Congestion Games. *Games and Economic Behavior*, 21(1): 85–101.

Huang, X.; and Lu, P. 2021. An Algorithmic Framework for Approximating Maximin Share Allocation of Chores. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, 630–631. New York, NY, USA: Association for Computing Machinery.

Huzhang, G.; Huang, X.; Zhang, S.; and Bei, X. 2017. Online Roommate Allocation Problem. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI 2017)*, 235–241. AAAI Press.

Igarashi, A.; Peters, D.; and Elkind, E. 2017. Group Activity Selection on Social Networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1).

Jang, I.; Shin, H.-S.; and Tsourdos, A. 2018. Anonymous Hedonic Game for Task Allocation in a Large-Scale Multiple Agent System. *IEEE Transactions on Robotics*, 34(6): 1534–1548.

Kurokawa, D.; Procaccia, A. D.; and Wang, J. 2018. Fair Enough: Guaranteeing Approximate Maximin Shares. *J. ACM*, 65(2).

Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, 125–131. New York, NY, USA: Association for Computing Machinery.

Manurangsi, P.; and Suksompong, W. 2017. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences*, 89: 100–108.

Milchtaich, I. 1996. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13(1): 111–124.

Moulin, H. 2004. *Fair Division and Collective Welfare*. Cambridge, MA: MIT Press.

Rosenthal, R. W. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1): 65–67.

Segal-Halevi, E.; and Nitzan, S. 2019. Fair cake-cutting among families. *Social Choice and Welfare*, 53(4): 709–740.

Steinhaus, H. 1948. The problem of fair division. *Econometrica*, 16: 101–104.

Suksompong, W. 2018. Approximate maximin shares for groups of agents. *Mathematical Social Sciences*, 92: 40–47.

Thomson, W. 1983. Problems of fair division and the Egalitarian solution. *Journal of Economic Theory*, 31(2): 211–226.