Epistemic Pairwise Maximin Share

Michal Feldman
Tel Aviv University
mfeldman@tauex.tau.ac.il

Amos Fiat
Tel Aviv University
fiat@tau.ac.il

Yael Nissan
Tel Aviv University
yaelshemesh@mail.tau.ac.il

Tomasz Ponitka
Tel Aviv University
tomaszp@mail.tau.ac.il

Abstract

We introduce *epistemic pairwise maximin share* (EPMMS), a new fairness notion for fair division of indivisible goods. Two fundamental notions in this setting are envy-freeness up to any item (EFX) and pairwise maximin share (PMMS), with PMMS being stronger than EFX. While EFX has been extensively studied, far less is known about PMMS. Recent work shows that relaxing EFX via an epistemic perspective leads to substantial progress on the EFX problem, raising the question of whether a similar approach can advance our understanding of PMMS. Motivated by this, we initiate the study of EPMMS, the epistemic relaxation of PMMS. EPMMS is more challenging than EEFX: the key approaches underlying recent progress on epistemic EFX inherently fail to extend to EPMMS. We establish the following results. First, for additive valuations, $4/5$-EPMMS allocations exist and can be efficiently computed. Second, for bivalued valuations, EPMMS allocations exist and can be efficiently computed; in fact, we obtain the stronger guarantee of epistemic groupwise maximin share (EGMMS), which also strengthens the existence of MMS allocations for this setting. Third, EPMMS allocations exist in two settings where MMS allocations need not exist: instances with three additive agents or two types of additive agents.

1 Introduction

The fair division problem asks how to allocate a set of items among n agents with differing preferences so that every agent is treated fairly. It arises in many settings, from allocating seats in university courses [23] to assigning papers to reviewers at conferences [48]. Our primary focus is on the setting with *indivisible* items (an allocation $X_1, \ldots, X_n$ forms a partition of the items into disjoint bundles) and *additive* valuations (agent i's value for a bundle X is $v_i(X) = \sum_{g \in X} v_i(\{g\})$).

For two agents, the canonical solution is the *cut-and-choose* mechanism: one agent divides the items into two bundles, and the other picks their preferred one. With more than two agents, however, the situation becomes more complex, and no obvious canonical solution emerges. A central challenge in fair division is to identify properties of cut-and-choose that extend beyond two agents.

Toward this goal, Caragiannis et al. [25] proposed two fairness notions. The first is *envy-freeness up to any good* (EFX). We say that agent i envies agent j if i prefers j's allocation to their own. Although cut-and-choose guarantees that the chooser does not envy the cutter, it is generally impossible to guarantee envy-freeness (EF) for both agents, as even the simple case of a single item demonstrates. EFX is a relaxation of EF requiring that any envy disappears after hypothetically removing any single good from the envied bundle: $v_i(X_i) \geq v_i(X_j \setminus \{g\})$ for all $i, j \in [n]$ and all $g \in X_j$.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

The second fairness notion is *pairwise maximin share* (PMMS). Consider any pair of agents i and j, and imagine running cut-and-choose over their combined items $X_i \cup X_j$. The best guarantee i could secure as the cutter is their *maximin share*, $\max_{(A,B)} \min\{v_i(A), v_i(B)\}$, taken over all partitions (A, B) of $X_i \cup X_j$. PMMS requires that every agent receives at least their maximin share within every pair they belong to. Under mild assumptions, PMMS is strictly stronger than EFX [24, 25, 34].

Both EFX and PMMS are satisfied by cut-and-choose when there are only two agents, but whether either can be satisfied for an arbitrary number of agents with additive valuations is a long-standing open problem. The existence of EFX is known in important special cases: for three agents [4, 28], for bivalued valuations [10, 24, 41], and for instances with few distinct valuation types [40, 50]. Despite being introduced in the same paper as EFX, PMMS remains significantly less explored, and as the stronger notion, much less is known about its existence.

As a successful line of attack on the EFX problem, Caragiannis et al. [26] introduced a compelling relaxation of EFX called *Epistemic EFX* (EEFX). An allocation is EEFX if each agent could hypothetically rearrange the items held by the other agents into bundles that, from their own perspective, constitute an EFX allocation. Notably, each agent can verify fairness on the basis of their own bundle alone, without reference to how the remaining items are actually partitioned among the other agents.

This relaxation has proven more tractable than EFX itself. Caragiannis et al. [26] established the existence of EEFX allocations under additive valuations via a reduction to the identical-rankings setting, followed by an application of Envy Cycle Elimination algorithm. Akrami and Rathi [2] subsequently extended this result to arbitrary monotone valuations via the Lone Divider approach.

Motivated by the progress on EFX enabled by the epistemic approach, and by the limited understanding of PMMS, we initiate the study of *Epistemic PMMS* (EPMMS), the analogous relaxation of PMMS. Notably, both of the two approaches underlying the existing EEFX results inherently fail to carry over to EPMMS, as we demonstrate through explicit counterexamples in Appendix A.

Beyond comparison-based fairness notions like EFX and PMMS, a substantial body of work concerns *share-based* fairness. The central share-based notion is the *maximin share* (MMS), defined as the maximum value a cutter can guarantee by partitioning the entire set of items into n bundles, and receiving the least valuable bundle. It is known that allocations giving every agent at least their MMS need not exist in general, even for three agents with additive valuations [32, 45]. The epistemic approach does not apply to MMS directly, as MMS is defined independently of how items are distributed among the remaining agents. Nonetheless, every MMS allocation satisfies EEFX [26]. By contrast, in Appendix A, we show that MMS does not imply EPMMS, which yields a further separation between EEFX and EPMMS. See Figure 1 for a broader picture of these relationships.

More broadly, a recurring theme in the literature on EFX and PMMS is the role of *PMMS-feasibility*[1], a structural condition that generalizes additivity and also captures budget-additive, unit-demand, and multiplicative valuations [4]; see Section 2 and Appendix C. Without any assumptions, both notions face strong impossibilities: EFX can fail to exist under monotone valuations for three or more agents [6], and PMMS fails to exist already for two agents [24]. PMMS-feasibility circumvents these impossibilities: EFX exists for three agents when at least one is PMMS-feasible [4], and PMMS exists for two agents when one of them is PMMS-feasible [24].

Until quite recently the question of if PMMS exists for $n > 2$ agents was wide open. However, some very recent papers (September 5 and September 8, 2026) [13, 36] show that PMMS need not exist even for three agents and additive valuations. These new papers are part of the ongoing research revolution using the new AI tools.

Notably, EEFX is guaranteed to exist for arbitrary monotone valuations [2], whereas EPMMS inherits the non-existence of PMMS for two monotone agents, since the two notions coincide in this case. This marks yet another contrast between EEFX and EPMMS.

[1] PMMS-feasibility is exactly the condition introduced by Akrami et al. [4] under the name MMS-feasibility. We refer to it as PMMS-feasibility for three reasons: first, to align the terminology with the GMMS-feasibility condition that we introduce in this paper (see Definition 2.11); second, to emphasize that PMMS-feasibility (resp. GMMS-feasibility) precisely characterizes the valuations for which envy-freeness implies PMMS (resp. GMMS), see Appendix C.3; and third, because the original name can potentially be misleading: although PMMS-feasibility guarantees the existence of an MMS allocation for two agents (where MMS coincides with PMMS), no such guarantee holds for three or more agents.

1.1 Our Results

We begin with additive valuations, for which we obtain a multiplicative approximation of EPMMS.

Main Theorem 1 (Additive Valuations; see Theorem 4.1)**.** *Every instance with additive valuations admits a $4/5$-EPMMS allocation. Moreover, this allocation can be computed in polynomial time.*

To establish the existence of $4/5$-EPMMS for additive valuations, we prove that $4/5$-PMMS allocations exist under the *identical-rankings* assumption[2]. For comparison, Kurokawa [43] proved the existence of 0.781-PMMS allocations for additive valuations; our result shows that the epistemic relaxation improves this factor to 0.8 (but for the weaker notion of EPMMS).

We next establish the existence of exact EPMMS allocations in several important settings. While EFX is known to exist in each of these settings [4, 10, 24, 28, 40, 41, 50], EPMMS and EFX are incomparable notions, so our results neither imply nor are implied by these EFX existence results.

First, we consider *bivalued* valuations, in which each agent i's valuation function v_i is additive and satisfies $v_i(\{g\}) \in \{a, b\}$ for every item g and some fixed $a, b \geq 0$. In this setting, we in fact establish a stronger guarantee: the existence of an *Epistemic GMMS* (EGMMS) allocation. The groupwise maximin share (GMMS) criterion strengthens both PMMS and MMS by requiring that each agent receive at least their maximin share within every subgroup of agents to which they belong; EGMMS is the analogous epistemic relaxation (see Definition 2.8).

Main Theorem 2 (Bivalued Valuations; see Theorem 5.1)**.** *Every instance with bivalued valuations admits an EGMMS, and hence an EPMMS, allocation. Moreover, this allocation can be computed in polynomial time.*

In particular, our result strengthens that of Feige [31], who proved the existence of MMS allocations for bivalued instances, since EGMMS implies MMS. We note that EGMMS is a strictly stronger notion: in fact, even PMMS and MMS simultaneously do not imply EGMMS; see Appendix B. We now turn to the case of three agents.

Main Theorem 3 (Three Agents; see Theorem 6.1)**.** *Every instance with three additive agents admits an EPMMS allocation. More generally, this holds whenever at least two of the agents are PMMS-feasible.*

Finally, we consider instances with two *types* of agents, meaning that the agents are partitioned into two groups within which all members share the same valuation function.

Main Theorem 4 (Two Types of Agents; see Theorem 6.2)**.** *Every instance with two types of additive agents admits an EPMMS allocation. More generally, this holds whenever at least one of the types is PMMS-feasible.*

For both of the last two settings, MMS allocations need not exist [32]. Consequently, since EGMMS implies MMS, neither theorem can be strengthened to EGMMS in general.

1.2 Related Works

In this subsection, we summarize the literature most closely related to our work. For a broader overview of the field of fair division, see the survey of Amanatidis et al. [11].

Epistemic-EFX. The epistemic perspective on fairness was initiated by Aziz et al. [14], who introduced *epistemic envy-freeness*. Building on this idea, Caragiannis et al. [26] introduced the analogous relaxation of EFX, namely EEFX, and established its existence for additive valuations. Akrami and Rathi [2] subsequently extended the existence result to arbitrary monotone valuations. More

[2]Both our work and that of Dai et al. [30] claim the existence of a $4/5$-PMMS allocation for instances with identical rankings via the Envy Cycle Elimination algorithm. However, the proof of Dai et al. [30] contains issues that, to the best of our knowledge, remain unresolved. Crucially, they do not argue that an envy cycle elimination step (Line 8 of Algorithm 1) preserves the $4/5$-PMMS guarantee. This does not follow automatically: in Appendix A.1, we show that in general a single envy cycle elimination step can only guarantee $2/3$-PMMS, even when starting from an exact PMMS allocation. We show that the $4/5$-PMMS guarantee can nevertheless be recovered for the Envy Cycle Elimination algorithm in the setting with identical rankings, but only under the additional assumption that the valuations are non-degenerate, which is missing from the analysis of Dai et al. [30]. In Appendix G.1, we provide examples showing that without non-degeneracy, the Envy Cycle Elimination algorithm does *not* guarantee $4/5$-PMMS, even under identical rankings and even in the absence of zero-valued items.

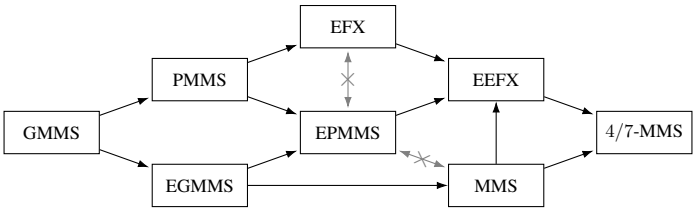


Figure 1: Implications between fairness notions under non-degenerate additive valuations. Most implications follow directly from the definitions; the following require proof: (1) PMMS $\Rightarrow$ EFX and EPMMS $\Rightarrow$ EEFX on non-degenerate instances [25]; (2) MMS $\Rightarrow$ EEFX [26]; and (3) EEFX $\Rightarrow$ $4/7$-MMS [8]. In Appendix A.3, we show that EPMMS is incomparable to both MMS and EFX.

recently, Akrami et al. [5] strengthened these guarantees by showing that EEFX can be achieved simultaneously with envy-freeness up to some item (EF1).

PMMS. Caragiannis et al. [25] showed that Nash-welfare-maximizing allocations are 0.618-PMMS. The best known approximation of PMMS for additive valuations is 0.781, due to Kurokawa [43]. Exact PMMS allocations exist in several restricted settings: Byrka et al. [24] established existence for binary-valued and pair-demand valuations, and Christodoulou and Mastrakoulis [29] showed existence under additive valuations whenever each item has non-negative value for at most two agents. PMMS has also been studied alongside other fairness notions: Amanatidis et al. [8] showed that EFX implies $2/3$-PMMS and EF1 implies $1/2$-PMMS; Amanatidis et al. [9] obtained allocations that are simultaneously 0.618-EFX and 0.717-PMMS; and Feldman et al. [33] achieved 0.618-EFX, 0.618-PMMS, and a 0.618-fraction of the maximum Nash welfare simultaneously.

GMMS. The GMMS notion was introduced by Barman et al. [18]. For additive valuations, the best known multiplicative approximation is $4/7$, established independently by Amanatidis et al. [9] and Chaudhury et al. [27]. Several works have further studied GMMS in conjunction with other fairness notions: Amanatidis et al. [9] established the existence of 0.55-GMMS and 0.618-EFX allocations, while Feldman et al. [33] obtained allocations that are simultaneously 0.44-GMMS, 0.618-EFX, and achieve a 0.618-fraction of the maximum Nash welfare. Beyond additive valuations, Barman and Verma [17] established a $1/6$-GMMS approximation for XOS valuations.

Bivalued Valuations. The class of bivalued valuations, where $v_i(\{g\}) \in \{a, b\}$ for every agent i and item g, has received significant attention as a natural restriction under which several fairness notions become tractable. Feige [31] established the existence of MMS allocations in this setting, in contrast to the additive case [32, 44]. Amanatidis et al. [10] showed that EFX allocations always exist under bivalued valuations, and this was extended to personalized bivalued valuations, where $v_i(\{g\}) \in \{a_i, b_i\}$, by Byrka et al. [24] and Jin and Tao [41]. For the important special case of binary additive valuations ($v_i(\{g\}) \in \{0, 1\}$), Byrka et al. [24] further showed that PMMS allocations exist.

Few Agents. A natural line of research approaches the EFX existence problem by restricting the number of agents. The case of two agents was settled by Plaut and Roughgarden [51], who showed that EFX allocations always exist; for three agents, existence was established by Chaudhury et al. [28] for additive valuations and subsequently extended to PMMS-feasible valuations by Akrami et al. [4]. For four agents, Berger et al. [21] showed that EFX allocations exist provided that one item can be left unallocated, and Ashuri et al. [12] established the existence of EF2X allocations, a relaxation of EFX in which envy must vanish after removing any two items from the envied bundle.

Few Types of Agents. A complementary line of research restricts the number of distinct valuation functions rather than the number of agents. Mahara [50] established the existence of EFX allocations when the agents partition into at most two types, and HV et al. [40] extended this to three types. For an arbitrary number of types k, Ghosal et al. [35] showed that EFX allocations exist provided that $k - 2$ items can be left unallocated.

2 Model and Preliminaries

We consider a set of n agents, $N = \{1, 2, \ldots, n\}$, and a set M of m indivisible goods. Each agent $i \in N$ is associated with a monotone valuation function $v_i : 2^M \to \mathbb{R}_{\geq 0}$. We will specify our

assumptions on the valuation classes later. For notational convenience, we write $v_i(g)$ instead of $v_i(\{g\})$ for each $g \in M$.

Our goal is to study allocations that assign all goods in M to the agents. Formally, such an allocation is a partition $X = (X_1, X_2, \ldots, X_n)$ of M, where $X_i \cap X_j = \emptyset$ for all $i \neq j$, and $\bigcup_{i \in N} X_i = M$, and X_i is the set of goods allocated to agent i. We denote by $\Pi_n(S)$ the set of all partitions of a set S into n bundles. We now introduce the fairness notions, beginning with the maximin share.

Definition 2.1 (Maximin Share, [23])**.** *Given n agents and a subset $S \subseteq M$ of goods, the n-maximin share of agent i with respect to S is defined as: $\mu_i(n, S) = \max_{X \in \Pi_n(S)} \min_{X_j \in X} v_i(X_j)$*

Fairness Notions. We now define the following fairness notions. We refer to Figure 1 for an illustration of the implications between those fairness notions.

Definition 2.2 (α-EFX, [25])**.** *An allocation $X = (X_1, ..., X_n)$ is α-envy-free up to any good (α-EFX) if, for any pair of agents $i, j \in N$, and any item $g \in X_j$, it holds that $v_i(X_i) \geq \alpha v_i(X_j \setminus \{g\})$.*

Definition 2.3 (α-MMS, [23])**.** *Let $\alpha \in [0, 1]$. An allocation $X = (X_1, ..., X_n)$ is α-maximin-share-fair (α-MMS) if, for every $i \in N$, it holds that $v_i(X_i) \geq \alpha \mu_i(n, M)$.*

Definition 2.4 (α-PMMS, [25])**.** *Let $\alpha \in [0, 1]$. An allocation X is α-pairwise-maximin-share-fair (α-PMMS) if, for any pair of agents $i, j \in N$, it holds that $v_i(X_i) \geq \alpha \mu_i(2, X_i \cup X_j)$.*

Whenever this inequality is violated for agents i and j, we say that agent i α-PMMS-envies agent j.

Definition 2.5 (α-GMMS, [18])**.** *An allocation X is a α-groupwise-maximin-share-fair (α-GMMS) if, for all $J \subseteq [n]$ with $i \in J$, it holds that $v_i(X_i) \geq \alpha \mu_i(|J|, \cup_{j \in J} X_j)$.*

Note that GMMS strictly generalizes both MMS and PMMS [18]. Moreover, for non-degenerate valuations (i.e., valuations under which no two bundles share the same value), PMMS implies EFX [24, 25, 34]. Next, we define the epistemic variants of the fairness notions above. The notion of EEFX has been defined by [26], and the notions of EPMMS and EGMMS are introduced herein.

Definition 2.6 (EEFX, [26])**.** *An allocation $X = (X_1, X_2, \ldots, X_n)$ is* Epistemic α-EFX (α-EEFX) *if for every $i \in [n]$, there exists an allocation $Y^i = (Y^i_1, Y^i_2, \ldots, Y^i_n)$ such that $Y^i_i = X_i$ and agent i is α-EFX-satisfied with allocation Y^i, i.e., $v_i(Y^i_i) \geq \alpha v_i(Y^i_j \setminus \{g\})$ for all j and $g \in Y^i_j$.*

Definition 2.7 (EPMMS)**.** *An allocation $X = (X_1, X_2, \ldots, X_n)$ is* Epistemic α-PMMS (α-EPMMS) *if for every $i \in [n]$, there exists an allocation $Y^i = (Y^i_1, Y^i_2, \ldots, Y^i_n)$ such that $Y^i_i = X_i$ and agent i is α-PMMS-satisfied with allocation Y^i, i.e., $v_i(Y^i_i) \geq \alpha \mu_i(2, Y^i_i \cup Y^i_j)$ for all j.*

Definition 2.8 (EGMMS)**.** *An allocation $X = (X_1, X_2, \ldots, X_n)$ is* Epistemic α-GMMS (α-EGMMS) *if for every $i \in [n]$, there exists an allocation $Y^i = (Y^i_1, Y^i_2, \ldots, Y^i_n)$ such that $Y^i_i = X_i$ and agent i is α-GMMS-satisfied with Y^i, i.e., $v_i(X_i) \geq \alpha \mu_i(|J|, \cup_{j \in J} Y^i_j)$ for all $i \in J \subseteq [n]$.*

For each notion above, the allocation Y^i is called a *certificate* of agent i for bundle X_i. When $\alpha = 1$, we omit it from the notation, writing, e.g., EPMMS rather than 1-EPMMS.

Valuation Classes. We next define the valuation classes considered in this work, beginning with the standard class of additive valuations.

Definition 2.9 (Additive Valuations)**.** *A valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is additive if for every subset of goods $S \subseteq M$, we have $v_i(S) = \sum_{g \in S} v_i(g)$.*

Beyond additive valuations, we consider two natural generalizations: the class of PMMS-feasible valuations, introduced by [4], and the class of GMMS-feasible valuations, which we introduce here.

Definition 2.10 (PMMS-feasible Valuations, [4])**.** *A valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is PMMS-feasible if for every subset of goods $S \subseteq M$ and partitions $A = (A_1, A_2)$ and $B = (B_1, B_2)$ of S, we have $\max\{v(B_1), v(B_2)\} \geq \min\{v(A_1), v(A_2)\}$.*

Definition 2.11 (GMMS-Feasible Valuations)**.** *A valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is GMMS-feasible if for any subset of goods $S \subseteq M$, every $k \leq m$, and any pair of partitions $A = (A_1, \ldots, A_k)$ and $B = (B_1, \ldots, B_k)$ of S, we have $\max\{v(B_1), \ldots, v(B_k)\} \geq \min\{v(A_1), \ldots, v(A_k)\}$.*

3 Structural Insights

Relations Between the Valuation Classes.

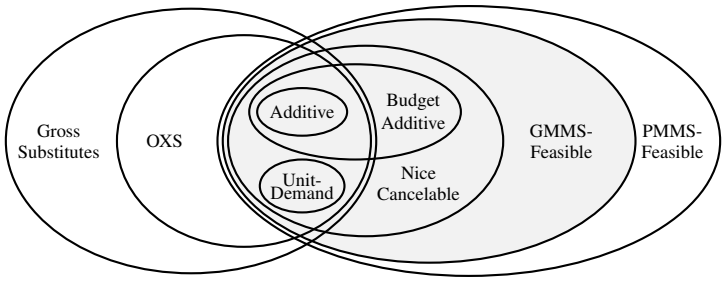


Figure 2: An illustration of the relationships between the valuation classes derived in Appendix C.

In Appendix C, we establish several structural properties of GMMS-feasible valuations. First, GMMS-feasible valuations form a *strict* subclass of PMMS-feasible valuations. Second, GMMS-feasible valuations form a *strict* superclass of nice cancelable valuations [21], which contain all additive, budget-additive, unit-demand, and multiplicative valuations. Third, PMMS-feasibility (resp. GMMS-feasibility) precisely characterizes valuations where envy-freeness implies PMMS (resp. GMMS). Finally, the class of gross substitutes (GS) valuations is incomparable with both MMS- and GMMS-feasible valuations: there exist GS valuations that are not PMMS-feasible, and GMMS-feasible valuations that are not GS. See Figure 2 for an illustration of these relations.

The Leximin Solution. The leximin (lexicographic maximin) solution is a foundational concept in fair division [51, 52] that selects the allocation which maximizes the minimum utility; then, if there are multiple allocations which achieve that minimum utility, it chooses among those the allocation which maximizes the second minimum utility, and so on. We adopt the following formal definition.

Definition 3.1 (Leximin Solution, [51])**.** *Fix a valuation function v. Let $\Pi_n(M)$ be the set of all possible allocations. For any $X \in \Pi_n(M)$, let $v(X) = (v(X_1), \ldots, v(X_n))$ be the utility profile, and let $v^{\uparrow}(X)$ be this profile sorted in non-decreasing order. An allocation X is a leximin solution under v if it lexicographically maximizes $v^{\uparrow}(X)$ over all $X \in \Pi_n(M)$. That is, X is strictly preferred to Y if at the first index k where $v_k^{\uparrow}(X) \neq v_k^{\uparrow}(Y)$, we have $v_k^{\uparrow}(X) > v_k^{\uparrow}(Y)$.*

Barman et al. [18] showed that any leximin allocation is GMMS when all agents have identical valuations. Since every GMMS allocation is, by definition, both a PMMS and an MMS allocation, it naturally follows that the leximin solution guarantees both properties for identical-valuation instances.

Proposition 3.2 ([18])**.** *For instances where all agents have identical monotone valuations (i.e., $v_i = v$ for all $i \in N$ and some v), any leximin allocation under v is a GMMS allocation.*

Identical-Ranking Instances. We now show how to reduce instances with arbitrary rankings over items to instances in which all agents share an identical ranking.

Definition 3.3 (Identical-Ranking Instances)**.** *An instance has identical rankings (or is an ordered instance) if there exists a common ordering of the goods $g_1, g_2, \ldots, g_m$ such that for every agent $i \in N$, it holds that $v_i(g_1) \geq v_i(g_2) \geq \ldots \geq v_i(g_m)$.*

The reduction to identical rankings is a powerful tool in fair division, as it enables one to establish existence results on a significantly simpler instance. This perspective has been particularly useful in the study of Maximin Share (MMS) allocations, where it has been leveraged to compute $2/3$-MMS allocations [7, 16], as well as several other MMS approximation results [22, 45]. More recently, [26] employed this framework to prove the existence of Epistemic EFX (EEFX) allocations for additive valuations. In this paper, we build on this approach to obtain guarantees for EPMMS approximation under additive valuations and EGMMS allocations under bivalued valuations.

Proposition 3.4 (Reduction to Identical Rankings)**.** *Let $0 < \alpha \leq 1$. First, if an α-EPMMS allocation exists for all additive instances with identical rankings, then an α-EPMMS allocation exists for all additive instances with arbitrary rankings. Second, if an α-EGMMS allocation exists for all bivalued instances with identical rankings, then an α-EGMMS allocation exists for all bivalued instances with arbitrary rankings. Moreover, in both cases, if the former allocations can be computed in polynomial time, then so can the latter.*

Algorithm 1: Envy Cycle Elimination Algorithm [49]

Input: An instance with non-degenerate additive valuations and identical rankings.
Output: A $4/5$-PMMS and EFX allocation.

Initialize $X \leftarrow (\emptyset, \ldots, \emptyset)$ and $U \leftarrow M$;
while $U \neq \emptyset$ **do**
  if $\exists\, i \in [n]$ *such that* $v_j(X_j) \geq v_j(X_i)$ *for all* $j \in [n]$ **then**
    Let $g \in \arg\max_{g' \in U} v_i(g')$ be the most valuable item in U;
    $X_i \leftarrow X_i \cup \{g\}$ and $U \leftarrow U \setminus \{g\}$; `// Assign` g `to an unenvied agent.`
  else
    Let $(i_1, \ldots, i_\ell)$ satisfy $v_{i_r}(X_{i_{r+1}}) > v_{i_r}(X_{i_r})\ \forall r \in [\ell]$, with indices taken mod ℓ;
    $X_{i_r} \leftarrow X_{i_{r+1}}$ for all $r \in [\ell]$; `// Eliminate an envy cycle.`
return X;

It is worth noting that the reduction for α-EGMMS extends to additive valuations as well. We state it here for bivalued instances, as this is the setting for which we establish EGMMS guarantees. While the reduction algorithm itself is essentially identical to that of [26] for EEFX (and similarly to [7] for $2/3$-MMS), the analysis is more involved. Unlike these settings, where one reasons about the removal of a single item from another agent's bundle, here we must handle substantially richer deviations: arbitrary repartitions of pairs of bundles for EPMMS, and combinations of multiple bundles for EGMMS. These features introduce new technical challenges that require a more delicate argument. The proof of Proposition 3.4 and the presentation of the ordering framework appear in Appendix F.

4 Approximating EPMMS for Additive Valuations

In this section, we show that 4/5-EPMMS allocation exists for additive valuations.

Theorem 4.1 (Additive Valuations)**.** *Every instance with additive valuations admits a* $4/5$*-EPMMS allocation. Moreover, this allocation can be computed in polynomial time.*

In fact, our proof yields a stronger guarantee: an allocation that is simultaneously $4/5$-EPMMS, EEFX and $2/3$-MMS. By the reduction in Proposition 3.4, it suffices to prove the following proposition in order to establish Theorem 4.1.

Proposition 4.1 (Additive Valuations with Identical Rankings)**.** *Every instance with additive valuations and identical rankings admits a* $4/5$*-PMMS, EFX and* $2/3$*-MMS allocation. Moreover, this allocation can be computed in polynomial time.*

Before proving Proposition 4.1, we state a useful lemma.

Lemma 4.2. *Fix agent i, and let P, Q be two disjoint sets of items. Then, for any $S \subseteq P \cup Q$,*

$$\mu_i(2, P \cup Q) \leq \mu_i(2, (P \cup Q) \setminus S) + v_i(S).$$

Proof. Assume that (A_1, A_2) is a partition of $P \cup Q$ such that $\min\{v(A_1), v(A_2)\} = \mu_i(2, P \cup Q)$. Without loss of generality, assume that $v_i(A_1 \setminus S) \leq v_i(A_2 \setminus S)$. Since $(A_1 \setminus S, A_2 \setminus S)$ is a partition of $(P \cup Q) \setminus S$, it follows that $v_i(A_1 \setminus S) \leq \mu_i(2, (P \cup Q) \setminus S)$. Thus,

$$\mu_i(2, P \cup Q) \leq v_i(A_1) \leq v_i(A_1 \setminus S) + v_i(S) \leq \mu_i(2, (P \cup Q) \setminus S) + v_i(S). \qquad \square$$

We give a partial proof of Proposition 4.1, deferring some details to Appendix E.1.

(Partial) Proof of Proposition 4.1. We assume without loss of generality that the valuations are non-degenerate, meaning $v_i(S) \neq v_i(T)$ for every agent i and every pair of distinct bundles S, T; see Proposition G.4 for a formal argument.

We claim that the Envy-Cycle Elimination algorithm (Algorithm 1, [49]) produces a $4/5$-PMMS and EFX allocation in polynomial time. Polynomial-time termination was established by Lipton et al. [49], and Plaut and Roughgarden [51], Barman and Krishnamurthy [16] showed that the resulting allocation satisfies EFX and $2/3$-MMS; it therefore remains to prove the $4/5$-PMMS guarantee.

Let X be an allocation produced by this procedure, and let $i, j \in N$. We want to show that $v_i(X_i) \geq \frac{4}{5}\mu_i(2, X_i \cup X_j)$.

Consider the point during the execution of Algorithm 1 at which bundle X_j is formed: some agent j' holds a bundle $\tilde{X}_{j'} = X_j \setminus \{g\}$, and item g is added to it in Line 5, producing X_j. Because agent i does not envy agent j' at the moment g is added, and since the value an agent holds for their own bundle can only increase over time, we establish that $v_i(X_i) \geq v_i(X_j \setminus \{g\})$. If $v_i(g) \leq \frac{1}{2}v_i(X_i)$,

$$\begin{aligned}
\mu_i(2, X_i \cup X_j) &\leq \tfrac{1}{2}(v_i(X_i) + v_i(X_j)) && \text{(since } \mu_i(2,S) \leq \tfrac{1}{2}v_i(S) \text{ for all } S\text{)} \\
&= \tfrac{1}{2}(v_i(X_i) + v_i(X_j \setminus \{g\}) + v_i(g)) && \text{(by additivity of } v_i\text{)} \\
&\leq \tfrac{1}{2}(v_i(X_i) + v_i(X_i) + \tfrac{1}{2}v_i(X_i)) && \text{(by the above)} \\
&= \tfrac{5}{4}v_i(X_i).
\end{aligned}$$

Otherwise, $v_i(g) > \frac{1}{2}v_i(X_i)$. When g is added to X_j, some agent i' holds a bundle $\tilde{X}_{i'} \subseteq X_i$. The full proof addresses the subcase $|X_j| = 1$ separately; for the subcase $|X_j| \geq 2$, a careful argument, included in the full proof, leveraging the identical-rankings and non-degeneracy assumptions shows that $\tilde{X}_{i'} = \{h\}$ and $X_j = \{g', g\}$, with the strict ordering $v_i(h) > v_i(g') > v_i(g)$, and we obtain

$$\begin{aligned}
\mu_i(2, X_i \cup X_j) &\leq \mu_i(2, (X_i \cup X_j) \setminus (X_i \setminus \{h\})) + v_i(X_i \setminus \{h\}) && \text{(by Lemma 4.2)} \\
&= \mu_i(2, \{g, g', h\}) + v_i(X_i \setminus \{h\}) && \text{(by substitution)} \\
&\leq v_i(h) + v_i(X_i \setminus \{h\}) && \text{(by the value hierarchy)} \\
&= v_i(X_i) && \text{(by additivity)} \qquad \square
\end{aligned}$$

Remark 4.3. *In fact, our proof of the existence of a $4/5$-PMMS allocation holds under the slightly weaker assumption that the agents agree on the identity and ranking of the top n items, i.e., there exist distinct items $g_1, \ldots, g_n$ such that $v_i(g_1) \geq v_i(g_2) \geq \cdots \geq v_i(g_n) \geq v_i(g)$ for every agent i and every $g \notin \{g_1, \ldots, g_n\}$. We note, however, that EFX requires the identical-rankings assumption.*

Remark 4.4. *In Appendix D, we show that this analysis is tight by exhibiting an instance on which Algorithm 1 achieves no better than a $4/5$ approximation of PMMS.*

5 EGMMS for Bivalued Valuations

In this section we establish the existence of EGMMS for instances with bivalued valuations.

Definition 5.1 (Bivalued Valuations)**.** *An instance has bivalued valuations if the valuations are additive and there exist common values $0 \leq a < b$ with $v_i(\{g\}) \in \{a, b\}$ for all $i \in N$ and $g \in M$.*

Theorem 5.1 (Bivalued Valuations)**.** *Every instance with bivalued valuations admits an EGMMS, and hence an EPMMS, allocation. Moreover, this allocation can be computed in polynomial time.*

By Proposition 3.4, to establish the existence of EGMMS for general additive valuations, it suffices to prove it for instances with identical rankings. In this setting, given the common item ordering $M = (g_1, \ldots, g_m)$, each agent i has an index $k_i \in \{0, 1, \ldots, m\}$ such that $v_i(g_j) = b$ for all $j \leq k_i$, and $v_i(g_j) = a$ for all $j > k_i$. We define agent i's *b-valued prefix*, denoted by $p_b(i)$, as this set of high-valued items: $p_b(i) := \{g_j \in M \mid 1 \leq j \leq k_i\}$.

Algorithm 2 builds upon the algorithmic framework of Feige [31], which computes MMS allocations for bivalued instances. The key difference lies in Line 3: Feige [31] select an arbitrary MMS partition according to agent i, which suffices to establish the existence of MMS allocations for bivalued valuations. To obtain the stronger guarantee of EGMMS, however, an arbitrary MMS partition does not suffice, as we show in Observation A.7. For this reason, we instead select the leximin partition.

We present a partial version of the proof of Theorem 5.1, omitting some details; the full version is deferred to Appendix E.2.

(Partial) Proof of Theorem 5.1. We claim the following is a valid EGMMS certificate for agent i:

$$Y^i = (X_1, X_2, \ldots, X_i, B^i_{i+1}, B^i_{i+2}, \ldots, B^i_n).$$

We need to show that, for every subset $J \subseteq [n]$ with $i \in J$, we have $v_i(X_i) \geq \mu_i(|J|, \mathcal{M}_J)$, where $\mathcal{M}_J = \bigcup_{j \in J} Y^i_j$. Let $j_1 = \min(J)$. We prove the claim by backward induction on j_1, starting from

Algorithm 2: Feige [31]'s algorithm with leximin partitions

Input: An instance with bivalued valuations and identical rankings.
Output: An EGMMS allocation.

$M^{(1)} \leftarrow M$; Index agents s.t. $v_1(M) \geq \cdots \geq v_n(M)$ and items s.t. $v_i(g_1) \geq \cdots \geq v_i(g_m)\ \forall i$;
for $i = 1, \ldots, n$ **do**
  Let $(B_i^i, \ldots, B_n^i)$ be the leximin partition of $M^{(i)}$ under v_i into $(n{-}i{+}1)$ bundles;
  Rearrange bundles so that B_i^i has minimum cardinality;
  Rearrange items so that B_i^i contains the $|B_i^i \cap p_b(i)|$ highest-indexed items of $p_b(i) \cap M^{(i)}$;
  $X_i \leftarrow B_i^i$ and $M^{(i+1)} \leftarrow M^{(i)} \setminus X_i$;
return $X = (X_1, \ldots, X_n)$;

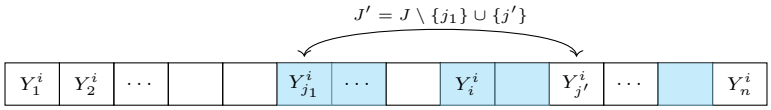


Figure 3: Illustration of the swapping argument used in the proof of Theorem 5.1. The bundles highlighted in blue correspond to the set J. Note that it might be that $j' < i$.

$j_1 = i$ (since $i \in J$) down to $j_1 = 1$. The base case ($j_1 = i$) follows from Proposition 3.2, since $\{Y_j^i \mid j \in J\}$ is a subset of bundles from i's leximin partition selected in Line 3.

For the inductive step, the case $Y_{j_1}^i \cap p_b(i) \neq \emptyset$ is handled in the full proof. In this partial proof, we only consider the case where $Y_{j_1}^i \cap p_b(i) = \emptyset$. In this case, according to v_i, all items in X_{j_1} have value a. If $|X_{j_1}| \leq |\mathcal{M}_J|/|J|$, we can show that $v_i(X_i) \geq \mu_i(|J| - 1, \mathcal{M}_J \setminus X_{j_1}) \geq \mu_i(|J|, \mathcal{M}_J)$, where the first inequality follows from the inductive hypothesis, and the second inequality reflects the fact that removing a bundle of least-valued items of below-average size cannot decrease the MMS share of the remaining agents; see Lemma E.2 for a formal statement.

We now turn to the remaining case, $|X_{j_1}| > |\mathcal{M}_J|/|J|$. Since X_{j_1} is the smallest bundle in a partition of $M^{(j_1)}$ (by Line 4 of Algorithm 2), we have $|X_{j_1}| \leq |M^{(j_1)}|/(n - j_1 + 1)$. Hence $|\mathcal{M}_J|/|J| < |M^{(j_1)}|/(n - j_1 + 1)$, i.e., the average bundle size over J is strictly smaller than the overall average bundle size on $M^{(j_1)}$. It follows that some $j' \in N_{j_1} \setminus J$ satisfies $|Y_{j'}^i| > |X_{j_1}|$.

Let $J' = (J \setminus \{j_1\}) \cup \{j'\}$, and note that $\min(J') > j_1$. See Figure 3 for an illustration. As before, we obtain $v_i(X_i) \geq \mu_i(|J'|, \mathcal{M}_{J'}) \geq \mu_i(|J|, \mathcal{M}_J)$, where the first inequality follows from the inductive hypothesis, and the second inequality holds because any MMS partition of $\mathcal{M}_J$ can be converted into a partition of $\mathcal{M}_{J'}$ by replacing X_{j_1} (which contains only a-valued items) with the strictly larger bundle $Y_{j'}^i$, all of whose items have value at least a. □

Finally, even though computing leximin allocations is NP-hard for general additive valuations [51], Algorithm 2 can be implemented in polynomial time, as a consequence of the following observation; its proof is deferred to Appendix E.2.

Observation 5.2. *There exists a polynomial time algorithm that computes the leximin allocation for any instance with identical bivalued valuations.*

6 EPMMS for Three Agents and for Two Types of Agents

In this section, we establish the existence of EPMMS allocations in two important restricted settings, with proofs deferred to Appendix E.3. We begin with the three-agent case.

Theorem 6.1 (Three Agents)**.** *Every instance with three additive agents admits an EPMMS allocation. More generally, this holds whenever at least two of the agents are PMMS-feasible.*

Remark 6.1. *Before sketching the proof of the theorem, we note that our proof actually establishes an additional guarantee alongside EPMMS. Specifically, if one of the valuations is GMMS-feasible (Definition 2.11), then the resulting allocation is individually MMS-satisfying or PMMS-satisfying*

(IMMP): each agent either (1) receives at least her MMS share, or (2) does not PMMS-envy any other agent. This is a natural generalization of the individually MMS-satisfying or EFX-satisfying (IMMX) property introduced by Babaioff and Frosh [15], obtained by replacing EFX with PMMS. They claimed that IMMX exists for three additive agents, and our analysis strengthens this to IMMP.

We describe the allocation rule used in the proof of Theorem 6.1. Suppose v_2 and v_3 are PMMS-feasible. Let (X_1, X_2, X_3) be agent 1's leximin partition, and assume without loss of generality that X_1 is the least preferred bundle according to agent 2. Let (X_2', X_3') be agent 2's leximin partition of $X_2 \cup X_3$ into two bundles. Agent 3 receives her most preferred bundle among X_1, X_2', X_3'. If X_1 remains, then agent 1 receives X_1 and agent 2 the remaining bundle; otherwise, agent 2 selects her preferred bundle from $\{X_2, X_3\}$ and agent 1 receives the remaining one. Figure 5 illustrates this rule. The full proof shows that this satisfies the desired guarantees.

We next consider the case of two agent types, for which we prove the following theorem.

Theorem 6.2 (Two Types of Agents)**.** *Every instance with two types of additive agents admits an EPMMS allocation. More generally, this holds whenever at least one of the types is PMMS-feasible.*

We describe the allocation rule. Partition the agents into N_α and N_β according to whether they have valuation v_α or v_β, and assume v_β is PMMS-feasible. Compute a leximin partition $(X_1, \ldots, X_n)$ of M with respect to v_α, and relabel the bundles so that $X_1, \ldots, X_{|N_\alpha|}$ are the $|N_\alpha|$ least preferred under v_β. Then compute a leximin partition $(X'_{|N_\alpha|+1}, \ldots, X'_n)$ of $\bigcup_{i>|N_\alpha|} X_i$ into $|N_\beta|$ bundles with respect to v_β. Allocate $X_1, \ldots, X_{|N_\alpha|}$ to the agents in N_α and $X'_{|N_\alpha|+1}, \ldots, X'_n$ to the agents in N_β. Crucially, agents in N_β receive bundles from a leximin partition under v_β and prefer them to every bundle assigned to N_α, while agents in N_α use $(X_1, \ldots, X_n)$ as their EPMMS certificate.

7 Conclusions and Open Problems

In this work, we initiate the study of epistemic PMMS allocations. We establish the existence of EPMMS in several important cases, along with a constant-factor approximation for general additive valuations. Whether exact EPMMS allocations exist for arbitrary additive valuations, even for four agents, remains an open problem. For the special case of bivalued valuations, we obtain the stronger guarantee of epistemic GMMS, which also implies the known existence result for MMS [31]. Two important questions remain open for bivalued valuations: whether PMMS allocations always exist (which would also imply the known existence result for EFX [10, 24, 41]), and the harder question of whether GMMS allocations always exist (which would imply both PMMS and EGMMS).

Acknowledgments and Disclosure of Funding

This project has been partially funded by the European Research Council (ERC) under the European Union's Horizon Europe Program (FACT, grant agreement No. 101170373), by an Amazon Research Award, by the NSF-BSF (grant number 2020788), and by the Israel Science Foundation Breakthrough Program (grant No. 2600/24).

References

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

[2] Hannaneh Akrami and Nidhi Rathi. Epistemic EFX allocations exist for monotone valuations. In *AAAI*, pages 13520–13528. AAAI Press, 2025.

[3] Hannaneh Akrami and Nidhi Rathi. Achieving maximin share and EFX/EF1 guarantees simultaneously. In *AAAI*, pages 13529–13537. AAAI Press, 2025.

[4] Hannaneh Akrami, Noga Alon, Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, and Ruta Mehta. EFX: A simpler approach and an (almost) optimal guarantee via rainbow cycle number. *Oper. Res.*, 73(2):738–751, 2025.

[5] Hannaneh Akrami, Ryoga Mahara, Kurt Mehlhorn, and Nidhi Rathi. Achieving EF1 and epistemic EFX guarantees simultaneously. *CoRR*, abs/2602.11732, 2026.

[6] Hannaneh Akrami, Alexander Mayorov, Kurt Mehlhorn, Shreyas Srinivas, and Christoph Weidenbach. A counterexample to efx; $n \geq 3$ agents, $m \geq n + 5$ items, monotone valuations; via sat-solving. *CoRR*, abs/2604.18216, 2026.

[7] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. Approximation algorithms for computing maximin share allocations. *ACM Trans. Algorithms*, 13(4):52:1–52:28, 2017.

[8] Georgios Amanatidis, Georgios Birmpas, and Vangelis Markakis. Comparing approximate relaxations of envy-freeness. In *IJCAI*, pages 42–48. ijcai.org, 2018.

[9] Georgios Amanatidis, Evangelos Markakis, and Apostolos Ntokos. Multiple birds with one stone: Beating 1/2 for EFX and GMMS via envy cycle elimination. *Theor. Comput. Sci.*, 841: 94–109, 2020.

[10] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A. Voudouris. Maximum nash welfare and other stories about EFX. *Theor. Comput. Sci.*, 863:69–85, 2021.

[11] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.*, 322:103965, 2023.

[12] Arash Ashuri, Vasilis Gkatzelis, and Alkmini Sgouritsa. EF2X exists for four agents. In *AAAI*, pages 13555–13563. AAAI Press, 2025.

[13] Haris Aziz. Pairwise maximin share allocations need not exist, 2026. URL `https://arxiv.org/abs/2609.06282`.

[14] Haris Aziz, Sylvain Bouveret, Ioannis Caragiannis, Ira Giagkousi, and Jérôme Lang. Knowledge, fairness, and social constraints. In *AAAI*, pages 4638–4645. AAAI Press, 2018.

[15] Moshe Babaioff and Gefen Frosh. Near-optimal best-of-both-worlds fairness for few agents. *CoRR*, abs/2602.14668, 2026.

[16] Siddharth Barman and Sanath Kumar Krishnamurthy. Approximation algorithms for maximin fair division, 2020. URL `https://arxiv.org/abs/1703.01851`.

[17] Siddharth Barman and Paritosh Verma. Approximating nash social welfare under binary XOS and binary subadditive valuations. In *WINE*, Lecture Notes in Computer Science, pages 373–390. Springer, 2021.

[18] Siddharth Barman, Arpita Biswas, Sanath Kumar Krishna Murthy, and Yadati Narahari. Group-wise maximin fair allocation of indivisible goods. In *AAAI*, pages 917–924. AAAI Press, 2018.

[19] Nawal Benabbou, Mithun Chakraborty, Edith Elkind, and Yair Zick. Fairness towards groups of agents in the allocation of indivisible items. In *IJCAI*, pages 95–101. ijcai.org, 2019.

[20] Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, and Yair Zick. Finding fair and efficient allocations when valuations don't add up. In *SAGT*, Lecture Notes in Computer Science, pages 32–46. Springer, 2020.

[21] Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. Almost full EFX exists for four agents. In *AAAI*, pages 4826–4833. AAAI Press, 2022.

[22] Sylvain Bouveret and Michel Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Auton. Agents Multi Agent Syst.*, 30(2):259–290, 2016.

[23] Eric Budish. The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. In *BQGT*, page 74:1. ACM, 2010.

[24] Jaroslaw Byrka, Franciszek Malinka, and Tomasz Ponitka. Probing EFX via PMMS: (non-)existence results in discrete fair division. In *AAAI*, pages 16735–16742. AAAI Press, 2026.

[25] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. *ACM Trans. Economics and Comput.*, 7(3):12:1–12:32, 2019.

[26] Ioannis Caragiannis, Jugal Garg, Nidhi Rathi, Eklavya Sharma, and Giovanna Varricchio. New fairness concepts for allocating indivisible items. In *IJCAI*, pages 2554–2562. ijcai.org, 2023.

[27] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. A little charity guarantees almost envy-freeness. *SIAM J. Comput.*, 50(4):1336–1358, 2021.

[28] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. EFX exists for three agents. *J. ACM*, 71(1):4:1–4:27, 2024.

[29] George Christodoulou and Symeon Mastrakoulis. Exact and approximate maximin share allocations in multi-graphs. In *AAAI*, pages 16761–16769. AAAI Press, 2026.

[30] Sijia Dai, Xinru Guo, Huahua Miao, Guichen Gao, Yicheng Xu, and Yong Zhang. The existence and efficiency of PMMS allocations. *Theor. Comput. Sci.*, 989:114388, 2024.

[31] Uriel Feige. Maximin fair allocations with two item values. *Preprint*, 2022. Version : February 22, 2022. Available at `https://www.wisdom.weizmann.ac.il/~feige/mypapers/MMSab.pdf`.

[32] Uriel Feige, Ariel Sapir, and Laliv Tauber. A tight negative example for MMS fair allocations. In *WINE*, Lecture Notes in Computer Science, pages 355–372. Springer, 2021.

[33] Michal Feldman, Simon Mauras, and Tomasz Ponitka. On optimal tradeoffs between EFX and nash welfare. In *AAAI*, pages 9688–9695. AAAI Press, 2024.

[34] Jugal Garg and Eklavya Sharma. Exploring relations among fairness notions in discrete fair division. *CoRR*, abs/2502.02815, 2025.

[35] Pratik Ghosal, Vishwa Prakash HV, Prajakta Nimbhorkar, and Nithin Varma. (almost full) EFX for three (and more) types of agents. In *AAAI*, pages 13889–13896. AAAI Press, 2025.

[36] Paul Gölz. Pmms allocations need not exist for 3 agents with additive valuations, 2026. URL `https://arxiv.org/abs/2609.08954`.

[37] Hadi Hosseini, Andrew Searns, and Erel Segal-Halevi. Ordinal maximin share approximation for goods. *J. Artif. Intell. Res.*, 74, 2022.

[38] Halvard Hummel. Maximin shares in hereditary set systems. *ACM Trans. Economics and Comput.*, 13(3):12:1–12:33, 2025.

[39] Halvard Hummel and Magnus Lie Hetland. Maximin shares under cardinality constraints. *Auton. Agents Multi Agent Syst.*, 40(1):9, 2026.

[40] Vishwa Prakash HV, Pratik Ghosal, Prajakta Nimbhorkar, and Nithin Varma. EFX exists for three types of agents. In *EC*, pages 101–128. ACM, 2025.

[41] Jiarong Jin and Biaoshuai Tao. On pareto-optimal and fair allocations with personalized bi-valued utilities. *CoRR*, abs/2507.18251, 2025.

[42] H. W. Kuhn. On games of fair division. In Martin Shubik, editor, *Essays in Mathematical Economics in Honor of Oskar Morgenstern*, pages 29–37. Princeton University Press, Princeton, NJ, 1967.

[43] David Kurokawa. *Fair Division in Game Theoretic Settings*. PhD thesis, Carnegie Mellon University, 2017.

[44] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. When can the maximin share guarantee be guaranteed? In *AAAI*, pages 523–529. AAAI Press, 2016.

[45] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. *J. ACM*, 65(2):8:1–8:27, 2018.

[46] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games Econ. Behav.*, 55(2):270–296, 2006.

[47] Renato Paes Leme. Gross substitutability: An algorithmic survey. *Games Econ. Behav.*, 106: 294–316, 2017.

[48] Jing Wu Lian, Nicholas Mattei, Renee Noble, and Toby Walsh. The conference paper assignment problem: Using order weighted averages to assign indivisible goods. In *AAAI*, pages 1138–1145. AAAI Press, 2018.

[49] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *EC*, pages 125–131. ACM, 2004.

[50] Ryoga Mahara. Existence of EFX for two additive valuations. *Discret. Appl. Math.*, 340: 115–122, 2023.

[51] Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. *SIAM J. Discret. Math.*, 34(2):1039–1068, 2020.

[52] John Rawls. *A Theory of Justice*. Harvard University Press, Cambridge, MA, 1971.

[53] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.

A Separations between EPMMS and EEFX

A.1 PMMS Is Not Preserved Under Envy-Cycle Elimination

We begin by highlighting a fundamental difference between PMMS and EFX. The central operation of the Envy-Cycle Elimination algorithm (Algorithm 1) is the envy-cycle elimination step (Line 8): we identify a cycle of agents in which each agent prefers the next agent's bundle, and we rotate the bundles along the cycle. It is well known that this step preserves the EFX guarantee, since the set of bundles is unchanged and each agent can only increase the value of her own bundle [51].

In contrast, PMMS is not preserved under envy-cycle elimination. We construct an instance in which a PMMS allocation, after a single application of the procedure, becomes only $\frac{2}{3}$-PMMS.

Proposition A.1. *For every $0 < \alpha \leq 1$, there exists an instance with additive valuations and an α-PMMS allocation such that a single envy-cycle elimination step (i.e., one execution of Line 8 in Algorithm 1) produces an allocation that fails to be $\frac{2\alpha+\epsilon}{2+\alpha}$-PMMS for any $\epsilon > 0$.*

Proof. Consider an instance with $n = 3$ agents and six goods, where agent 1 has the additive valuation

$$v_1(g_1) = 2\alpha, \quad v_1(g_2) = 2, \quad v_1(g_3) = 2, \quad v_1(g_4) = \alpha, \quad v_1(g_5) = \alpha, \quad v_1(g_6) = \epsilon/2,$$

and $\epsilon > 0$ is arbitrarily small. Consider the initial allocation X given by

$$X_1 = \{g_1\}, \quad X_2 = \{g_2, g_3\}, \quad X_3 = \{g_4, g_5, g_6\}.$$

We first verify that X satisfies the α-PMMS condition for agent 1, whose current value is $v_1(X_1) = 2\alpha$. Comparing her bundle against each of the other agents', we have

$$\begin{aligned} \alpha \cdot \mu_1(2, X_1 \cup X_2) &= \alpha \cdot \mu_1(2, \{g_1, g_2, g_3\}) = \alpha \cdot \min\{2,\, 2 + 2\alpha\} = 2\alpha = v_1(X_1), \\ \mu_1(2, X_1 \cup X_3) &= \mu_1(2, \{g_1, g_4, g_5, g_6\}) = \min\{2\alpha,\, 2\alpha + \epsilon/2\} = 2\alpha = v_1(X_1). \end{aligned}$$

Hence agent 1's α-PMMS condition is satisfied.

Now suppose an envy cycle is resolved (e.g., $1 \to 3 \to 1$), and as a result agent 1 receives X_3 in place of X_1. Her value increases to $v_1(X_3) = 2\alpha + \epsilon/2$, and the maximin share again agent 2 becomes

$$\mu_1(2, X_2 \cup X_3) = \mu_1(2, \{g_2, g_3, g_4, g_5, g_6\}) = \min\{2 + \alpha,\, 2 + \alpha + \epsilon/2\} = 2 + \alpha.$$

The resulting ratio is

$$\frac{v_1(X_3)}{\mu_1(2, X_2 \cup X_3)} = \frac{2\alpha + \epsilon/2}{2+\alpha} < \frac{2\alpha+\epsilon}{2+\alpha},$$

so agent 1's bundle violates the $\frac{2\alpha+\epsilon}{2+\alpha}$-PMMS guarantee after the envy-cycle elimination step. □

We next show that this bound is tight: eliminating an envy cycle in a PMMS allocation always results in an allocation that is $\frac{2}{3}$-PMMS.

Proposition A.2. *Under additive valuations and for $0 < \alpha \leq 1$, a single envy-cycle elimination step (i.e., one execution of Line 8 in Algorithm 1) applied to an α-PMMS allocation produces a $\frac{2\alpha}{2+\alpha}$-PMMS allocation.*

Proof. Let $X = (X_1, \ldots, X_n)$ be the allocation immediately before Line 8 of Algorithm 1 is executed, and let $X' = (X'_1, \ldots, X'_n)$ denote the allocation after its execution. Fix an agent i. Since X is α-PMMS, we have $v_i(X_i) \geq \alpha \cdot \mu_i(2, X_i \cup X_j)$ for all j. The envy-cycle elimination step only permutes bundles among agents, so the set of bundles is unchanged; hence $v_i(X_i) \geq \alpha \cdot \mu_i(2, X_i \cup X'_j)$ for all j as well. Our goal is to show that

$$v_i(X'_i) \geq \frac{2\alpha}{2+\alpha} \cdot \mu_i(2, X'_i \cup X'_j) \quad \text{for all } j. \tag{1}$$

Fix an agent j, set $\mu' = \mu_i(2, X'_i \cup X'_j)$, and let (A, B) be a partition of $X'_i \cup X'_j$ such that $v_i(A) \geq \mu'$ and $v_i(B) \geq \mu'$. Since $v_i(X'_i) = v_i(A \cap X'_i) + v_i(B \cap X'_i)$, we may assume without loss of generality that $v_i(B \cap X'_i) \leq \frac{1}{2} v_i(X'_i)$. Observe that

$$(B \setminus X'_i) \cup X_i \cup (X'_j \setminus B) = X_i \cup X'_j,$$

so $(B \setminus X'_i, X_i \cup (X'_j \setminus B))$ is a partition of $X_i \cup X'_j$. It follows that

$$\begin{aligned}\mu_i(2, X_i \cup X'_j) &\geq \min\big\{v_i(B \setminus X'_i),\ v_i(X_i \cup (X'_j \setminus B))\big\} \\ &= \min\big\{v_i(B) - v_i(B \cap X'_i),\ v_i(X_i) + v_i(A) - v_i(A \cap X'_i)\big\}.\end{aligned}$$

We distinguish two cases according to which term attains this minimum.

Case 1: $\mu_i(2, X_i \cup X'_j) \geq v_i(B) - v_i(B \cap X'_i)$. We have:

$$\begin{aligned} v_i(X'_i) &\geq v_i(X_i) && \text{(by the design of the algorithm)} \\ &\geq \alpha \mu_i(2, X_i \cup X'_j) && \text{(since } X \text{ is } \alpha\text{-PMMS)} \\ &\geq \alpha \left(v_i(B) - v_i(B \cap X'_i)\right) && \text{(by the assumption of Case 1)} \\ &\geq \alpha \left(\mu' - \tfrac{1}{2} v_i(X'_i)\right) && \text{(since } v_i(B) \geq \mu' \text{ and } v_i(B \cap X'_i) \leq \tfrac{1}{2} v_i(X'_i)) \end{aligned}$$

Rearranging gives

$$\frac{1}{\alpha} v_i(X'_i) + \frac{1}{2} v_i(X'_i) \geq \mu' \qquad \Longrightarrow \qquad v_i(X'_i) \geq \frac{2\alpha}{2+\alpha} \mu'$$

establishing Equation (1).

Case 2: $\mu_i(2, X_i \cup X'_j) \geq v_i(X_i) + (v_i(A) - v_i(A \cap X'_i))$. We have:

$$\begin{aligned} v_i(X_i) &\geq \alpha \mu_i(2, X_i \cup X'_j) && \text{(since } X \text{ is } \alpha\text{-PMMS)} \\ &\geq \alpha \left(v_i(X_i) + (v_i(A) - v_i(A \cap X'_i))\right) && \text{(by the assumption of Case 2)} \\ &\geq \alpha \left(v_i(X_i) + \mu' - v_i(X'_i)\right) && \text{(since } v_i(A) \geq \mu' \text{ and } v_i(A \cap X'_i) \leq v_i(X'_i)) \end{aligned}$$

Since $v_i(X'_i) \geq v_i(X_i)$ by the design of the algorithm, rearranging gives

$$(1-\alpha) v_i(X'_i) \geq \alpha(\mu' - v_i(X'_i)) \qquad \Longrightarrow \qquad v_i(X'_i) \geq \alpha \mu'.$$

Since $\alpha \geq \frac{2\alpha}{2+\alpha}$, this implies Equation (1), completing the proof. □

A.2 The Lone Divider Approach Does Not Guarantee EPMMS

The lone divider approach [42, 53] is a common technique for achieving share-based fairness notions, including approximations of the maximin share [1, 3, 37–39]. The algorithm proceeds in rounds. In each round, a designated "divider" agent partitions the remaining items into n' bundles—where n' is the number of remaining agents—such that each bundle meets her own fair share threshold. Next, a bipartite graph is constructed between the remaining agents and these bundles, with an edge connecting an agent to a bundle whenever the bundle meets that agent's fair share threshold. A (partial or perfect) matching is then computed in this graph, chosen so that no unmatched agent has an edge to a matched bundle. Finally, each matched agent is assigned the bundle to which she is paired, and these agents and bundles are removed before proceeding to the next round.

To establish the existence of EEFX allocations for general monotone valuations, Akrami and Rathi [2] elegantly adapt this approach by defining desirability through the EEFX property rather than a numerical share. Specifically, for any partition of the items, an *EEFX-graph* is constructed in which an edge exists between an agent and a bundle if and only if the bundle is EEFX-feasible for that agent; see the following definition.

Definition A.3 (EEFX-Feasibility, [2])**.** *Let k be a positive integer, $i \in [n]$ an agent, and $S \subseteq M$ a subset of items. A bundle $A \subseteq S$ is* k-EEFX-feasible *for agent i with respect to S if there exists a partition of $S \setminus A$ into $k-1$ bundles $C_1, C_2, \ldots, C_{k-1}$ such that*

$$v_i(A) \geq v_i(C_j \setminus \{g\}) \quad \textit{for every } j \in [k-1] \textit{ and every } g \in C_j.$$

We then define

$$\mathrm{EEFX}_i^k(S) = \{A \subseteq S \mid A \textit{ is } k\textit{-EEFX-feasible for agent } i \textit{ with respect to } S\}.$$

The lone divider approach succeeds in this setting because of the following structural property: if a bundle is EEFX-feasible for an agent once some items have been removed, then it was already EEFX-feasible before their removal. Consequently, after each round it suffices to verify the EEFX condition on the reduced instance, without revisiting bundles assigned in earlier rounds. The following lemma of Akrami and Rathi [2] formalizes this property.

Lemma A.4 ([2, Lemma 3.1])**.** *Consider any fair division instance, an agent $i \in [n]$, and a bundle $A \subseteq M$ with $A \notin \mathrm{EEFX}_i^n(M)$. Then every $B \in \mathrm{EEFX}_i^{n-1}(M \setminus A)$ satisfies $B \in \mathrm{EEFX}_i^n(M)$.*

While this property enables the modified lone divider method for EEFX, we now show that it fails when EEFX is replaced with EPMMS, ruling out a direct adaptation of the same argument. We begin with the analogous notion of EPMMS-feasibility.

Definition A.5 (EPMMS-Feasibility)**.** *Let k be a positive integer, $i \in [n]$ an agent, and $S \subseteq M$ a subset of items. A bundle $A \subseteq S$ is* k-EPMMS-feasible *for agent i with respect to S if there exists a partition of $S \setminus A$ into $k-1$ bundles $C_1, C_2, \ldots, C_{k-1}$ such that*

$$v_i(A) \geq \mu_i(2, A \cup C_j) \quad \textit{for every } j \in [k-1].$$

We then define

$$\mathrm{EPMMS}_i^k(S) = \{A \subseteq S \mid A \textit{ is } k\textit{-EPMMS-feasible for agent } i \textit{ with respect to } S\}.$$

The following lemma shows that the analogue of Lemma A.4 fails for EPMMS.

Lemma A.6. *There exists a fair division instance, an agent $i \in [n]$, and a bundle $A \subseteq M$ with $A \notin \mathrm{EPMMS}_i^n(M)$ such that some $B \in \mathrm{EPMMS}_i^{n-1}(M \setminus A)$ satisfies $B \notin \mathrm{EPMMS}_i^n(M)$.*

Proof. Consider an instance with $n = 3$ agents and six items, and let i be an agent with an additive valuation assigning values $1, 4, 4, 6, 6, 6$ to the six items. Throughout the proof, we represent bundles by the multisets of their item values, and we drop the subscript i since only one agent is involved.

Consider the partition

$$X = (X_1, X_2, X_3) = (\{6\}, \{4, 4, 1\}, \{6, 6\}).$$

We first show that $X_2 \notin \mathrm{EPMMS}^3(M)$. Suppose for contradiction that some partition (C_1, C_2) of $M \setminus X_2 = \{6, 6, 6\}$ into two bundles satisfies $v(X_2) \geq \mu(2, X_2 \cup C_j)$ for both $j \in \{1, 2\}$. Without loss of generality, C_1 contains at least two items of value 6, so

$$\mu(2, X_2 \cup C_1) \geq \mu(2, \{1, 4, 4, 6, 6\}) = \min\{6 + 4 + 1,\ 6 + 4\} = 10 > 9 = v(X_2),$$

a contradiction. Hence $X_2 \notin \text{EPMMS}^3(M)$.

Now consider $B = \{6\}$ in the reduced instance $M \setminus X_2 = \{6, 6, 6\}$. The partition $(\{6\}, \{6, 6\})$ certifies that $B \in \text{EPMMS}^2(M \setminus X_2)$, since

$$v(B) = 6 \geq \mu(2, \{6, 6, 6\}) = \min\{6 + 6,\ 6\} = 6.$$

We now argue that $B \notin \text{EPMMS}^3(M)$. Suppose for contradiction that some partition (D_1, D_2) of $M \setminus B = \{1, 4, 4, 6, 6\}$ into two bundles satisfies $v(B) \geq \mu(2, B \cup D_j)$ for both $j \in \{1, 2\}$. Since $|M \setminus B| = 5$, at least one of D_1, D_2 contains three or more items; without loss of generality, assume $|D_1| \geq 3$. Then moving the lowest-valued item of D_1 into B yields two bundles each of value strictly greater than $v(B) = 6$, so $\mu(2, B \cup D_1) > v(B)$, a contradiction. Hence $B \notin \text{EPMMS}^3(M)$, completing the proof. □

A.3 EPMMS Is Incomparable to Both MMS and EFX

In this section, we show that EPMMS neither implies nor is implied by MMS or EFX. This stands in contrast to EEFX, which is implied by both MMS and EFX [26]. We first show that EPMMS is incomparable to MMS.

Observation A.7. *Neither EPMMS nor MMS implies the other, even when agents have identical bivalued valuations (a special case of additive valuations).*

Proof. (EPMMS $\not\Rightarrow$ MMS) Consider an instance with 6 items and agents having identical bivalued valuations that assign value 1 to three items and value 2 to the remaining three. Consider the allocation

$$X = (X_1, X_2, X_3) = (\{2\}, \{2, 2\}, \{1, 1, 1\}),$$

where each bundle is represented by the values of the items it contains. Note that the 3-maximin share is

$$\mu(3, M) = 3 = \min\{2 + 1,\ 2 + 1,\ 2 + 1\}.$$

Hence X is not an MMS allocation, as agent 1 receives value $2 < 3$.

We claim, however, that X is PMMS, and therefore EPMMS. Since agent 2 does not envy agents 1 and 3, and agent 3 does not envy agent 1, the PMMS condition is automatically satisfied for these pairs (Proposition C.5). For the remaining pairs, we verify directly:

$$\begin{aligned} v(X_1) &= 2 = \mu(2, X_1 \cup X_2) = \mu(2, \{2, 2, 2\}), \\ v(X_1) &= 2 = \mu(2, X_1 \cup X_3) = \mu(2, \{2, 1, 1, 1\}), \\ v(X_3) &= 3 = \mu(2, X_2 \cup X_3) = \mu(2, \{2, 2, 1, 1, 1\}). \end{aligned}$$

This completes the proof of the first part of the observation.

(MMS $\not\Rightarrow$ EPMMS). Consider an instance with 9 items and agents having identical bivalued valuations that assign value 5 to two items and value 3 to the remaining seven. Consider the allocation

$$X = (X_1, X_2, X_3) = (\{3, 3, 3\}, \{5, 5\}, \{3, 3, 3, 3\}),$$

where bundles are represented by the values of the items they contain. Observe that the 3-maximin share equals

$$\begin{aligned} \mu(3, M) = 9 &= \min\{3 + 3 + 3,\ 5 + 5,\ 3 + 3 + 3 + 3\} \\ &= \min\{3 + 3 + 3,\ 5 + 3 + 3,\ 5 + 3 + 3\}. \end{aligned}$$

Therefore, X is an MMS allocation, since each agent receives value at least 9.

However, X is not an EPMMS allocation. Suppose, for the sake of contradiction, that there exists an allocation $X' = (X'_1, X'_2, X'_3)$ such that $X'_2 = X_2$ and agent 2 does not PMMS-envy any other agent under X'. Notice that either X'_1 or X'_3 must contain at least four items of value 3. Without loss of generality, suppose this holds for X'_1. Then

$$\mu(2, X'_1 \cup X'_2) \geq \mu(2, \{3, 3, 3, 3\} \cup \{5, 5\}) = 11 = \min\{5 + 3 + 3,\ 5 + 3 + 3\}.$$

Hence agent 2 PMMS-envies agent 1, a contradiction. We conclude that MMS does not imply EPMMS. □

We now show that EPMMS is incomparable to EFX.

Observation A.8. *Neither EPMMS nor EFX implies the other, even when agents have identical bivalued valuations (a special case of additive valuations).*

Proof. (EPMMS $\not\Rightarrow$ EFX) Consider an instance with 2 items and 3 agents, where every agent values both items at 1. Consider the allocation

$$X = (X_1, X_2, X_3) = (\{1, 1\}, \emptyset, \emptyset),$$

where bundles are represented by the values of the items they contain. This allocation is not EFX, since agents 2 and 3 envy agent 1 even after removing any single good. However, X is EPMMS: both agents 2 and 3 could hypothetically repartition the two items assigned to agent 1, splitting them equally between the remaining agents to obtain a PMMS allocation.

(EFX $\not\Rightarrow$ EPMMS) Consider an instance with 5 items and two agents having identical additive valuations that assign value 3 to two items and value 2 to the remaining three. Consider the allocation

$$X = (X_1, X_2) = (\{3, 2, 2\}, \{3, 2\}),$$

where bundles are represented by the values of the items they contain. This allocation is EFX, since agent 2 does not envy agent 1 after the removal of any single item. However, X is not EPMMS: in the two-agent case, EPMMS coincides with PMMS, and

$$\mu(2, \{3, 2, 2\} \cup \{3, 2\}) = \min\{3 + 3,\ 2 + 2 + 2\} = 6,$$

which is strictly larger than $3 + 2 = 5$. □

B Separation between EGMMS and EPMMS

In this section, we show that although EGMMS directly implies both EPMMS and MMS, the converse does not hold. In fact, we prove a stronger separation: even simultaneously satisfying MMS and PMMS does not imply EGMMS.

Proposition B.1. *There exists an instance with identical additive valuations admitting an allocation that is both MMS and PMMS, yet* not *EGMMS.*

Proof. (PMMS $\wedge$ MMS $\not\Rightarrow$ EGMMS) Consider four agents with identical additive valuations over eight items of values

$$5,\ 4,\ 4,\ 3,\ 2,\ 2,\ 2,\ 1,$$

and the allocation

$$X_1 = \{1, 4\}, \qquad X_2 = \{5\}, \qquad X_3 = \{2, 2, 2\}, \qquad X_4 = \{3, 4\},$$

where each bundle is written as the multiset of item values.

We first argue that the allocation is MMS. The total value is 23, so the MMS value is at most $\lfloor 23/4 \rfloor = 5$. Every agent receives value at least 5, so the allocation is MMS.

Next, we argue that the allocation is PMMS. Since the valuations are additive, for any pair i, j with $v(X_i) \geq v(X_j)$ the PMMS condition for agent i holds (Proposition C.5). It therefore suffices, for each pair, to verify the condition for the agent with the smaller bundle:

$$\begin{aligned}
&\mu(2,\ X_1 \cup X_3) = \mu(2, \{1, 2, 2, 2, 4\}) \leq \lfloor 11/2 \rfloor = 5 = v(X_1),\\
&\mu(2,\ X_1 \cup X_4) = \mu(2, \{1, 3, 4, 4\}) = \min\{4 + 1,\ 4 + 3\} = 5 = v(X_1),\\
&\mu(2,\ X_2 \cup X_3) = \mu(2, \{2, 2, 2, 5\}) \leq \lfloor 11/2 \rfloor = 5 = v(X_2),\\
&\mu(2,\ X_2 \cup X_4) = \mu(2, \{3, 4, 5\}) = \min\{3 + 4,\ 5\} = 5 = v(X_2),\\
&\mu(2,\ X_3 \cup X_4) = \mu(2, \{2, 2, 2, 3, 4\}) \leq \lfloor 13/2 \rfloor = 6 = v(X_3).
\end{aligned}$$

Hence the allocation is PMMS.

We now argue that the allocation is not EGMMS. We show that agent 1 admits no EGMMS-certificate. Suppose, for contradiction, that

$$Y = (Y_1, Y_2, Y_3, Y_4), \qquad Y_1 = X_1 = \{1, 4\},$$

is an EGMMS-certificate for agent 1; that is, Y is a partition of all items with $Y_1 = X_1$, and for every subgroup S with $1 \in S$, it holds that

$$v(X_1) \;\geq\; \mu\big(|S|,\, \bigcup_{j\in S} Y_j\big).$$

We claim that the pairwise (i.e., $|S| = 2$) conditions force the partition of the remaining items $\{2, 2, 2, 3, 4, 5\}$ uniquely (up to relabeling of Y_2, Y_3, Y_4).

Assume without loss of generality that $5 \in Y_2$. We claim $Y_2 = \{5\}$. If instead some item of value $a \neq 5$ also belongs to Y_2, then $a \geq 2$, and the partition $(\{4, a\}, \{1, 5\})$ of $Y_1 \cup Y_2$ shows

$$\mu(2,\, Y_1 \cup Y_2) \;\geq\; \min\{4 + a,\, 5 + 1\} \;\geq\; 6 \;>\; 5 \;=\; v(X_1),$$

a contradiction. Hence $Y_2 = \{5\}$.

Next, assume without loss of generality that $3 \in Y_3$. We claim that $Y_3 = \{3, 4\}$.

Suppose first that $|Y_3| \geq 3$. Then Y_3 contains 3 and two further items a, b, each of value at least 2. The partition $(\{4, a\}, \{1, 3, b\})$ of $Y_1 \cup Y_3$ gives

$$\mu(2,\, Y_1 \cup Y_3) \;\geq\; \min\{4 + a,\, 1 + 3 + b\} \;\geq\; 6 \;>\; v(X_1),$$

a contradiction. Suppose instead that $|Y_3| = 1$. Then $Y_3 = \{3\}$ and $Y_4 = \{2, 2, 2, 4\}$, so the partition $(\{4, 2, 1\}, \{2, 2, 4\})$ of $Y_1 \cup Y_4$ gives

$$\mu(2,\, Y_1 \cup Y_4) \;\geq\; \min\{4 + 2 + 1,\, 2 + 2 + 4\} \;=\; 7 \;>\; v(X_1),$$

a contradiction. Finally, suppose $Y_3 = \{3, 2\}$. Then $Y_4 = \{2, 2, 4\}$, and the partition $(\{4, 2, 1\}, \{2, 4\})$ of $Y_1 \cup Y_4$ gives

$$\mu(2,\, Y_1 \cup Y_4) \;\geq\; \min\{4 + 2 + 1,\, 2 + 4\} \;=\; 6 \;>\; v(X_1),$$

a contradiction.

This confirms that the unique certificate (up to permutation) is

$$Y_2 = \{5\}, \qquad Y_3 = \{3, 4\}, \qquad Y_4 = \{2, 2, 2\}.$$

Now consider the subgroup $S = \{1, 3, 4\}$. The corresponding union is

$$Y_1 \cup Y_3 \cup Y_4 \;=\; \{1, 2, 2, 2, 3, 4, 4\},$$

which can be partitioned into three bundles of value 6 each:

$$\{2, 4\}, \qquad \{2, 4\}, \qquad \{1, 2, 3\}.$$

Therefore $\mu(3,\, Y_1 \cup Y_3 \cup Y_4) \geq 6 > 5 = v(X_1)$, violating the groupwise requirement for S.

Since the unique pairwise-feasible certificate fails the groupwise condition, no EGMMS-certificate for agent 1 exists, and the allocation is not EGMMS. □

C Relations Between the Valuation Classes

C.1 GMMS-Feasible Valuations Are A Strict Subclass of PMMS-feasible Valuations

We prove that PMMS-feasible valuations form a strict subclass of GMMS-feasible valuations. That every GMMS-feasible valuation is also PMMS-feasible follows directly from the definitions. The converse, however, does not hold, as we demonstrate below.

Proposition C.1. *There exists a monotone PMMS-feasible valuation function that is not GMMS-feasible.*

Proof. (PMMS-feasibility $\not\Rightarrow$ GMMS-feasibility) Consider a set of 9 items. Define a valuation function v such that $v(S) = 0$ for all subsets S with $|S| \leq 2$ and $v(S) = 1$ for all subsets with $|S| \geq 4$. For subsets of size 3, the value is defined as follows:

$$v(S) = 0 \text{ for } S \in \{\{1, 2, 5\}, \{1, 2, 9\}, \{1, 3, 4\}, \{1, 3, 5\}, \{1, 3, 6\}, \{1, 3, 7\}, \{1, 3, 8\},$$

$$\{1,3,9\},\{1,4,5\},\{1,4,7\},\{1,4,8\},\{1,4,9\},\{1,5,6\},\{1,5,7\},$$
$$\{1,5,8\},\{1,5,9\},\{1,6,9\},\{1,7,8\},\{1,7,9\},\{1,8,9\},\{2,3,5\},$$
$$\{2,5,8\},\{2,5,9\},\{3,4,5\},\{3,4,9\},\{3,5,6\},\{3,5,7\},\{3,5,8\},$$
$$\{3,5,9\},\{3,6,9\},\{3,7,9\},\{3,8,9\},\{4,5,7\},\{4,5,8\},\{4,5,9\},$$
$$\{4,8,9\},\{5,6,8\},\{5,6,9\},\{5,7,8\},\{5,7,9\},\{5,8,9\}\}$$

and $v(S) = 1$ for all other subsets S of size 3. This valuation is not GMMS-feasible, since:

$$\min\{v(\{1,2,3\}), v(\{4,5,6\}), v(\{7,8,9\})\} = 1,$$
$$\max\{v(\{1,4,7\}), v(\{2,5,8\}), v(\{3,6,9\})\} = 0$$

It can, however, be verified that v is PMMS-feasible. To check this, we verify that there is no pair of partitions (X, Y) and (A, B) satisfying $X \cap Y = \emptyset$, $A \cap B = \emptyset$, $X \cup Y = A \cup B$, and $\min\{v(X), v(Y)\} > \max\{v(A), v(B)\}$. The critical case to inspect is when all sets have size 3, i.e. $|X| = |Y| = |A| = |B| = 3$, because for $\min\{v(X), v(Y)\} > \max\{v(A), v(B)\}$ to hold, we must have $v(X) = v(Y) = 1$ (implying $|X|, |Y| \geq 3$) and $v(A) = v(B) = 0$ (implying $|A|, |B| \leq 3$). The only possible configuration satisfying these size constraints is when all sets have size 3, and these cases can be verified computationally. There are $\binom{9}{6} = 84$ subsets of size 6 in a 9-element set. Each such subset contains $\binom{6}{3} = 20$ subsets of size 3, which form $\binom{20}{2} = 190$ distinct pairs to verify. Hence, there are in total $84 \times 190 = 15{,}960$ cases to check, which we have verified computationally. □

C.2 Nice Cancelable Valuations Are A Strict Subclass of GMMS-Feasible Valuations

We prove that GMMS-feasible valuations form a strict superset of nice cancelable valuations, a class known to include additive, budget-additive, unit-demand, and multiplicative valuations [21]. We begin by defining the class of nice cancelable valuations.

Definition C.2 (Nice Cancelable Valuations, [21])**.** *A valuation $v_i : 2^M \to \mathbb{R}_{\geq 0}$ is* cancelable *if $v_i(S \cup \{g\}) > v_i(T \cup \{g\})$ implies $v_i(S) > v_i(T)$ for all $S, T \subseteq M$ and every $g \in M \setminus (S \cup T)$. A valuation v_i is* nice cancelable *if it is cancelable and there exists a cancelable valuation $v_i' : 2^M \to \mathbb{R}_{\geq 0}$ that is* non-degenerate*, meaning $v_i'(S) \neq v_i'(T)$ for all distinct $S, T \subseteq M$, and* respects v_i*, meaning $v_i(S) > v_i(T)$ implies $v_i'(S) > v_i'(T)$ for all $S, T \subseteq M$.*

Next, we show that nice-cancelability implies GMMS-feasibility.

Proposition C.3. *Every nice cancelable valuation is GMMS-feasible.*

Proof. (Nice-cancelability ⇒ GMMS-feasibility) Let v be a nice cancelable valuation, and let v' be a non-degenerate cancelable valuation that respects v. Fix any $S \subseteq M$, any $k \geq 2$, and any two partitions $A = (A_1, \ldots, A_k)$ and $B = (B_1, \ldots, B_k)$ of S. Suppose, for contradiction, that

$$\min\{v(A_1), \ldots, v(A_k)\} > \max\{v(B_1), \ldots, v(B_k)\}.$$

Since v' respects v, this implies

$$\min\{v'(A_1), \ldots, v'(A_k)\} > \max\{v'(B_1), \ldots, v'(B_k)\}. \tag{2}$$

We show by induction on ℓ that $v'(A_1 \cup \cdots \cup A_\ell) \geq v'(B_1 \cup \cdots \cup B_\ell)$ for all $1 \leq \ell \leq k-1$. The base case $\ell = 1$ is immediate from Equation (2). Assume the claim holds for some $1 \leq \ell < k-1$. By Equation (2), $v'(A_{\ell+1}) > v'(B_{\ell+1})$. Let $X = A_{\ell+1} \setminus (B_1 \cup \cdots \cup B_\ell)$ and $Y = (B_1 \cup \cdots \cup B_\ell) \setminus A_{\ell+1}$. Applying the contrapositive of cancelability of v' twice yields

$$v'(A_1 \cup \cdots \cup A_\ell) \geq v'(B_1 \cup \cdots \cup B_\ell) \implies v'(A_1 \cup \cdots \cup A_\ell \cup X) \geq v'(B_1 \cup \cdots \cup B_\ell \cup X),$$
$$v'(A_{\ell+1}) \geq v'(B_{\ell+1}) \implies v'(Y \cup A_{\ell+1}) \geq v'(Y \cup B_{\ell+1}).$$

Chaining these inequalities gives $v'(A_1 \cup \cdots \cup A_\ell \cup X) \geq v'(Y \cup B_{\ell+1})$. Letting $Z = A_{\ell+1} \cap (B_1 \cup \cdots \cup B_\ell)$, and applying the contrapositive of cancelability of v' once more yields

$$v'(A_1 \cup \cdots \cup A_{\ell+1}) = v'(A_1 \cup \cdots \cup A_\ell \cup X \cup Z) \geq v'(Y \cup B_{\ell+1} \cup Z) = v'(B_1 \cup \cdots \cup B_{\ell+1})$$

completing the induction.

Taking $\ell = k-1$, we obtain $v'(A_1 \cup \cdots \cup A_{k-1}) \geq v'(B_1 \cup \cdots \cup B_{k-1})$. Since A and B are partitions of the same set S, we have $A_1 \cup \cdots \cup A_{k-1} = S \setminus A_k$ and $B_1 \cup \cdots \cup B_{k-1} = S \setminus B_k$. Because $A_k \neq B_k$ (since, by Equation (2), $v'(A_k) > v'(B_k)$), these two sets are distinct, and non-degeneracy of v' implies strict inequality

$$v'(S \setminus A_k) > v'(S \setminus B_k). \tag{3}$$

Writing $S \setminus A_k = (B_k \setminus A_k) \cup (S \setminus (A_k \cup B_k))$ and $S \setminus B_k = (A_k \setminus B_k) \cup (S \setminus (A_k \cup B_k))$, and applying cancelability of v' to Equation (3) to remove the common elements of $S \setminus (A_k \cup B_k)$ one at a time, we obtain

$$v'(B_k \setminus A_k) > v'(A_k \setminus B_k).$$

Applying cancelability once more to add the common elements of $A_k \cap B_k$ one at a time gives

$$v'(B_k) = v'((B_k \setminus A_k) \cup (A_k \cap B_k)) \geq v'((A_k \setminus B_k) \cup (A_k \cap B_k)) = v'(A_k).$$

As before, $A_k \neq B_k$, so non-degeneracy of v' implies strict inequality $v'(B_k) > v'(A_k)$. Since v' respects v, this gives $v(B_k) \geq v(A_k)$, contradicting Equation (2). □

Finally, we show that GMMS-feasibility does not necessarily imply nice-cancelability.

Proposition C.4. *Not every GMMS-feasible valuation is nice cancelable.*

Proof. (GMMS-feasibility $\not\Rightarrow$ nice-cancelability) Akrami et al. [4, Example 1] provide an example of a set of items $M = \{g_1, g_2, g_3\}$ and a monotone PMMS-feasible valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ that is not nice cancelable. We claim that the constructed valuation also satisfies GMMS-feasibility. We note that the condition of Definition 2.11 holds for any partition with $k = 2$ directly from PMMS-feasibility, and the case $k = 1$ is trivial. Consider the case $k = 3$, and let $A = (A_1, A_2, A_3)$ and $B = (B_1, B_2, B_3)$ be any two partitions of a subset $S \subseteq M$. Either (1) one of the subsets B_1, B_2, B_3 is empty, or (2) each of B_1, B_2, B_3 is a singleton. In case (1), we have $\min\{v(B_1), v(B_2), v(B_3)\} = 0$, so the condition holds trivially. In case (2), we have $\min\{v(B_1), v(B_2), v(B_3)\} = \min\{v(g_1), v(g_2), v(g_3)\}$, and the singleton of minimal value must belong to one of the sets A_1, A_2, A_3, which by monotonicity has value at least that of the singleton. Hence the condition holds. □

C.3 When Does Envy-Freeness Imply GMMS or PMMS?

We first show that PMMS-feasibility is both sufficient and necessary for envy-freeness to imply PMMS, so it precisely characterizes the valuations under which this implication holds.

Proposition C.5. *When v_i is an PMMS-feasible valuation, any allocation that is envy-free for agent i is also PMMS for agent i. Conversely, when v_i is not PMMS-feasible, there exists a (possibly partial) allocation that is envy-free for agent i but not PMMS for agent i.*

Proof. (EF $\Rightarrow$ PMMS under PMMS-feasibility) Let $X = (X_1, \ldots, X_n)$ be an allocation that is envy-free for agent i; that is, $v_i(X_i) \geq v_i(X_j)$ for every other agent j. Fix an agent $j \neq i$, and let (X_i', X_j') be a partition of $X_i \cup X_j$ that attains the pairwise maximin share, so that

$$\mu_i(2, X_i \cup X_j) = \min\{v_i(X_i'),\, v_i(X_j')\}.$$

Then

$$\begin{aligned} v_i(X_i) &= \max\{v_i(X_i),\, v_i(X_j)\} && \text{(by envy-freeness)} \\ &\geq \min\{v_i(X_i'),\, v_i(X_j')\} && \text{(by PMMS-feasibility)} \\ &= \mu_i(2, X_i \cup X_j) && \text{(by the choice of } (X_i', X_j')\text{).} \end{aligned}$$

Since j was arbitrary, X is PMMS for agent i.

(EF $\not\Rightarrow$ PMMS without PMMS-feasibility) Assume v_i is not PMMS-feasible. Then there exist two disjoint bundles $X_i, X_j \subseteq M$ and a partition (X_i', X_j') of $X_i \cup X_j$ such that

$$\max\{v_i(X_i), v_i(X_j)\} < \min\{v_i(X_i'), v_i(X_j')\}.$$

Without loss of generality, assume $v_i(X_i) \geq v_i(X_j)$, and consider the (possibly partial) two-agent allocation $X = (X_i, X_j)$ among agents i and j. This allocation is envy-free for agent i, since $v_i(X_i) \geq v_i(X_j)$. It is not, however, PMMS for agent i:

$$v_i(X_i) = \max\{v_i(X_i), v_i(X_j)\} < \min\{v_i(X_i'), v_i(X_j')\} \leq \mu_i(2, X_i \cup X_j),$$

where the final inequality holds because (X_i', X_j') is a particular partition of $X_i \cup X_j$, while $\mu_i(2, X_i \cup X_j)$ is the maximum of $\min\{v_i(A), v_i(B)\}$ over all such partitions. □

Similarly, we show that GMMS-feasibility is both sufficient and necessary for envy-freeness to imply GMMS, so it precisely characterizes the valuations under which this implication holds.

Proposition C.6. *When v_i is a GMMS-feasible valuation, any allocation that is envy-free for agent i is also GMMS for agent i. Conversely, when v_i is not GMMS-feasible, there exists a (possibly partial) allocation that is envy-free for agent i but not GMMS for agent i.*

Proof. (EF $\Rightarrow$ GMMS under GMMS-feasibility) Let $X = (X_1, \ldots, X_n)$ be an allocation that is envy-free for agent i; that is, $v_i(X_i) \geq v_i(X_j)$ for every other agent j. Fix a subset of agent $i \in J \subseteq N$, and let $(X_j')_{j \in J}$ be a partition of $\bigcup_{j \in J} X_j$ that attains the groupwise maximin share, so that

$$\mu_i(|J|, \bigcup_{j \in J} X_j) = \min_{j \in J} v_i(X_j').$$

Then

$$\begin{aligned} v_i(X_i) &= \max_{j \in J} v_i(X_j) && \text{(by envy-freeness)} \\ &\geq \min_{j \in J} v_i(X_j') && \text{(by GMMS-feasibility)} \\ &= \mu_i(|J|, \bigcup_{j \in J} X_j) && \text{(by the choice of } (X_j')_{j \in J}\text{).} \end{aligned}$$

Since J was arbitrary, X is GMMS for agent i.

(EF $\not\Rightarrow$ GMMS without GMMS-feasibility) Assume v_i is not GMMS-feasible. Then there exist a set of indices J with $i \in J$, disjoint bundles $(X_j)_{j \in J}$, and a partition $(X_j')_{j \in J}$ of $\bigcup_{j \in J} X_j$ such that

$$\max_{j \in J} v_i(X_j) < \min_{j \in J} v_i(X_j').$$

Without loss of generality, assume $v_i(X_i) \geq v_i(X_j)$ for all $j \in J$, and consider the (possibly partial) allocation $X = (X_j)_{j \in J}$ among agents indexed by J, including agent i. This allocation is envy-free for agent i, since $v_i(X_i) \geq v_i(X_j)$ for all $j \in J$ by assumption. It is not, however, GMMS for i:

$$v_i(X_i) = \max_{j \in J} v_i(X_j) < \min_{j \in J} v_i(X_j') \leq \mu_i(|J|, \bigcup_{j \in J} X_j),$$

where the final inequality holds because $(X_j')_{j \in J}$ is a particular partition of $\bigcup_{j \in J} X_j$, while $\mu_i(|J|, \bigcup_{j \in J} X_j)$ is the maximum of $\min_{j \in J} v_i(X_j'')$ over all such partitions X''. □

C.4 GMMS-Feasible Valuations Are Incomparable with Gross Substitutes Valuations

We show that (G)PMMS-feasible valuations are incomparable with the class of gross substitutes valuations [47], an important and well-studied member of the hierarchy of complement-free valuations [46]. We begin by defining the class of OXS valuations [19, 20, 46], which forms a special case of gross substitutes valuations [47].

Definition C.7 (OXS Valuations, [46])**.** *A valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ over a set of items M is an OXS valuation if there exists a finite set X and a weight function $w : M \times X \to \mathbb{R}_{\geq 0}$ such that for every subset $S \subseteq M$:*

$$v(S) = \max_{\pi \in \mathcal{M}(S,X)} \sum_{(j,k) \in \pi} w(j,k)$$

where $\mathcal{M}(S, X)$ is the set of all matchings between the items in S and the elements of X.

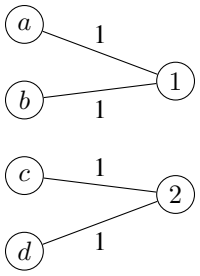


Figure 4: The weight function w used in the proof of Proposition C.8. Edges represent pairs with weight 1; all omitted pairs have weight 0.

We first show that not every gross substitutes valuation is PMMS-feasible.

Proposition C.8. *There exists an OXS valuation function, and consequently a gross substitutes valuation function, that is not PMMS-feasible.*

Proof. (OXS $\not\Rightarrow$ PMMS-feasibility) Let $M = \{a, b, c, d\}$ denote the set of items, and let $X = \{1, 2\}$. Define the weight function $w : M \times X \to \mathbb{R}_{\geq 0}$ by setting $w(a, 1) = w(b, 1) = w(c, 2) = w(d, 2) = 1$ and $w(i, j) = 0$ for all remaining pairs; see Figure 4 for the corresponding bipartite graph. Let v be the OXS valuation function induced by w. This valuation is not PMMS-feasible, since

$$\min\{v(\{a, c\}), v(\{b, d\})\} = 2 > 1 = \max\{v(\{a, b\}), v(\{c, d\})\}.$$

This completes the proof. □

Finally, the analysis of Leme [47, Section 2] shows that not every GMMS-feasible valuation is a gross substitutes valuation.

Proposition C.9 ([47])**.** *There exists a budget-additive valuation function, and consequently a GMMS-feasible valuation, that is not a gross substitutes valuation.*

D Tightness of Our Analysis of the Envy Cycles Procedure

In this section, we show that our analysis of the Envy Cycle Elimination algorithm (Proposition 4.1) is tight: the algorithm guarantees no better than $4/5$-PMMS.

Proposition D.1. *There exists an instance where the output of Algorithm 1 does not satisfy α-PMMS for any $\alpha > \frac{4}{5}$.*

Proof. Assume there are three agents and nine items. Let $\epsilon > 0$ be arbitrarily small, and define $\epsilon_k = \epsilon \cdot \frac{1}{2^{8-k}}$. The value functions are as follows:

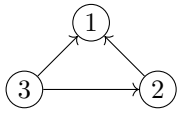
	g_1	g_2	g_3	g_4	g_5	g_6	g_7	g_8	g_9
Agent 1	3	$2+\epsilon_8$	$2+\epsilon_7$	$2+\epsilon_6$	ϵ_5	ϵ_4	ϵ_3	ϵ_2	ϵ_0
Agent 2	$2+\epsilon_8$	$2+\epsilon_7$	$2+\epsilon_6$	$2+\epsilon_5$	$2+\epsilon_4$	$2+\epsilon_3$	$1+\epsilon_2$	1	ϵ_0
Agent 3	$2+\epsilon_8$	$2+\epsilon_6+\epsilon_1$	$2+\epsilon_6$	$2+\epsilon_5$	$2+\epsilon_4$	$2+\epsilon_3$	$1+\epsilon_2$	1	ϵ_0

The perturbations ϵ_k are introduced to ensure that the instance is non-degenerate (as defined above).

We now simulate the algorithm on this instance. We also include the envy graph (whose nodes are the agents and which contains an edge $i \to j$ whenever $v_i(X_j) > v_i(X_i)$) at the end of each step.

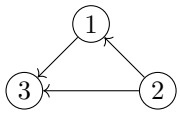
After the first three iterations of the loop, the allocation will be as follows:

$$\begin{aligned} X_1 &= \{g_1\}, \quad v_1(X_1) = 3, \\ X_2 &= \{g_2\}, \quad v_2(X_2) = 2 + \epsilon_7, \\ X_3 &= \{g_3\}, \quad v_3(X_3) = 2 + \epsilon_6 \end{aligned}$$



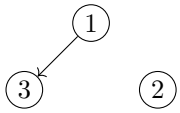
Next, agent 3 receives an item:

$$\begin{aligned} X_1 &= \{g_1\}, & v_1(X_1) &= 3, \\ X_2 &= \{g_2\}, & v_2(X_2) &= 2 + \epsilon_7, \\ X_3 &= \{g_3, g_4\}, & v_3(X_3) &= 4 + \epsilon_6 + \epsilon_5 \end{aligned}$$



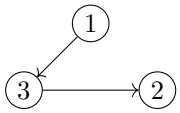
Next, agent 2 receives an item:

$$\begin{aligned} X_1 &= \{g_1\}, & v_1(X_1) &= 3, \\ X_2 &= \{g_2, g_5\}, & v_2(X_2) &= 4 + \epsilon_7 + \epsilon_4, \\ X_3 &= \{g_3, g_4\}, & v_3(X_3) &= 4 + \epsilon_6 + \epsilon_5 \end{aligned}$$



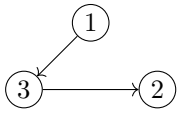
Next, agent 2 receives another item:

$$\begin{aligned} X_1 &= \{g_1\}, & v_1(X_1) &= 3, \\ X_2 &= \{g_2, g_5, g_6\}, & v_2(X_2) &= 6 + \epsilon_7 + \epsilon_4 + \epsilon_3, \\ X_3 &= \{g_3, g_4\}, & v_3(X_3) &= 4 + \epsilon_6 + \epsilon_5 \end{aligned}$$



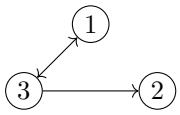
Next, agent 1 receives an item:

$$\begin{aligned} X_1 &= \{g_1, g_7\}, & v_1(X_1) &= 3 + \epsilon_3, \\ X_2 &= \{g_2, g_5, g_6\}, & v_2(X_2) &= 6 + \epsilon_7 + \epsilon_4 + \epsilon_3, \\ X_3 &= \{g_3, g_4\}, & v_3(X_3) &= 4 + \epsilon_6 + \epsilon_5 \end{aligned}$$



Next, agent 1 receives another item:

$$\begin{aligned} X_1 &= \{g_1, g_7, g_8\}, & v_1(X_1) &= 3 + \epsilon_3 + \epsilon_2, \\ X_2 &= \{g_2, g_5, g_6\}, & v_2(X_2) &= 6 + \epsilon_7 + \epsilon_4 + \epsilon_3, \\ X_3 &= \{g_3, g_4\}, & v_3(X_3) &= 4 + \epsilon_6 + \epsilon_5 \end{aligned}$$

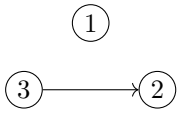


At this point, agent 3 begins to envy agent 1, since

$$v_3(X_1) = 4 + \epsilon_8 + \epsilon_2 \geq 4 + \epsilon_6 + \epsilon_5 = v_3(X_3).$$

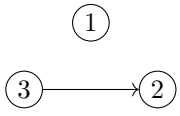
We therefore resolve the cycle between agents 1 and 3:

$$\begin{aligned} X_1 &= \{g_3, g_4\}, & v_1(X_1) &= 4 + \epsilon_7 + \epsilon_6, \\ X_2 &= \{g_2, g_5, g_6\}, & v_2(X_2) &= 6 + \epsilon_7 + \epsilon_4 + \epsilon_3, \\ X_3 &= \{g_1, g_7, g_8\}, & v_3(X_3) &= 4 + \epsilon_8 + \epsilon_2 \end{aligned}$$



Finally, agent 3 receives the last item:

$$\begin{aligned} X_1 &= \{g_3, g_4\}, & v_1(X_1) &= 4 + \epsilon_7 + \epsilon_6, \\ X_2 &= \{g_2, g_5, g_6\}, & v_2(X_2) &= 6 + \epsilon_7 + \epsilon_4 + \epsilon_3, \\ X_3 &= \{g_1, g_7, g_8, g_9\}, & v_3(X_3) &= 4 + \epsilon_8 + \epsilon_2 + \epsilon_0 \end{aligned}$$



The algorithm now terminates. Observe that

$$v_3(X_3) = 4 + \epsilon_8 + \epsilon_2 + \epsilon_0,$$

while

$$\mu_3(2, X_2 \cup X_3) = \min\left(v_3(\{g_1, g_6, g_8\}),\, v_3(\{g_2, g_5, g_7, g_9\})\right) = 5 + \epsilon_6 + \epsilon_4 + \epsilon_2 + \epsilon_1 + \epsilon_0.$$

Thus, we cannot guarantee better than a $\frac{4}{5}$-PMMS allocation using Algorithm 1. □

E Missing Proofs

E.1 Missing Proofs from Section 4

Proposition 4.1 (Additive Valuations with Identical Rankings)**.** *Every instance with additive valuations and identical rankings admits a $4/5$-PMMS, EFX and $2/3$-MMS allocation. Moreover, this allocation can be computed in polynomial time.*

Proof of Proposition 4.1 . We assume that the valuations are non-degenerate, meaning that $v_i(S) \neq v_i(T)$ for every agent i and every pair of distinct bundles S, T. This assumption is without loss of generality; see Proposition G.4 for a formal argument.

We claim that the Envy Cycle Elimination algorithm (Algorithm 1, [49]) produces an $4/5$-PMMS and EFX allocation in polynomial time. Polynomial-time termination was established by Lipton et al. [49], and Plaut and Roughgarden [51] showed that the resulting allocation satisfies EFX; it therefore remains to prove that it also satisfies $4/5$-PMMS.

Let X be an allocation produced by the above procedure, and let $i, j \in N$. We show that

$$v_i(X_i) \geq \tfrac{4}{5}\mu_i(2, X_i \cup X_j). \tag{4}$$

Case 1: $|X_j| = 1$. Denote $X_j = \{g\}$. Using Lemma 4.2 we obtain

$$\mu_i(2, X_i \cup X_j) \leq \mu_i(2, \{g\}) + v_i(X_i) = v_i(X_i),$$

implying Equation 4, as desired.

Case 2: $|X_j| \geq 2$. Consider the point during the execution of Algorithm 1 at which bundle X_j is formed: some agent j' holds a bundle $\tilde{X}_{j'} = X_j \setminus \{g\}$, and item g is added to it in Line 5, producing X_j. Note that j' need not equal j, since the bundle may be reassigned to agent j at a later stage via the envy-cycle operation in Line 8.

Let $\tilde{X}_i$ denote the bundle held by agent i at that point. Since agent i does not envy agent j' at the moment g is added to $\tilde{X}_{j'}$, we have $v_i(\tilde{X}_i) \geq v_i(\tilde{X}_{j'})$. Moreover, by the design of the algorithm, the value each agent has for her own bundle can only increase over time. In particular, $v_i(X_i) \geq v_i(\tilde{X}_i) \geq v_i(\tilde{X}_{j'}) = v_i(X_j \setminus \{g\})$.

We distinguish between two further cases. If $v_i(g) \leq \frac{1}{2}v_i(X_i)$, then

$$\begin{aligned}
\mu_i(2, X_i \cup X_j) &\leq \tfrac{1}{2}(v_i(X_i) + v_i(X_j)) && \text{(since } \mu_i(2,S) \leq \tfrac{1}{2}v_i(S) \text{ for all } S\text{)} \\
&= \tfrac{1}{2}(v_i(X_i) + v_i(X_j \setminus \{g\}) + v_i(g)) && \text{(by additivity of } v_i\text{)} \\
&\leq \tfrac{1}{2}(v_i(X_i) + v_i(X_i) + \tfrac{1}{2}v_i(X_i)) && \text{(by the assumption and observation above)} \\
&= \tfrac{5}{4}v_i(X_i),
\end{aligned}$$

implying Equation 4, as desired.

Consider next the case where $v_i(g) > \frac{1}{2}v_i(X_i)$. Note that at the moment g is added to X_j, the bundle X_i may not yet have been fully formed; however, there must be some agent i' holding a bundle $\tilde{X}_{i'} \subseteq X_i$ at that point, with all items in $X_i \setminus \tilde{X}_{i'}$ still unallocated. Since items are allocated in descending order of value, for every $g' \in \tilde{X}_{i'} \cup (X_j \setminus \{g\})$ we have

$$v_i(g') \geq v_i(g) > \tfrac{1}{2}v_i(X_i). \tag{5}$$

We now show that $|\tilde{X}_{i'}| = 1$. On one hand, since $v_i(\tilde{X}_{i'}) \leq v_i(X_i)$ and each item in $\tilde{X}_{i'}$ is worth at least $\frac{1}{2}v_i(X_i)$ to agent i by Equation (5), we have $|\tilde{X}_{i'}| \leq 1$. On the other hand, since item g is assigned to agent j', agent i' does not envy the bundle $\tilde{X}_{j'} = X_j \setminus \{g\}$; hence, by the non-degeneracy of the valuations, agent i' must hold at least one item.

We next show that $|X_j| = 2$. By assumption, we already have $|X_j| \geq 2$, and strict inequality is impossible since $v_i(X_i) \geq v_i(X_j \setminus \{g\})$ and, by Equation (5), each item in $X_j \setminus \{g\}$ is worth at least $\frac{1}{2}v_i(X_i)$ to agent i.

Denote $X_j = \{g', g\}$ and $\tilde{X}_{i'} = \{h\}$. We next show that

$$v_i(h) > v_i(g') > v_i(g). \tag{6}$$

Note that at the moment item g is first allocated, we have $\tilde{X}_{j'} = \{g'\}$ and $\tilde{X}_{i'} = \{h\}$. Since items are allocated in descending order and, by assumption, no two items share the same value, it follows that $v_i(g') > v_i(g)$ and $v_i(h) > v_i(g')$. Moreover, since item g is assigned to agent j', agent i' does not envy agent j' at that point, so $v_{i'}(h) > v_{i'}(g')$; by the identical ranking assumption, this implies $v_i(h) > v_i(g')$. We get

$$\begin{aligned}
\mu_i(2, X_i \cup X_j) &\leq \mu_i(2, (X_i \cup X_j) \setminus (X_i \setminus \{h\})) + v_i(X_i \setminus \{h\}) && \text{(by Lemma 4.2)} \\
&= \mu_i(2, \{g, g', h\}) + v_i(X_i \setminus \{h\}) && \text{(by the above)} \\
&\leq v_i(h) + v_i(X_i \setminus \{h\}) && \text{(by Equation (6))} \\
&= v_i(X_i) && \text{(by additivity)}
\end{aligned}$$

implying Equation 4, as desired. □

E.2 Missing Proofs from Section 5

Lemma E.1 (Leximin Substitution)**.** *Let S be a set of items, $k \geq 2$ be an integer, and v be an additive valuation function. Let $P = \{A_1, A_2, \ldots, A_k\}$ be an optimal leximin partition of S into k bundles. If $P^\star$ is any optimal leximin partition of $S \setminus A_1$ into $k-1$ bundles under v, then $\{A_1\} \cup P^\star$ is also an optimal leximin partition of S into k bundles under v.*

Proof for Lemma E.1. First, observe that the sub-partition $P \setminus \{A_1\} = \{A_2, \ldots, A_k\}$ must itself be an optimal leximin partition of $S \setminus A_1$ into $k-1$ bundles. Otherwise, there would exist a partition of $S \setminus A_1$ with a strictly larger lexicographically sorted value vector, which, when combined with A_1, would yield a partition of S that strictly dominates P, contradicting the optimality of P.

Since both $P^\star$ and $P \setminus \{A_1\}$ are optimal leximin partitions of the same instance, their sorted valuation vectors must be identical.

Now consider the partition $\{A_1\} \cup P^\star$. Its multiset of bundle values consists of $v(A_1)$ together with the values induced by $P^\star$. Because the sorted valuation vector of $P^\star$ matches that of $P \setminus \{A_1\}$, it follows that the sorted valuation vector of $\{A_1\} \cup P^\star$ is identical to that of P.

Since P is an optimal leximin partition of S, any partition with the same lexicographically sorted valuation vector is also optimal. Therefore, $\{A_1\} \cup P^\star$ is an optimal leximin partition of S. □

Lemma E.2 (MMS Monotonicity)**.** *Let M be a set of m items, $n \geq 2$ be an integer, and v be an additive valuation function. Let $X \subseteq M$ be a subset with $|X| \leq \frac{m}{n}$, consisting of the least-valued items in M. Then, $\mu(n-1, M \setminus X) \geq \mu(n, M)$.*

Proof for Lemma E.2. Let $P = \{A_1, \ldots, A_n\}$ be an MMS partition of M into n bundles, i.e., $v(A_j) \geq \mu(n, M)$ for all $j \in [n]$. By the pigeonhole principle, since m items are distributed among n bundles, there exists a bundle $B \in P$ such that $|B| \geq \frac{m}{n}$.

Since $P \setminus \{B\}$ is a valid partition of $M \setminus B$ into $n-1$ bundles, and for every $A \in P \setminus \{B\}$ we have $v(A) \geq \mu(n, M)$, it must hold that

$$\mu(n-1, M \setminus B) \geq \mu(n, M). \tag{7}$$

In addition, since $|X| \leq \frac{m}{n} \leq |B|$, it follows that $|X| \leq |B|$. Moreover, because X consists of the least-valued items in M, removing X can only increase (or at least not decrease) the MMS compared to removing B. Therefore,

$$\mu(n-1, M \setminus X) \geq \mu(n-1, M \setminus B). \tag{8}$$

Combining (7) and (8) yields

$$\mu(n-1, M \setminus X) \geq \mu(n, M).$$ □

Theorem 5.1 (Bivalued Valuations)**.** *Every instance with bivalued valuations admits an EGMMS, and hence an EPMMS, allocation. Moreover, this allocation can be computed in polynomial time.*

Proof of Theorem 5.1. We claim that Algorithm 2 produces an EGMMS allocation. Since the algorithm orders the agents by their total valuation of the grand bundle ($v_1(M) \geq v_2(M) \geq \cdots \geq v_n(M)$), and valuations are bivalued with identical rankings, this ordering coincides with the lengths of their b-valued prefixes. Thus, $|p_b(1)| \geq |p_b(2)| \geq \cdots \geq |p_b(n)|$, and in particular, for any $j < i$ we have $p_b(i) \subseteq p_b(j)$. We now show that the resulting allocation X is EGMMS. Fix an agent i. We construct a valid EGMMS certificate Y^i for agent i:

$$Y^i = (X_1, X_2, \ldots, X_i, B^i_{i+1}, B^i_{i+2}, \ldots, B^i_n).$$

Namely, $Y^i_j = X_j$ for $j \leq i$, and $Y^i_j = B^i_j$ for $j > i$. We claim that this partition forms an EGMMS certificate for agent i. Since it forms a partition of the entire set of items M, and $Y^i_i = X_i$, it remains to show that agent i is GMMS-satisfied with Y^i. That is, we need to show that for every subset $J \subseteq [n]$ with $i \in J$, it holds that:

$$v_i(X_i) \geq \mu_i(|J|, \mathcal{M}_J), \quad \text{where} \quad \mathcal{M}_J = \bigcup_{j \in J} Y^i_j. \tag{9}$$

Let $j_1 = \min(J)$. Because $i \in J$, it holds that $j_1 \leq i$. We prove Equation (9) by (backward) induction on the minimum index j_1, starting from $j_1 = i$ down to $j_1 = 1$.

Base Case ($j_1 = i$): If $\min(J) = i$, then $J \subseteq \{i, i+1, \ldots, n\}$. The subset $\{Y^i_j \mid j \in J\} = \{B^i_j \mid j \in J\}$ consists of bundles corresponding to a subset of agent i's leximin partition of $M^{(i)}$. By Proposition 3.2, this is a GMMS allocation for the instance in which all agents have identical valuations v_i. Thus $v_i(X_i) \geq \mu_i(|J|, \mathcal{M}_J)$, as desired.

Inductive Step ($j_1 < i$): Assume the GMMS condition holds for all valid subsets J' where $i \in J'$ and $\min(J') > j_1$, i.e., $v_i(X_i) \geq \mu_i(|J'|, \bigcup_{j \in J'} Y^i_j)$.

Case 1: $Y^i_{j_1} \cap p_b(i) \neq \emptyset$ (some item in $p_b(i)$ is allocated to agent j_1).
Since $j_1 < i$, it holds that $p_b(i) \subseteq p_b(j_1)$. In the algorithm, agent j_1 takes the *highest-indexed* available items from $p_b(j_1)$. If agent j_1 is forced to take an item from $p_b(i)$ (which is a prefix of $p_b(j_1)$), it implies that all available items in the disagreement zone, $p_b(j_1) \setminus p_b(i)$, have already been allocated either to agent j_1 or to an earlier agent.

Consequently, for every remaining item $t \in M^{(j_1+1)}$, the valuations of agents j_1 and i are identical. Since agent i has the shortest prefix among all agents between j_1 and i by the assumption on the ordering of agents, it holds that for every agent $\ell \in \{j_1, \ldots, i\}$, we have $p_b(i) \subseteq p_b(\ell) \subseteq p_b(j_1)$, and thus, for every remaining item $t \in M^{(j_1+1)}$, the valuations are identical: $v_\ell(t) = v_{j_1}(t)$.

We now prove that the certificate sequence $(Y^i_{j_1}, Y^i_{j_1+1}, \ldots, Y^i_n)$ forms a valid leximin partition of $M^{(j_1)}$ under valuation v_{j_1}. We show this by working backward from agent i and applying Lemma E.1.

First, consider the items $M^{(i)}$ remaining at agent i's turn, which are partitioned into the suffix $(Y^i_i, \ldots, Y^i_n)$. By definition of the certificate, this sequence corresponds to an optimal leximin partition under v_i. Since v_i and v_{j_1} are identical over these items, $(Y^i_i, \ldots, Y^i_n)$ is also an optimal leximin partition of $M^{(i)}$ under v_{j_1}.

Now, assume inductively that for some index $k \in \{j_1 + 1, \ldots, i\}$, the suffix $(Y^i_k, \ldots, Y^i_n)$ is an optimal leximin partition of $M^{(k)}$ under v_{j_1}. At step $k-1$, agent $k-1$ computes an optimal leximin partition of $M^{(k-1)}$ under v_{k-1} and is assigned bundle $Y^i_{k-1} = X_{k-1}$. Observe that v_{k-1} and v_{j_1} agree on $M^{(k-1)}$: this follows from the observation above when $k - 1 > j_1$, and trivially when $k - 1 = j_1$. Thus, the bundle Y^i_{k-1} is also part of a valid leximin partition of $M^{(k-1)}$ under v_{j_1}.

By Lemma E.1, since Y^i_{k-1} is a bundle from a valid leximin partition under v_{j_1}, and the subsequent sequence $(Y^i_k, \ldots, Y^i_n)$ forms a valid leximin partition of the remaining items $M^{(k)} = M^{(k-1)} \setminus Y^i_{k-1}$ under v_{j_1}, joining them yields an optimal leximin partition. Thus, $(Y^i_{k-1}, Y^i_k, \ldots, Y^i_n)$ is a valid leximin partition of $M^{(k-1)}$ under v_{j_1}.

Applying this logic inductively backward from i down to j_1, we conclude that the entire constructed sequence $(Y^i_{j_1}, \ldots, Y^i_n)$ forms a valid leximin partition of $M^{(j_1)}$ under valuation v_{j_1}.

Because a leximin partition is a GMMS allocation for identical agents (Proposition 3.2), it holds that:

$$v_{j_1}(X_i) = v_{j_1}(Y_i^i) \geq \mu_{j_1}(|J|, \mathcal{M}_J). \tag{10}$$

Moreover, since $X_i \subseteq M^{(j_1+1)}$, agent i and agent j_1 value its items identically:

$$v_i(X_i) = v_{j_1}(X_i). \tag{11}$$

Finally, because $j_1 < i$, agent j_1 has at least as many b-valued items in $\mathcal{M}_J$ as agent i does. Consequently, agent j_1's valuation weakly dominates agent i's. Thus, the maximin share for agent j_1 is at least as large as that for agent i:

$$\mu_{j_1}(|J|, \mathcal{M}_J) \geq \mu_i(|J|, \mathcal{M}_J). \tag{12}$$

Chaining (11), (10), and (12) together, we conclude that

$$v_i(X_i) \geq \mu_i(|J|, \mathcal{M}_J).$$

Case 2: $Y_{j_1}^i \cap p_b(i) = \emptyset$ (No item of $p_b(i)$ is allocated to agent j_1).
In this case, according to v_i, all items in X_{j_1} have value a.

We distinguish between two further cases. Consider first the case where $|X_{j_1}| \leq \frac{|\mathcal{M}_J|}{|J|}$.
Let $J' = J \setminus \{j_1\}$. Because j_1 is the unique minimum index of J, we have $\min(J') > j_1$. By the inductive hypothesis, we have $v_i(X_i) \geq \mu_i(|J'|, \mathcal{M}_{J'})$.

By Lemma E.2, since $|X_{j_1}| \leq \frac{|\mathcal{M}_J|}{|J|}$ and X_{j_1} consists only of items of minimum value a (under v_i), it follows that

$$\mu_i(|J'|, \mathcal{M}_{J'}) = \mu_i(|J| - 1, \mathcal{M}_J \setminus X_{j_1}) \geq \mu_i(|J|, \mathcal{M}_J).$$

Chaining this with the inductive hypothesis yields $v_i(X_i) \geq \mu_i(|J|, \mathcal{M}_J)$, completing the proof for this subcase.

Now, let us examine the complementary case where $|X_{j_1}| > \frac{|\mathcal{M}_J|}{|J|}$.
In Algorithm 2, X_{j_1} is chosen as the minimum cardinality bundle of a partition of $M^{(j_1)}$ into $n - j_1 + 1$ bundles. Thus, $|X_{j_1}| \leq \frac{|M^{(j_1)}|}{n-j_1+1}$.

Combining this with our current assumption, we have:

$$\frac{|M^{(j_1)}|}{n - j_1 + 1} > \frac{|\mathcal{M}_J|}{|J|}. \tag{13}$$

Denote $N_{j_1} = \{j_1, j_1 + 1, \ldots, n\}$ as the suffix of agents from j_1 to n. Notice that the bundles $\{Y_k^i \mid k \in N_{j_1}\}$ form a partition of the items in $M^{(j_1)}$. The left-hand side of Inequality 13 is the average bundle size across this entire partition, while the right-hand side of Inequality 13 is the average bundle size of the subset of bundles indexed by $J \subseteq N_{j_1}$.

Because the average size of the subset J is strictly less than the overall average, the average size of the remaining bundles indexed by $N_{j_1} \setminus J$ must be strictly greater than the overall average. Therefore, there must exist at least one bundle in the complement set, indexed by some $j' \in N_{j_1} \setminus J$, whose size is strictly greater than the overall average:

$$|Y_{j'}^i| > \frac{|M^{(j_1)}|}{n - j_1 + 1} \geq |X_{j_1}|. \tag{14}$$

Let $J' = (J \setminus \{j_1\}) \cup \{j'\}$. Because j_1 is the unique minimum element of J, and $j' \in N_{j_1} \setminus J$ implies $j' > j_1$, it holds that $\min(J') > j_1$. By the inductive hypothesis, we apply the GMMS condition to J':

$$v_i(X_i) \geq \mu_i(|J'|, \mathcal{M}_{J'}). \tag{15}$$

We now compare the maximin share for agent i of $\mathcal{M}_{J'}$ to that of $\mathcal{M}_J$. Notice that $\mathcal{M}_{J'} = (\mathcal{M}_J \setminus X_{j_1}) \cup Y_{j'}^i$. From the condition of Case 2, every item in X_{j_1} has a value of exactly a according to v_i. Meanwhile, every item in $Y_{j'}^i$ has a value of at least a. Also, $|Y_{j'}^i| > |X_{j_1}|$.

Consider the MMS partition of $\mathcal{M}_J$ into $|J|$ bundles based on v_i. If we replace the items of X_{j_1} in this partition with exactly $|X_{j_1}|$ items from $Y^i_{j'}$, and distribute the remaining items of $Y^i_{j'}$ arbitrarily, the value of every bundle in the partition weakly increases. This forms a valid partition of $\mathcal{M}_{J'}$ into $|J'|$ bundles where the minimum bundle is at least as valuable as the minimum bundle in the optimal partition of $\mathcal{M}_J$. Thus:

$$\mu_i(|J'|, \mathcal{M}_{J'}) \geq \mu_i(|J|, \mathcal{M}_J). \tag{16}$$

Chaining (15) and (16) yields:

$$v_i(X_i) \geq \mu_i(|J|, \mathcal{M}_J).$$

This satisfies the GMMS condition for this subcase, completing the inductive step.

Therefore, agent i is GMMS-satisfied with Y^i for all valid subsets J. Since this holds for any arbitrary agent i, the allocation X is EGMMS for the ordered instance. By Proposition 3.4, mapping this allocation back to the original items yields an EGMMS allocation for the general instance, completing the proof. □

Observation 5.2. *There exists a polynomial time algorithm that computes the leximin allocation for any instance with identical bivalued valuations.*

Proof. Let n be the number of agents, N_a the number of items of value a, and N_b the number of items of value b. Consider the following dynamic programming solution. For all $1 \leq k \leq n$, $0 \leq n_a \leq N_a$, and $0 \leq n_b \leq N_b$, let

$$\mathrm{DP}[k][n_a][n_b]$$

denote the sorted (non-decreasing) value vector of a leximin allocation of the sub-instance consisting of k agents with identical bivalued valuations, n_a items of value a, and n_b items of value b.

For the base case $k = 1$, the unique allocation gives the single agent every item, so we set $\mathrm{DP}[1][n_a][n_b] = (n_a \cdot a + n_b \cdot b)$.

For $k > 1$, assume inductively that $\mathrm{DP}[k'][n'_a][n'_b]$ has already been computed for every triple (k', n'_a, n'_b) with $k' \leq k$, $n'_a \leq n_a$, $n'_b \leq n_b$, and at least one of these inequalities strict. We iterate over all pairs (i, j) with $0 \leq i \leq n_a$ and $0 \leq j \leq n_b$, where (i, j) corresponds to assigning a bundle of i items of value a and j items of value b to one agent. For each such pair, we form a candidate vector $c[i][j]$ by inserting the value $i \cdot a + j \cdot b$ into the vector $\mathrm{DP}[k-1][n_a - i][n_b - j]$ and re-sorting in non-decreasing order. We then set $\mathrm{DP}[k][n_a][n_b]$ to be the lexicographically maximal vector among all such candidates $c[i][j]$.

Since the number of states is nN_aN_b and the number of pairs iterated over in each state is at most N_aN_b, the algorithm runs in polynomial time. □

E.3 Missing Proofs from Section 6

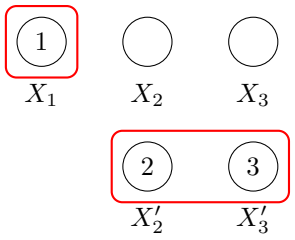
Theorem 6.1 (Three Agents)**.** *Every instance with three additive agents admits an EPMMS allocation. More generally, this holds whenever at least two of the agents are PMMS-feasible.*

Proof of Theorem 6.1. Consider the case of three agents, and suppose that v_2 and v_3 are PMMS-feasible. Let (X_1, X_2, X_3) be agent 1's leximin partition. Assume without loss of generality that X_1 is the least preferred bundle according to agent 2, i.e., $v_2(X_2) \geq v_2(X_1)$ and $v_2(X_3) \geq v_2(X_1)$.

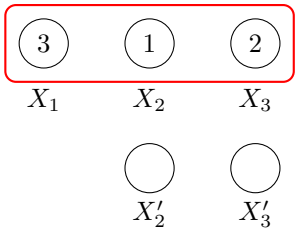
Let (X'_2, X'_3) be agent 2's leximin partition of $X_2 \cup X_3$ into two bundles. We then allocate to agent 3 her most preferred bundle among X_1, X'_2, X'_3.

We next describe the allocation of agents 1 and 2. If X_1 is still available, then agent 1 takes X_1, and agent 2 receives the remaining bundle. Otherwise, if X_1 is taken by agent 3, then agent 2 chooses her preferred bundle among X_2, X_3 (the bundles from agent 1's partition), and agent 1 receives the remaining bundle. See Figures 5a and 5b for an illustration of the allocation.

We first analyze agent 3's allocation. Agent 3 receives the maximal bundle in the partition (X_1, X'_2, X'_3). Thus, the partition (X_1, X'_2, X'_3) serves as her EPMMS certificate: by Proposition C.5, since she receives the maximal bundle in this partition, she does not PMMS-envy any other bundle in the partition since v_3 is PMMS-feasible. Moreover, if we assume that her valuation is also GMMS-feasible, then she receives at least her MMS value.



(a) Case 1: X_1 remains available. Agent 3 selects from $\{X_1, X'_2, X'_3\}$; agent 1 receives X_1, and agent 2 receives the remaining bundle among X'_2, X'_3.



(b) Case 2: Agent 3 takes X_1. Agent 2 then chooses her preferred bundle among X_2 and X_3, and agent 1 receives the remaining bundle.

Figure 5: The two cases in the proof of Theorem 6.1.

Next, we argue that agent 2 does not PMMS-envy any other agent.

If she receives one of the bundles in (X'_2, X'_3), then, by the choice of (X'_2, X'_3),

$$\min\{v_2(X'_2), v_2(X'_3)\} \geq \min\{v_2(X_2), v_2(X_3)\} \geq v_2(X_1).$$

Thus, she does not PMMS-envy the bundle X_1. Moreover,

$$\min\{v_2(X'_2), v_2(X'_3)\} \geq \mu_2(2, X'_2 \cup X'_3),$$

and therefore she does not PMMS-envy the remaining bundle among X'_2 and X'_3.

If she receives her most preferred bundle in (X_2, X_3), then she receives the maximal bundle in the allocation; thus, she does not PMMS-envy any other agent. Since the final allocation is a PMMS allocation from the perspective of agent 2, it also serves as her EPMMS certificate.

Finally, since agent 1 receives one of X_1, X_2, X_3, The partition (X_1, X_2, X_3) serves as her EPMMS certificate, since it is agent 1 leximin partition. Moreover, since the leximin partition is also an MMS partition, agent 1 receives at least her MMS share. □

Theorem 6.2 (Two Types of Agents)**.** *Every instance with two types of additive agents admits an EPMMS allocation. More generally, this holds whenever at least one of the types is PMMS-feasible.*

Proof of Theorem 6.2. We consider the case where there are only two types of valuation functions, v_α and v_β. That is, for each agent i, either $v_i = v_\alpha$ or $v_i = v_\beta$. Let N_α be the set of agents whose valuation functions are v_α, and let N_β be the set of agents whose valuation functions are v_β. Assume that v_β is an PMMS-feasible valuation.

Let $(X_1, X_2, \ldots, X_n)$ be a leximin partition of M into n bundles according to v_α. Assume without loss of generality that $X_1, \ldots, X_{|N_\alpha|}$ are the $|N_\alpha|$ least-preferred bundles according to v_β.

Let $(X'_{|N_\alpha|+1}, \ldots, X'_n)$ be a leximin partition of $\bigcup_{i=|N_\alpha|+1}^{n} X_i$ into $|N_\beta|$ bundles according to v_β.

We now argue that the allocation $(X_1, \ldots, X_{|N_\alpha|}, X'_{|N_\alpha|+1}, \ldots, X'_n)$, in which the agents in N_α receive bundles $X_1, \ldots, X_{|N_\alpha|}$ and the agents in N_β receive bundles $X'_{|N_\alpha|+1}, \ldots, X'_n$, is an EPMMS allocation.

First, we show that agents in N_β satisfy the PMMS condition. By the assumption, $X_1, \ldots, X_{|N_\alpha|}$ are the $|N_\alpha|$ bundles of minimum value according to v_β. Hence every bundle in $(X_{|N_\alpha|+1}, \ldots, X_n)$ has value at least

$$\max_{j \leq |N_\alpha|} v_\beta(X_j).$$

Now consider the leximin partition $(X'_{|N_\alpha|+1}, \ldots, X'_n)$ of $\bigcup_{i=|N_\alpha|+1}^{n} X_i$ with respect to v_β. By the definition of the leximin solution, the value of the minimum bundle in this partition is at least the minimum value in the original partition $(X_{|N_\alpha|+1}, \ldots, X_n)$. Therefore, for every $i > |N_\alpha|$,

$$v_\beta(X'_i) \geq \min_{j > |N_\alpha|} v_\beta(X'_j) \geq \min_{j > |N_\alpha|} v_\beta(X_j) \geq \max_{j \leq |N_\alpha|} v_\beta(X_j).$$

Thus each agent in N_β prefers their bundle to any bundle allocated to an agent in N_α, and since v_β is an PMMS-feasible valuation, they do not PMMS-envy agents in N_α.

Among the agents in N_β, the bundles $X'_{|N_\alpha|+1}, \ldots, X'_n$ form a leximin partition under the identical valuation function v_β. Since leximin allocations with identical valuations are known to satisfy PMMS by Proposition 3.2, no pair of agents in N_β exhibits PMMS envy.

Finally, we show that agents in N_α satisfy EPMMS. The partition $(X_1, X_2, \ldots, X_n)$ is an EPMMS certificate for the agents in N_α, since it is a leximin partition under the identical valuation function v_α, and leximin allocations with identical valuations are known to satisfy PMMS by Proposition 3.2. □

F EPMMS and EGMMS Reductions to Identical Rankings

F.1 The Ordering Framework and Picking Sequence

The ordering framework is a powerful tool in fair division, typically used to design approximation algorithms for Maximin Share (MMS) allocations. Recently, Caragiannis et al. [26] utilized this framework to prove the existence of Epistemic EFX (EEFX) allocations for additive valuations. In this paper, we adapt their approach to establish guarantees for EPMMS approximations and EGMMS allocations under bivalued valuations. The core idea is to transform an instance with arbitrary valuations into a simplified instance where all agents share an identical ranking over the items, solve the allocation problem there, and then map the solution back to the original instance.

Let $\mathcal{I}$ be a fair division instance. For each agent $i \in [n]$, let $(\ell_1, \ell_2, \ldots, \ell_m)$ be an ordering of the items in M sorted by i's valuation in non-increasing order, such that:

$$v_i(\ell_1) \geq v_i(\ell_2) \geq \cdots \geq v_i(\ell_m).$$

Intuitively, ℓ_t represents agent i's t-th favorite item.

Definition F.1 (Item Rank)**.** *For an agent $i \in [n]$ and an item $j \in M$, the* rank *of j, denoted by $r_i(j)$, is the index t such that $j = \ell_t$. For a set of items $S \subseteq M$, we define the set of their ranks as:*

$$r_i(S) := \{\, r_i(j) \mid j \in S \,\}.$$

Definition F.2 (Ordered Valuation and Instance)**.** *The* ordered valuation function *of agent i, denoted by $\hat{v}_i$, is defined for any subset of ranks $T \subseteq [m]$ as:*

$$\hat{v}_i(T) := v_i(\{\, \ell_t \mid t \in T \,\}).$$

We denote by $\hat{\mathcal{I}} = Order(\mathcal{I})$ the ordered fair division instance, which contains the same agents and the same number of items as $\mathcal{I}$, but where each agent i has the valuation function $\hat{v}_i$ over the items $[m]$.

Claim F.3. *For every subset of items $S \subseteq M$ and every agent $i \in [n]$, it holds that:*

$$v_i(S) = \hat{v}_i\big(r_i(S)\big).$$

Example F.4. *Suppose there are four items, $M = \{1, 2, 3, 4\}$, and agent i's valuation is:*

$$v_i(1) = 6, \quad v_i(2) = 11, \quad v_i(3) = 8, \quad v_i(4) = 9.$$

Their ranks are:

$$r_i(1) = 4, \quad r_i(2) = 1, \quad r_i(3) = 3, \quad r_i(4) = 2.$$

The ordered valuation function $\hat{v}_i$ evaluates the t-th best items directly:

$$\hat{v}_i(1) = 11, \quad \hat{v}_i(2) = 9, \quad \hat{v}_i(3) = 8, \quad \hat{v}_i(4) = 6$$

F.1.1 The Reduction Algorithm

To map allocations from the ordered instance $\hat{\mathcal{I}}$ back to the original instance $\mathcal{I}$, we employ a picking sequence algorithm.

1. **Ordering:** Construct the ordered instance $\hat{\mathcal{I}} = \text{Order}(\mathcal{I})$.
2. **Base Allocation:** Compute the target allocation $\hat{X} = (\hat{X}_1, \ldots, \hat{X}_n)$ for the identical rankings instance $\hat{\mathcal{I}}$.

3. **Picking Sequence:** Define the sequence $L = (a_1, a_2, \ldots, a_m)$ such that $a_k = i$ if the k-th item in the identical ranking is assigned to agent i in $\hat{X}$.
4. **Actual Allocation:** Execute the $\textsc{Pick}(\mathcal{I}, L)$ routine: in each step k, agent a_k picks their favorite remaining item according to their original valuation v_i. Let X be the resulting allocation.

To analyze how values shift when translating $\hat{X}$ to X, we rely on a structural lemma introduced by Caragiannis et al [26].

Lemma F.5 (Lemma 2 in [26])**.** *For every agent $i \in [n]$, there exists a bijection $\pi_i : [m] \to M$ such that:*

- *For every $j \in \hat{X}_i$, we have $\pi_i(j) \in X_i$ and $r_i(\pi_i(j)) \leq j$.*
- *For every $j \in [m] \setminus \hat{X}_i$, we have $\pi_i(j) \in M \setminus X_i$ and $r_i(\pi_i(j)) \geq j$.*

Intuitively, Lemma F.5 guarantees that when an agent transitions from their assigned bundle $\hat{X}_i$ to their actual bundle X_i, every item they receive *weakly improves* in rank, while every item they do not receive *weakly degrades* in rank.

Having established our ordering framework and the reduction algorithm, we are now ready to prove Proposition 3.4:

Proof for Proposition 3.4. Let $\mathcal{I}$ be an arbitrary fair division instance, and let $\hat{\mathcal{I}}$ be the corresponding identical rankings instance obtained via the ordering framework. Let $\hat{X}$ be the α-EPMMS (resp., α-EGMMS) allocation in $\hat{\mathcal{I}}$, and let X be the resulting allocation generated by the picking sequence algorithm.

Fix an arbitrary agent $i \in [n]$. Let $\pi_i : [m] \to [m]$ be the bijection guaranteed by lemma F.5, mapping items from the ordered instance to the original instance.

Let $\hat{Y}^i = (\hat{Y}^i_1, \ldots, \hat{Y}^i_n)$ be the α-EPMMS (resp., α-EGMMS) certificate for agent i with respect to the bundle $\hat{X}_i$ in $\hat{\mathcal{I}}$. By definition, $\hat{Y}^i_i = \hat{X}_i$. We construct the corresponding certificate $Y^i = (Y^i_1, \ldots, Y^i_n)$ in the original instance $\mathcal{I}$ by mapping the items through π_i:

$$Y^i_k = \{\, \pi_i(g) \mid g \in \hat{Y}^i_k \,\} \quad \text{for all } k \in [n].$$

Because π_i is a bijection, Y^i is a well-defined allocation of all items, and by the properties of π_i, $Y^i_i = X_i$. We must now show that Y^i is a valid α-EPMMS (resp., α-EGMMS) certificate for agent i.

Let $J \subseteq [n]$ be a subset of agents such that $i \in J$.

- For α-EPMMS, we restrict J to subsets of size 2 (i.e., $J = \{i, j\}$ for some $j \neq i$).
- For α-EGMMS, J can be any subset containing i.

Let $k = |J|$ and define the combined bundle $S = \bigcup_{j \in J} Y^i_j$. Correspondingly, let $\hat{S} = \bigcup_{j \in J} \hat{Y}^i_j$ be the combined bundle in the ordered instance.

Take an arbitrary k-partition $\mathcal{P} = (P_1, P_2, \ldots, P_k)$ of S. This partition naturally defines a corresponding k-partition $\hat{\mathcal{P}} = (\hat{P}_1, \hat{P}_2, \ldots, \hat{P}_k)$ of $\hat{S}$ where $\hat{P}_m = \{\, \pi_i^{-1}(g) \mid g \in P_m \,\}$.

Because $\hat{Y}^i$ is an α-EPMMS (resp., α-EGMMS) certificate in the ordered instance, agent i's valuation of $\hat{X}_i$ must be at least α times the value of the minimum piece in the partition $\hat{\mathcal{P}}$. Without loss of generality, assume $\hat{P}_1$ is this minimal piece. Therefore, we have:

$$\hat{v}_i(\hat{X}_i) \geq \alpha \hat{v}_i(\hat{P}_1). \tag{17}$$

Let $R = X_i \cap P_1$, which implies $\hat{R} = \hat{X}_i \cap \hat{P}_1$. By the properties of the bijection π_i, items mapped into X_i weakly improve in rank, while items mapped outside X_i weakly degrade. This yields the following inequalities:

- Inside X_i: $v_i(R) \geq \hat{v}_i(\hat{R})$ and $v_i(X_i \setminus R) \geq \hat{v}_i(\hat{X}_i \setminus \hat{R})$.
- Outside X_i: Because $P_1 \setminus R \cap X_i = \emptyset$, it follows that $v_i(P_1 \setminus R) \leq \hat{v}_i(\hat{P}_1 \setminus \hat{R})$.

We now apply these bounds to the fairness condition. By additivity of valuations:

$$\begin{aligned} &\hat{v}_i(\hat{X}_i \setminus \hat{R}) + \hat{v}_i(\hat{R}) \geq \alpha\hat{v}_i(\hat{P}_1 \setminus \hat{R}) + \alpha\hat{v}_i(\hat{R}) \\ \implies &\hat{v}_i(\hat{X}_i \setminus \hat{R}) + (1-\alpha)\hat{v}_i(\hat{R}) \geq \alpha\hat{v}_i(\hat{P}_1 \setminus \hat{R}). \end{aligned}$$

Since $0 < \alpha \leq 1$, the coefficient $(1-\alpha)$ is non-negative. We can therefore substitute $\hat{v}_i(\hat{R})$ with the larger term $v_i(R)$ while preserving the inequality:

$$\hat{v}_i(\hat{X}_i \setminus \hat{R}) + (1-\alpha)v_i(R) \geq \alpha\hat{v}_i(\hat{P}_1 \setminus \hat{R}).$$

Adding $\alpha v_i(R)$ to both sides, we obtain:

$$\hat{v}_i(\hat{X}_i \setminus \hat{R}) + v_i(R) \geq \alpha\hat{v}_i(\hat{P}_1 \setminus \hat{R}) + \alpha v_i(R).$$

Finally, we substitute the remaining ordered valuations with their original instance bounds:

$$\begin{aligned} v_i(X_i \setminus R) + v_i(R) &\geq \alpha\hat{v}_i(\hat{P}_1 \setminus \hat{R}) + \alpha v_i(R) && (\text{since } v_i(X_i \setminus R) \geq \hat{v}_i(\hat{X}_i \setminus \hat{R})) \\ &\geq \alpha v_i(P_1 \setminus R) + \alpha v_i(R) && (\text{since } v_i(P_1 \setminus R) \leq \hat{v}_i(\hat{P}_1 \setminus \hat{R})) \\ \implies v_i(X_i) &\geq \alpha v_i(P_1). && (\text{by additivity}) \end{aligned}$$

Because P_1 was chosen such that $\hat{v}_i(\hat{X}_i) \geq \alpha\hat{v}_i(\hat{P}_1)$, and our derivation proves $v_i(X_i) \geq \alpha v_i(P_1)$, it trivially holds that:

$$v_i(X_i) \geq \alpha \min_{P \in \mathcal{P}} v_i(P).$$

Since the k-partition $\mathcal{P}$ of the combined bundle S was chosen arbitrarily, this lower bound holds for all possible k-partitions. Therefore, agent i receives at least their α-Maximin Share of the bundle S. Because this holds for an arbitrary valid subset $J \subseteq [n]$, the allocation Y^i satisfies the requirements for an α-EPMMS (if $|J| = 2$) and an α-EGMMS (if $|J|$ is arbitrary) certificate. □

G Non-Degenerate Instances

G.1 The Envy Cycle Elimination Does Not Work in Degenerate Instances

We show that the non-degeneracy assumption on the input valuations to Algorithm 1 is necessary; without it, the $4/5$-PMMS guarantee fails. We first show that without additional assumptions on valuations (beyond identical rankings), Envy Cycle Elimination can guarantee at most $2/3$-PMMS; this example includes zero-valued items. We then show that ruling out zero-valued items is still insufficient to recover the $4/5$-PMMS guarantee: we provide an example without zero-valued items in which Envy Cycle Elimination can guarantee at most $3/4$-PMMS.

We begin with the first example, which permits zero-valued items.

Observation G.1. *Without the assumption that the valuations are non-degenerate, the Envy Cycle Elimination algorithm (Algorithm 1) does not guarantee better than $2/3$-PMMS on additive valuations with identical rankings.*

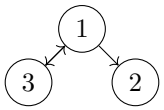
Proof. Fix $\epsilon > 0$, and consider an instance with 3 agents and 7 items, with the following additive valuations:

	g_1	g_2	g_3	g_4	g_5	g_6	g_7
Agent 1	2	2	2	1	1	ϵ	0
Agent 2	2	2	2	1	1	ϵ	0
Agent 3	1	0	0	0	0	0	0

Note that all three agents share the same ranking $g_1, g_2, \ldots, g_7$, so this is an identical-rankings instance. Consider the following execution of Algorithm 1:

1. Agent 1 receives item g_1. Agent 1 is unenvied before this assignment since her bundle is empty.
2. Agent 2 receives items g_2 and g_3. Agent 2 is unenvied before each of these assignments: initially her bundle is empty, and after receiving g_2, agent 1 values $\{g_2\}$ at 2 (equal to her own bundle) and agent 3 values $\{g_2\}$ at 0, so neither envies agent 2.
3. Agent 3 receives items g_4, g_5, and g_6. Before each assignment, agent 3 is unenvied since agents 1 and 2 value any subset of $\{g_4, g_5, g_6\}$ at no more than their own bundles (2 and 4, respectively).

At this point, the envy graph, whose nodes are the agents and which contains an edge $i \to j$ whenever $v_i(X_j) > v_i(X_i)$, is as follows:



No agent has in-degree 0, so the algorithm cannot proceed without resolving the cycle $1 \to 3 \to 1$, which it does by swapping the bundles of agents 1 and 3. After the swap, agent 1 holds $\{g_4, g_5, g_6\}$, valued at $2 + \epsilon$, while agent 2 holds $\{g_2, g_3\}$. We calculate:

$$\mu_1(2, \{g_2, g_3, g_4, g_5, g_6\}) \geq \min\{2 + 1 + \epsilon, 2 + 1\} = 3.$$

As $\epsilon \to 0$, the ratio $\frac{2+\epsilon}{3} \to \frac{2}{3}$, completing the proof. □

We now turn to the second example, which shows that excluding zero-valued items still does not suffice to recover the $4/5$-PMMS guarantee.

Observation G.2. *Even when all item values are strictly positive, the Envy Cycle Elimination algorithm (Algorithm 1) does not guarantee better than $3/4$-PMMS on additive valuations with identical rankings.*

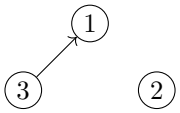
Proof. Assume there are three agents and seven items. Let $\epsilon > 0$ be a sufficiently small constant. The valuation functions are given by:

	g_1	g_2	g_3	g_4	g_5	g_6	g_7
Agent 1	3	3	2	2	1	ϵ	ϵ
Agent 2	3	3	2	2	1	ϵ	ϵ
Agent 3	10	1	1	1	1	1	ϵ

Note that all three agents share the same ranking $g_1, g_2, \ldots, g_7$, so this is an identical-rankings instance. Consider the following execution of Algorithm 1:

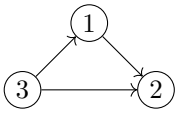
After the first three iterations, every agent holds an item:

$$X_1 = \{g_1\}, \quad v_1(X_1) = 3,$$
$$X_2 = \{g_2\}, \quad v_2(X_2) = 3,$$
$$X_3 = \{g_3\}, \quad v_3(X_3) = 1$$



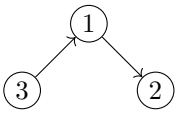
Next, agent 2 receives an item:

$$X_1 = \{g_1\}, \quad v_1(X_1) = 3,$$
$$X_2 = \{g_2, g_4\}, \quad v_2(X_2) = 5,$$
$$X_3 = \{g_3\}, \quad v_3(X_3) = 1$$



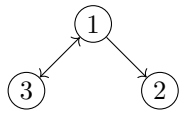
Next, agent 3 receives an item:

$$X_1 = \{g_1\}, \quad v_1(X_1) = 3,$$
$$X_2 = \{g_2, g_4\}, \quad v_2(X_2) = 5,$$
$$X_3 = \{g_3, g_5\}, \quad v_3(X_3) = 2$$



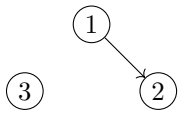
Next, agent 3 receives another item:

$$\begin{aligned} X_1 &= \{g_1\}, & v_1(X_1) &= 3, \\ X_2 &= \{g_2, g_4\}, & v_2(X_2) &= 5, \\ X_3 &= \{g_3, g_5, g_6\}, & v_3(X_3) &= 3 \end{aligned}$$



At this point, a cycle is created and we resolve it:

$$\begin{aligned} X_1 &= \{g_3, g_5, g_6\}, & v_1(X_1) &= 3 + \epsilon, \\ X_2 &= \{g_2, g_4\}, & v_2(X_2) &= 5, \\ X_3 &= \{g_1\}, & v_3(X_3) &= 10 \end{aligned}$$



Finally, suppose agent 3 receives the last item. The algorithm terminates with this allocation.

Observe that:

$$v_1(X_1) = 3 + \epsilon,$$

while

$$\mu_1(2, X_1 \cup X_2) = \min\left(v_1(\{g_2, g_5\}),\ v_1(\{g_3, g_4, g_6\})\right) = \min\{4, 4 + \epsilon\} = 4.$$

Thus, agent 1 cannot obtain more than $\frac{3+\epsilon}{4} \to \frac{3}{4}$ of her pairwise maximin share as $\epsilon \to 0$, completing the proof. □

G.2 Reduction to Non-Degenerate Instances

We adopt the following definition of non-degenerate instances from Chaudhury et al. [28].

Definition G.3 (Non-Degenerate Instances, [28])**.** *An instance is called* non-degenerate *if no agent assigns equal value to two distinct sets; that is, for every $i \in N$, $v_i(S) \neq v_i(T)$ whenever $S \neq T$.*

We now show that it suffices to restrict attention to non-degenerate instances.

In the following proposition, we show that to establish Proposition 4.1, it is enough to consider non-degenerate instances. Chaudhury et al. [28] proved an analogous reduction for EFX, and we adapt their argument to our setting.

Proposition G.4 (Reduction to Non-Degenerate Instances)**.** *If a polynomial-time algorithm exists for computing a 4/5-PMMS and EFX allocation on every instance with* non-degenerate *additive valuations and identical rankings, then it can be extended to handle every instance with (possibly degenerate) additive valuations and identical rankings.*

Proof. Fix $\alpha \leq 1$, and let $M = \{g_1, g_2, \ldots, g_m\}$ be ordered from least to most valuable. Given any instance $\mathcal{I}$ with valuation profile $(v_1, \ldots, v_n)$, we construct a perturbed instance $\mathcal{I}(\varepsilon)$ with valuation profile $(v_1', \ldots, v_n')$, where each v_i' is defined by

$$v_i'(g_j) = v_i(g_j) + \varepsilon \cdot 2^j.$$

For any $k < \ell$ and any $i \in N$, we have

$$v_i'(g_k) = v_i(g_k) + \varepsilon 2^k < v_i(g_\ell) + \varepsilon 2^\ell = v_i'(g_\ell).$$

Hence $\mathcal{I}(\varepsilon)$ preserves the original ordering of items for every agent, so the perturbed instance also has identical rankings.

For every $i \in N$, define

$$\delta_1(i) = \min_{S,T:\alpha v_i(S) \neq v_i(T)} |\alpha v_i(S) - v_i(T)|, \quad \delta_2(i) = \min_{S,T:v_i(S) \neq v_i(T)} |v_i(S) - v_i(T)|.$$

Set $\delta_1 = \min_i \delta_1(i)$, $\delta_2 = \min_i \delta_2(i)$, and $\delta = \min(\delta_1, \delta_2)$. Choose $\varepsilon > 0$ small enough that

$$\varepsilon \cdot 2^{m+1} < \delta.$$

We first show that for any agent i and any $S, T \subseteq M$,

$$v_i(S) > v_i(T) \Longrightarrow v_i'(S) > v_i'(T). \tag{18}$$

Suppose $v_i(S) > v_i(T)$. Then

$$\begin{aligned} v_i'(S) - v_i'(T) &= v_i(S) - v_i(T) + \varepsilon \left(\sum_{g_j \in S} 2^j - \sum_{g_j \in T} 2^j \right) \\ &\geq \delta_2 - \varepsilon \sum_{g_j \in T} 2^j \\ &\geq \delta - \varepsilon(2^{m+1} - 1) \\ &> 0. \end{aligned}$$

By the same argument, for any agent i and any $S, T \subseteq M$,

$$\alpha v_i(S) > v_i(T) \Longrightarrow \alpha v_i'(S) > v_i'(T). \tag{19}$$

Indeed, if $\alpha v_i(S) > v_i(T)$, then

$$\begin{aligned} \alpha v_i'(S) - v_i'(T) &= \alpha v_i(S) - v_i(T) + \varepsilon \left(\sum_{g_j \in S} \alpha 2^j - \sum_{g_j \in T} 2^j \right) \\ &\geq \delta_1 - \varepsilon \sum_{g_j \in T} 2^j \\ &\geq \delta - \varepsilon(2^{m+1} - 1) \\ &> 0. \end{aligned}$$

We next verify that $\mathcal{I}(\varepsilon)$ is non-degenerate. Take any two distinct sets $S, T \subseteq M$ and any agent $i \in N$. If $v_i(S) \neq v_i(T)$, then $v_i'(S) \neq v_i'(T)$ by eq. (18). If instead $v_i(S) = v_i(T)$, then

$$v_i'(S) - v_i'(T) = \varepsilon \left(\sum_{g_j \in S \setminus T} 2^j - \sum_{g_j \in T \setminus S} 2^j \right) \neq 0,$$

since $S \neq T$ and the sums of distinct subsets of powers of two differ. Hence $\mathcal{I}(\varepsilon)$ is non-degenerate.

It remains to show that any α-PMMS allocation in $\mathcal{I}(\varepsilon)$ is also an α-PMMS allocation in $\mathcal{I}$. Suppose, for contradiction, that X is an α-PMMS allocation in $\mathcal{I}(\varepsilon)$ but not in $\mathcal{I}$. Then there exist agents i, j such that $\alpha\mu_i(2, X_i \cup X_j) > v_i(X_i)$, so there is a partition (A_1, A_2) of $X_i \cup X_j$ satisfying $\alpha v_i(A_1) > v_i(X_i)$ and $\alpha v_i(A_2) > v_i(X_i)$. Applying eq. (19), we obtain $\alpha v_i'(A_1) > v_i'(X_i)$ and $\alpha v_i'(A_2) > v_i'(X_i)$, which gives

$$\alpha\mu_{v_i'}(2, X_i \cup X_j) \geq \alpha \min(v_i'(A_1), v_i'(A_2)) > v_i'(X_i),$$

contradicting the assumption that X is an α-PMMS allocation in $\mathcal{I}(\varepsilon)$.

Therefore, whenever every non-degenerate instance admits an α-PMMS allocation, any instance $\mathcal{I}$ can be handled by passing to its non-degenerate perturbation $\mathcal{I}(\varepsilon)$, and any α-PMMS allocation in $\mathcal{I}(\varepsilon)$ is also an α-PMMS allocation in $\mathcal{I}$. □

Remark G.5 (Polynomial-time implementation)**.** *The non-degeneracy assumption can be avoided by specifying tie-breaking in Algorithm 1. Fix a common order $g_1, \ldots, g_m$ satisfying $v_i(g_1) \geq \cdots \geq v_i(g_m)$ for every agent i, and allocate items in this order. Among sources, prioritize empty bundles, then singleton bundles, choosing the singleton with the largest item index; resolve remaining ties arbitrarily. Under these rules, the algorithm still terminates in polynomial time.*

Non-degeneracy is used twice in the proof of proposition 4.1. Consider the case $v_i(g) > \frac{1}{2} v_i(X_i)$ and the moment when g is added to the singleton $\{g'\}$, forming $X_j = \{g', g\}$. At that moment, some

agent i' holds the bundle $\tilde{X}_{i'} \subseteq X_i$ consisting of the already allocated items of X_i. The first use of non-degeneracy establishes that $|\tilde{X}_{i'}| \geq 1$, and the second establishes

$$v_i(h) \geq v_i(g') \geq v_i(g),$$

where $\tilde{X}_{i'} = \{h\}$.

Both uses of non-degeneracy can be replaced by the tie-breaking rules above. For the first, observe that every empty bundle is a source. Hence, whenever an item is added to a nonempty bundle, all bundles are nonempty. This establishes $|\tilde{X}_{i'}| \geq 1$.

For the second, write $\tilde{X}_{i'} = \{h\}$. We claim that h precedes g' in the common order. Suppose instead that h occurs later than g'. Since the holder of $\{g'\}$ is a source, every agent weakly prefers her own bundle to $\{g'\}$. Moreover, every agent values h at most as much as g', by the common ordering. Thus, no agent envies $\{h\}$ either, so its holder is also a source. However, the tie-breaking rule would then prefer $\{h\}$ to $\{g'\}$, a contradiction.

Therefore, h precedes g'. Since both items were allocated before g, we obtain

$$v_i(h) \geq v_i(g') \geq v_i(g).$$

These weak inequalities suffice for $\mu_i(2, \{g, g', h\}) \leq v_i(h)$, so the remainder of the proof applies unchanged. Thus, the same guarantees hold without computing a non-degenerate perturbation.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: As stated in the abstract and introduction, the paper introduces a new fairness notion and proves several theorems about it.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper presents a series of mathematical theorems with formal proofs in the appendix, and concludes with a list of open problems.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The model and all assumptions are clearly stated in the Model and Preliminaries section, and each theorem explicitly states the assumptions it relies on.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [N/A]

   Justification: The paper does not include experimental results.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [N/A]

Justification: The paper does not include any data or code.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [N/A]

   Justification: The paper does not include experimental results.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [N/A]

   Justification: The paper does not include experimental results.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [N/A]

   Justification: The paper does not include experimental results.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: This is a purely theoretical work that does not involve any experiments, data, or code. The authors confirm that the work conforms with the Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: This paper investigates a fair division problem with potential applications that could yield positive societal benefits in resource allocation settings. Since the contribution is purely theoretical, negative societal impacts are unlikely.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: This work is purely theoretical and does not involve any experiments, data, or code.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [N/A]

    Justification: This work is purely theoretical and does not involve any experiments, data, or code.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: This work is purely theoretical and does not involve any experiments, data, or code.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: This work does not involve any experiments.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: This work does not involve any experiments.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLMs were used solely for writing, editing, and formatting purposes.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.