

# Probing EFX via PMMS: (Non-)Existence Results in Discrete Fair Division

Jarosław Byrka<sup>1</sup>, Franciszek Malinka<sup>1</sup>, Tomasz Ponitka<sup>2</sup>

<sup>1</sup>University of Wrocław

<sup>2</sup>Tel Aviv University

## Abstract

We study the fair division of indivisible items and provide new insights into the EFX problem, which is widely regarded as the central open question in fair division, and the PMMS problem, a strictly stronger variant of EFX. Our first result constructs a three-agent instance with two monotone valuations and one additive valuation in which no PMMS allocation exists. Since EFX allocations are known to exist under these assumptions, this establishes a formal separation between EFX and PMMS.

We prove existence of fair allocations for three important special cases. We show that EFX allocations exist for personalized bivalued valuations, where for each agent  $i$  there exist values  $a_i > b_i$  such that agent  $i$  assigns value  $v_i(\{g\}) \in \{a_i, b_i\}$  to each good  $g$ . We establish an analogous existence result for PMMS allocations when  $a_i$  is divisible by  $b_i$ . We also prove that PMMS allocations exist for binary-valued MMS-feasible valuations, where each bundle  $S$  has value  $v_i(S) \in \{0, 1\}$ . Notably, this result holds even without assuming monotonicity of valuations and thus applies to the fair division of chores and mixed manna. Finally, we study a class of valuations called pair-demand valuations, which extend the well-studied unit-demand valuations to the case where each agent derives value from at most two items, and we show that PMMS allocations exist in this setting. Our proofs are constructive, and we provide polynomial-time algorithms for all three existence results.

**Extended version** — <https://arxiv.org/abs/2507.14957>

## 1 Introduction

The fair division problem involves allocating a set of resources among a group of agents to meet certain fairness criteria. When there are only two agents, the standard fair division protocol is the *cut-and-choose* method—where one agent proposes a division of the resources into two parts, and the other selects the preferred part—a practice that dates back to the story of Abraham and Lot in the Book of Genesis. The fair division literature, beginning with Steinhaus (1948), explores the fairness guarantees that can be achieved for more than two agents.

The initial focus of this literature has been on the cake-cutting scenario, where the resources to be allocated are assumed to be infinitely divisible. A key property of cut-and-choose in cake-cutting scenarios with two agents, formalized by Foley (1967), is *envy-freeness*: no agent prefers another agent’s allocation to their own. A central achievement of the cake-cutting literature is that this guarantee can be extended to any number of agents, under mild assumptions (Stromquist 1980).

This work focuses on the discrete fair division problem, where the goods are *indivisible*. Specifically, we consider a setting with  $n$  agents and  $m$  items, where each agent  $i \in N$  has a valuation function  $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$  that assigns a non-negative value to each subset of the items. The goal is to compute a fair partition  $(X_1, \dots, X_n)$  such that  $X_i \cap X_j = \emptyset$  for all  $i \neq j$  and  $X_1 \cup \dots \cup X_n = M$ . It is easy to see that even for two agents, an envy-free allocation may not exist—for instance, when there is only one item.

**EFX.** A key property of cut-and-choose in discrete settings with two agents, identified by Caragiannis et al. (2019), is *envy-freeness up to any good* (EFX): no agent prefers another’s allocation after the removal of any single item from that bundle, i.e., for all  $i, j \in N$ , it holds that  $v_i(X_i) \geq v_i(X_j \setminus \{g\})$  for all  $g \in X_j$ . Extending the existence of EFX allocations to an arbitrary number of agents—analogueous to the existence of envy-free allocations in cake cutting—remains a major open problem, even for additive valuations; notably, no counterexamples are known, even with arbitrary monotone valuations.

Nevertheless, significant progress has been made on the EFX problem for various special cases. Amanatidis et al. (2021) addressed the case of *bivalued valuations*, where each agent assigns one of two values to each item, i.e., there are  $a, b \geq 0$  such that  $v_i(\{g\}) \in \{a, b\}$  for all  $i \in N$  and  $g \in M$ , and  $v_i(S) = \sum_{g \in S} v_i(\{g\})$  for all  $S \subseteq M$ . They provided two distinct proofs of the existence of EFX allocations for bivalued valuations: (1) an existential proof based on maximizing Nash welfare (i.e., the product of agents’ valuations), and (2) a constructive proof via the Match-and-Freeze algorithm. Both proofs crucially rely on the parameters  $a$  and  $b$  being fixed across all agents’ valuations.

Furthermore, Chaudhury, Garg, and Mehlhorn (2024); Akrami et al. (2025) established the existence of EFX allocations for *three agents* where two agents have arbitrary

monotone valuations and the third agent’s valuation satisfies a natural condition known as MMS-feasibility, which covers important classes of valuations, including additive; see Definition 2.6. See also the extended version for additional discussion on the EFX problem.

**PMMS.** While EFX and its relaxations have received significant attention in the fair division literature, Caragiannis et al. (2019) also formalized another compelling fairness criterion known as the *pairwise maximin share* (PMMS), a natural generalization of the guarantees of cut-and-choose. Specifically, for any pair of agents  $i$  and  $j$ , it requires that agent  $i$  receives at least as much value as they would in the maximin partition of their combined items, i.e., the partition  $(A, B)$  of  $X_i \cup X_j$  that maximizes  $\min\{v_i(A), v_i(B)\}$ . This corresponds to the value one could guarantee as the cutter in the cut-and-choose protocol.

While PMMS allocations may not exist for arbitrary monotone valuations (see Proposition 3.1), it remains an open question whether they exist under MMS-feasible valuations, including additive ones. Notably, under mild assumptions, any PMMS allocation is also EFX, so proving the existence of PMMS allocations would simultaneously resolve the EFX problem; see the extended version for more details. However, PMMS is significantly less understood, and little is known about its existence in special cases.

An important advantage of PMMS over EFX is that its definition naturally extends to the fair division of *chores*, where agents have negative valuations that decrease as the number of items increases. In contrast, the definition of EFX must be adapted for chores by requiring that  $v_i(X_i \setminus \{g\}) \geq v_i(X_j)$  for all  $g \in X_i$ . Moreover, Christoforidis and Santorinaios (2024) recently showed that EFX for chores does not exist when there are three agents with arbitrary monotonically decreasing valuations; notably, these valuations violate MMS-feasibility.

The goal of this work is to advance our understanding of the existence of both EFX and PMMS allocations in key special cases of discrete fair division.

## Our Contribution

In this work, we provide new insights into the existence of EFX and PMMS allocations, including one non-existence result and three existence results.

We begin by presenting our negative result for PMMS in the case of three agents.

**Theorem 1** (Separation between EFX and PMMS for Three Agents). *There exists an instance with three agents, one with an additive valuation and two with arbitrary monotone valuations, for which no PMMS allocation exists.*

This result establishes a formal separation between the existence of EFX and PMMS allocations for the case of three agents. Akrami et al. (2025) have shown that EFX is always guaranteed under the assumptions of the theorem. In contrast, we show that PMMS may fail to exist. This indicates that proving the existence of PMMS for the case of three agents might require fundamentally different techniques.

Our first positive result concerns the existence of EFX in the important case of *personalized* bivalued valuations (e.g.,

Aziz et al. 2023; Bu et al. 2023; Amanatidis et al. 2025), where each agent  $i \in N$  has values  $a_i, b_i \geq 0$  such that  $v_i(\{g\}) \in \{a_i, b_i\}$  for all  $g \in M$ . We also consider the existence of PMMS allocations under *factored* personalized bivalued valuations (e.g., Ebadian, Peters, and Shah 2022; Akrami et al. 2022), where each  $a_i$  is an integer multiple of  $b_i$ . We establish the following theorem for this setting.

**Theorem 2** (Personalized Bivalued Valuations). *For any instance with personalized bivalued valuations, there exists an allocation that satisfies EFX. Moreover, if the valuations are also factored, then this allocation is guaranteed to satisfy PMMS.*

This generalizes the result of Amanatidis et al. (2021) in two ways. First, we consider the more general setting of personalized bivalued valuations. Second, we also establish the existence of a PMMS allocation under an additional assumption, rather than just an EFX allocation.

Our result is obtained by adapting the Match-and-Freeze algorithm to the setting of personalized bivalued valuations (see Algorithm 1). Notably, we also show that the argument based on maximizing Nash welfare does not extend to this setting (see Proposition 4.6).

We note that the existence of EFX allocations for personalized bivalued valuations has also been independently discovered by Jin and Tao (2025).

Our second positive result addresses the natural class of *binary-valued valuations*, where each bundle of items is either desirable or non-desirable to an agent; i.e.,  $v_i(S) \in \{0, 1\}$  for all  $i \in N$  and  $S \subseteq M$ . Binary-valued valuations are sometimes referred to as dichotomous preferences in the literature (e.g. Bogomolnaia and Moulin 2004; Bogomolnaia, Moulin, and Stong 2005; Kurokawa, Procaccia, and Shah 2018). We prove the following:

**Theorem 3** (Binary-Valued Valuations). *For any instance with binary-valued MMS-feasible valuations, there exists a PMMS allocation.*

To prove our result, we introduce a novel Cut-and-Choose-Graph procedure (see Algorithm 2), inspired by the Envy-Graph procedure of Lipton et al. (2004).

A key feature of our proof is that it does not rely on the monotonicity of valuations. Therefore, it applies not only to instances involving goods, but also to those involving chores or mixed manna, i.e., a combination of goods and chores.

Finally, we study the class of pair-demand valuations, where each agent desires to receive at most two items. This class naturally generalizes unit-demand valuations and is a special case of the broader and well-studied class of  $k$ -demand (or multi-demand) valuations (e.g., Berger, Eden, and Feldman 2020; Zhang and Conitzer 2020; Deligkas, Melissourgos, and Spirakis 2021), with  $k = 2$ . We establish the following theorem for this setting.

**Theorem 4** (Pair-Demand Valuations). *For any instance with pair-demand valuations, there exists a PMMS allocation.*

All of our positive results are not only existential but also constructive: every algorithm presented in this work runs in polynomial time (see Corollaries 4.5, 5.3 and 6.1).

## 2 Preliminaries

In this work we consider instances of the fair division problem where  $N$  is a set of  $n$  agents,  $N = \{1, \dots, n\}$ , and  $M$  is a set of  $m$  indivisible items,  $M = \{1, \dots, m\}$ . Each agent  $i \in N$  has a valuation function  $v_i : 2^M \rightarrow \mathbb{R}$ .

Unless stated otherwise, we assume that each agent  $i$ 's valuation is normalized, so that  $v_i(\emptyset) = 0$ , and monotone, so that  $v_i(S) \leq v_i(T)$  for any  $S \subseteq T \subseteq M$ . Notably, we do not make these assumptions in Section 5.

An allocation  $X = \langle X_1, \dots, X_n \rangle$  is a partition of the items among the agents, where each item is assigned to exactly one agent. Formally, we require that  $X_1 \cup \dots \cup X_n = M$  and  $X_i \cap X_j = \emptyset$  for all  $i \neq j$ .

### Fairness Notions

We now formally define the EFX condition.

**Definition 2.1** (EFX Allocations). *We say that an allocation  $X$  is envy-free up to any good (EFX) if, for every pair of agents  $i$  and  $j$ , and for every item  $g \in X_j$ , we have  $v_i(X_i) \geq v_i(X_j \setminus \{g\})$ . Otherwise, if there exists some item  $g \in X_j$  such that  $v_i(X_i) < v_i(X_j \setminus \{g\})$ , we say that agent  $i$  EFX-envies agent  $j$ .*

To introduce the notion of PMMS, we first define the fair share function  $\mu_i$ .

**Definition 2.2** (Fair Share). *Let  $k \in \{2, \dots, n\}$ . The fair share of agent  $i$  for a bundle  $S \subseteq M$  divided among  $k$  agents is defined as*

$$\mu_i(S, k) = \max_{X_1, \dots, X_k \text{ partition of } S} \min\{v_i(X_1), \dots, v_i(X_k)\}.$$

*We refer to any partition  $X_1, \dots, X_k$  that achieves the maximum above as a maximin partition. If only a single agent is being considered and it is clear from context, we omit the subscript and write  $\mu$  instead of  $\mu_i$ . Similarly, when  $k = 2$ , we omit the parameter and write  $\mu_i(S)$  to denote  $\mu_i(S, 2)$ .*

Using the notion of fair shares, we define PMMS, first introduced by Caragiannis et al. (2019).

**Definition 2.3** (PMMS Allocations). *An allocation  $X$  is pairwise maximin share fair (PMMS) if, for every pair of agents  $i$  and  $j$ , it holds that  $v_i(X_i) \geq \mu_i(X_i \cup X_j)$ . Otherwise, if  $v_i(X_i) < \mu_i(X_i \cup X_j)$ , we say that agent  $i$  PMMS-envies agent  $j$ .*

We also study the related concept of MMS, and compare it to PMMS in Section 3.

**Definition 2.4** (MMS Allocations). *An allocation  $X$  is maximin share fair (MMS) if, for every agent  $i$ , it holds that  $v_i(X_i) \geq \mu_i(M, n)$ .*

### Valuation Classes

An important class of valuations is the class of *additive* valuations, where for every  $S \subseteq M$  and any  $g \in M \setminus S$ , we have  $v(S \cup \{g\}) = v(S) + v(\{g\})$ .

We now formally define personalized bivalued valuations.

**Definition 2.5** (Personalized Bivalued Valuations). *Agent  $i$ 's valuation is said to be personalized bivalued if  $v_i$  is additive and there exist values  $a_i > b_i \geq 0$  such that*

*$v_i(\{g\}) \in \{a_i, b_i\}$  for every item  $g \in M$ . Furthermore, a personalized bivalued valuation is called factored if either  $a_i$  is divisible by  $b_i$ , or  $b_i = 0$ .*

Whenever we write  $a_i/b_i$  and  $b_i = 0$ , we interpret this fraction as a sufficiently large number  $K$ .

We also define the class of MMS-feasible valuations, following the definition of Akrami et al. (2025).

**Definition 2.6** (MMS-Feasible Valuations). *Agent  $i$ 's valuation is MMS-feasible if, for every subset  $S \subseteq M$  and for any two partitions  $(X_1, X_2)$  and  $(Y_1, Y_2)$  of  $S$  such that  $X_1 \cup X_2 = Y_1 \cup Y_2 = S$  and  $X_1 \cap X_2 = Y_1 \cap Y_2 = \emptyset$ , it holds that  $\max\{v_i(X_1), v_i(X_2)\} \geq \min\{v_i(Y_1), v_i(Y_2)\}$ .*

We define binary-valued valuations as follows.

**Definition 2.7** (Binary-Valued Valuations). *Agent  $i$ 's valuation is binary if for every subset  $S \subseteq M$ , it holds that  $v_i(S) \in \{0, 1\}$ .*

Finally, we define the class of pair-demand valuations.

**Definition 2.8** (Pair-Demand Valuations). *Agent  $i$ 's valuation is called pair-demand if there exist non-negative values  $v_{i,1}, \dots, v_{i,m} \geq 0$  such that, for every subset  $S \subseteq M$ , it holds that:*

$$v_i(S) = \max_{T \subseteq S, |T| \leq 2} \sum_{j \in T} v_{i,j}.$$

### Implications of PMMS

We note that, although PMMS does not imply EFX in general, it does imply a slightly weaker variant of EFX. Moreover, PMMS implies the full EFX condition if all valuations are non-degenerate; i.e., if  $v_i(S) \neq v_i(T)$  for all  $S \neq T$  and  $i$ . These implications are discussed in more detail in the extended version.

## 3 Non-Existence of PMMS

In this section, we examine the conditions necessary for the existence of PMMS allocations. We begin by discussing the two-agent case and then proceed to the case of three agents.

In the case of two agents, the PMMS condition is equivalent to MMS (Definition 2.4). It has been noted by Akrami et al. (2025) that MMS allocations for two agent exist when one agent has an MMS-feasible valuation and the other has an arbitrary valuation. This condition is necessary, as an MMS allocation may not exist for two agents if both valuations are arbitrary.

While one might expect that a non-existence result for two agents would automatically extend to settings with three or more agents, this extension is not immediate. For example, simply adding a dummy agent who values every subset at 0 to the two-agent counterexample is not sufficient. Intuitively, adding such a dummy agent is equivalent to allowing some items to go to charity, as considered by Caragiannis, Gravin, and Huang (2019). We present a non-trivial extension of the non-existence result beyond the two-agent case. To the best of our knowledge, this is the first known construction demonstrating the non-existence of a PMMS allocation for more than two agents.

**Proposition 3.1** (Non-Existence of PMMS). *For any  $n \geq 2$ , there exists an instance with  $n$  agents, each with monotone valuations, and  $n + \mathcal{O}(\log n)$  items, in which no PMMS allocation exists. Moreover, an MMS allocation exists in this instance.*

Our construction consists of an instance with  $n - 2$  star items and  $\mathcal{O}(\log n)$  common items. By carefully designing the valuation functions, we ensure that any PMMS allocation must assign exactly one star item to  $n - 2$  agents and partition the common items among the remaining two agents. Furthermore, by appropriately setting the valuations over the common items, we ensure that no pair of agents can partition these items among themselves and satisfy the PMMS condition. The full proof appears in the extended version.

The construction in the proof of Proposition 3.1 uses  $n + \mathcal{O}(\log n)$  items. An interesting open problem is to determine the minimal number of items required for such a construction; specifically, whether the result can be improved to include only  $n + \mathcal{O}(1)$  items.

As shown by Akrami et al. (2025), an EFX allocation always exists for three agents if at least one of them has an MMS-feasible valuation and the other two have monotone valuations. We show that these assumptions are not sufficient to guarantee the existence of a PMMS allocation. Specifically, we present an example with two agents having monotone valuations and one agent with an additive (hence MMS-feasible) valuation, where no PMMS allocation exists. This establishes a separation between the two fairness notions and suggests that the proof technique used to establish the existence of EFX allocations for three agents does not extend to PMMS. The proof of the following theorem is deferred to the extended version.

**Theorem 1** (Separation between EFX and PMMS for Three Agents). *There exists an instance with three agents, one with an additive valuation and two with arbitrary monotone valuations, for which no PMMS allocation exists.*

The instance used in the proof of Theorem 1 was discovered through a computational search over the space of possible valuations. Considerable effort was devoted to simplifying the counterexample, as the initial instances had complex and unintuitive structure. The counterexample presented here is notably simpler, using few distinct values and elegantly handling uneven bundle sizes.

## 4 Personalized Bivalued Valuations

In this section, we analyze the case of personalized bivalued valuations, where the valuations are additive and, for each agent  $i$ , there exist some  $a_i > b_i \geq 0$  such that  $v_i(\{g\}) \in \{a_i, b_i\}$  for all items  $g$ . Our main theorem for this setting is as follows.

**Theorem 2** (Personalized Bivalued Valuations). *For any instance with personalized bivalued valuations, there exists an allocation that satisfies EFX. Moreover, if the valuations are also factored, then this allocation is guaranteed to satisfy PMMS.*

The rest of this section is devoted to proving Theorem 2. We show that Algorithm 1 produces an EFX allocation for

---

### Algorithm 1: Personalized Match-and-Freeze Alg.

---

**Input** : A personalized bivalued instance  $\langle N, M, \mathcal{V} \rangle$   
**Output**: An EFX allocation  $X = \langle X_1, \dots, X_n \rangle$

- 1 Set  $P \leftarrow M$
- 2 Set  $w_i \leftarrow 0$  for  $i \in \{1, \dots, n\}$
- 3 Set  $L_r \leftarrow N$  for  $r \in \{1, \dots, m\}$
- 4 Set  $X_i \leftarrow \emptyset$  for  $i \in \{1, \dots, n\}$
- 5 **while**  $P \neq \emptyset$  **do**
- 6 Set  $r$  to be the current round (iteration) number
- 7 Construct a weighted bipartite graph  $G$  between  $L_r$  and  $P$
- 8 **for each**  $i \in L_r$  **and**  $g \in P$  **with**  $v_i(\{g\}) = a_i$  **do**
- 9 Add an edge from  $i$  to  $g$  with weight  $a_i/b_i$
- 10 Find a maximal (in terms of size) matching in  $G$  that maximizes the total weight
- 11 **for each matched pair**  $(i, g)$  **do**
- 12 Add  $g$  to  $X_i$
- 13 Remove  $g$  from  $P$
- 14 **for each connected component**  $C$  **of**  $G$  **do**
- 15 Let  $U$  be the set of unmatched agents in  $C$
- 16 **if**  $U \neq \emptyset$  **then**
- 17 Let  $t$  be the maximum of  $a_i/b_i$  for  $i \in U$
- 18 Remove all matched agents of  $C$  from  $L_{r+j}$  for  $j \in \{1, \dots, \lfloor t - 1 \rfloor\}$
- 19  $w_i \leftarrow r$  for each matched  $i$  in  $C$
- 20 **for each unmatched agent**  $i$  **in increasing**  $w_i$  **do**
- 21 Add any remaining good  $g$  to  $X_i$
- 22 Remove  $g$  from  $P$
- 23 **return**  $X = \langle X_1, \dots, X_n \rangle$

---

any personalized bivalued instance, and a PMMS allocation if the instance is factored. Algorithm 1 is a modified version of the Match-and-Freeze algorithm proposed by Amanatidis et al. (2021), which was designed to compute EFX allocations for non-personalized bivalued valuations. See Table 1 for an example of the execution of Algorithm 1.

We begin by stating the following property of the allocation produced by the algorithm; see the extended version for a proof.

**Lemma 4.1.** *Consider an execution of Algorithm 1, and fix a round  $r$ . Let  $C$  be any connected component of the graph  $G$  constructed during round  $r$ . If the set  $U$  of unmatched agents in  $C$  is nonempty and  $U \neq C$ , then it holds that:*

$$\max\{a_i/b_i : i \in U\} \leq \min\{a_i/b_i : i \in C \setminus U\}.$$

Next, we use Lemma 4.1 to derive the following technical lemma, which plays a key role in our analysis. The lemma and its proof are similar to the analysis of the Match-and-Freeze algorithm for non-personalized bivalued valuations, as presented by Amanatidis et al. (2021, Lemma 4.2). The proof is deferred to the extended version.

**Lemma 4.2.** *Consider an execution of Algorithm 1, and fix an agent  $i$ . Let  $r_i$  denote the last round in which an item  $g$  with  $v_i(\{g\}) = a_i$  was allocated to some agent. Let  $X_{i,r}$*

Agents	Round 1	Round 2	Round 3	Round 4	Round 5	Round 6
Agent 1 ( $a_1 = 2.5$ )	1	1	1	1	1	1
Agent 2 ( $a_2 = 3$ )	③	—	1	1	1	
Agent 3 ( $a_3 = 4$ )	1	1	1	1	1	1
Agent 4 ( $a_4 = 5$ )	⑤	—	—	—	1	

Table 1: An example execution of Algorithm 1 is presented for an instance with 4 agents and 18 items. The agents have bivalued valuations with  $a_1 = 2.5$ ,  $a_2 = 3$ ,  $a_3 = 4$ ,  $a_4 = 5$ , and  $b_i = 1$  for  $i \in \{1, 2, 3, 4\}$ . The items are  $x, y, z_1, \dots, z_{16}$ . Agents 1 and 2 assign the higher value to item  $x$  and no other item, while agents 3 and 4 assign the higher value to item  $y$  and no other item. Each cell represents the value of the item that a given agent received in a particular round. Cells with a circled value indicate rounds in which an agent received a high-value item. Cells corresponding to rounds in which an agent is frozen contain a dash, and cells for rounds after the last item has been allocated are left empty.

denote the singleton set allocated to agent  $i$  in round  $r$ , or the empty set if agent  $i$  does not receive any item in that round. Let  $F_i$  denote the set of rounds during which agent  $i$  is not active. Then:

- (i) We have  $v_i(X_{i,r}) = a_i$  for all  $r \in \{1, 2, \dots, r_i - 1\}$ .
- (ii) If  $F_i \neq \emptyset$ , then  $v_i(X_{i,r_i}) = a_i$  and  $F_i \subseteq \{r_i + 1, \dots, r_i + \lfloor a_i/b_i - 1 \rfloor\}$ .
- (iii) If  $v_i(X_{i,r_i}) = a_i$  and  $v_i(X_{j,r_i}) = a_i$  for some agent  $j$ , then  $F_i = F_j$ .

The proof of Theorem 2 also relies on the following claim, proved in the extended version.

**Lemma 4.3.** *Let  $X$  be an allocation, and let  $i$  and  $j$  be any agents. Suppose that  $v_i(X_i) \geq v_i(X_j) - b_i$ . Then agent  $i$  does not EFX-envy agent  $j$ . Moreover, if agent  $i$ 's valuation is factored, then agent  $i$  also does not PMMS-envy agent  $j$ .*

We now present a simplified version of the proof of the main theorem, under the additional assumption that the considered agent has received a low-value item. For the full proof, see the extended version.

*Simplified Proof of Theorem 2.* Consider an execution of Algorithm 1, and fix two agents  $i$  and  $j$ . Let  $R$  denote the index of the last round, and let  $r_i \leq R$  denote the last round in which an item  $g$  with  $v_i(\{g\}) = a_i$  was allocated to any agent. Let  $X_{i,r}$  denote the singleton set allocated to agent  $i$  in round  $r$ , or the empty set if agent  $i$  does not receive any item in that round. Similarly, let  $X_{j,r}$  denote the set allocated to agent  $j$  in round  $r$ . Let  $F_i \subseteq \{1, \dots, R\}$  denote the set of rounds during which agent  $i$  is inactive.

Our goal is to prove that  $i$  does not EFX-envy  $j$ , and if the valuation is factored, then  $i$  also does not PMMS-envy  $j$ . In the simplified version of the proof, we assume that  $v_i(X_{i,r_i}) = b_i$ . By Lemma 4.3, it suffices to show that  $v_i(X_i) \geq v_i(X_j) - b_i$ . We use this observation to complete the proof in most of the cases considered in our analysis.

By Lemma 4.2, we know that  $v_i(X_{i,r}) = a_i$  for all  $r \in \{1, \dots, r_i - 1\}$ , and  $v_i(X_{i,r}) \geq b_i$  for all  $r \in \{r_i + 1, \dots, R - 1\}$ . Moreover,  $v_i(X_{j,r}) \leq a_i$  for all  $r \in \{1, \dots, r_i - 1\}$  and  $v_i(X_{j,r}) \leq b_i$  for all  $r \in \{r_i + 1, \dots, R\}$ .

First, suppose that  $v_i(X_{j,r_i}) = b_i$ . In this case, by Lemma 4.2, we have  $F_i = \emptyset$ , and:

$$v_i(X_i) \geq (r_i - 1) \cdot a_i + (R - r_i) \cdot b_i \geq v_i(X_j) - b_i.$$

Next, suppose that  $v_i(X_{j,r_i}) = a_i$ . In this case, by Lemma 4.2, we know that agent  $i$  was not matched during round  $r_i$ , and thus  $F_i = \emptyset$ . Moreover, agent  $j$  must have been matched to item  $X_{j,r_i}$  during round  $r_i$ . Indeed, it cannot be that agent  $j$  received item  $X_{j,r_i}$  through the operation in Line 21, because if that were the case, then agent  $i$  would have been matched to item  $X_{j,r_i}$  during round  $r_i$ . Therefore, during round  $r_i$ , agents  $i$  and  $j$  belong to the same connected component, as they both have an edge to item  $X_{j,r_i}$ . Furthermore, by the choice of freezing time in Line 17, we have  $v_i(X_{j,r}) = 0$  for all  $r \in \{r_i + 1, \dots, r_i + \lfloor a_i/b_i - 1 \rfloor\}$ , since agent  $i$  remains unmatched during round  $r_i$ .

First, suppose that  $R > r_i + \lfloor a_i/b_i - 1 \rfloor$ . Since agent  $j$  becomes frozen, her priority is updated in Line 18, so that agent  $i$  always chooses an item before agent  $j$ . Hence,  $v_i(X_{i,R}) \geq v_i(X_{j,R})$ . It follows that:

$$\begin{aligned} v_i(X_i) &= (r_i - 1) \cdot a_i + (R - r_i) \cdot b_i + v_i(X_{i,R}) \\ &= r_i \cdot a_i + (R - a_i/b_i - r_i) \cdot b_i + v_i(X_{i,R}) \\ &\geq r_i \cdot a_i + (R - 1 - \lfloor a_i/b_i \rfloor - r_i) \cdot b_i + v_i(X_{j,R}) \\ &\geq v_i(X_j) - b_i. \end{aligned}$$

Now, suppose that  $R \leq r_i + \lfloor a_i/b_i - 1 \rfloor$ . In this case,  $X_j$  consists of exactly  $r_i$  items, and  $X_i$  consists of at least  $r_i - 1$  high-value items and possibly some low-value items. The EFX condition holds since  $v_i(X_i) \geq (r_i - 1) \cdot a_i \geq v_i(X_j \setminus \{g\})$  for any  $g \in X_j$ . For the PMMS condition, observe that the total value of low-value items in  $X_i$  is at most  $(\lfloor a_i/b_i - 1 \rfloor + 1) \cdot b_i \leq a_i$ , and therefore:

$$\mu_i(X_i \cup X_j) \leq (r_i - 1) \cdot a_i + (|X_i| - r_i + 1) \cdot b_i = v_i(X_i).$$

This completes the simplified proof of the theorem.  $\square$

In the following remark, we explain the importance of the order in which unmatched agents receive items in Line 20 of Algorithm 1.

**Remark 4.4.** *Consider the example described in Table 1. In the final round, only two items remain, and all agents are active. Therefore, two agents must be left without an item in this round. According to the priority ordering  $w$ , the items are given to agents 1 and 3, as they were not matched at any point during the execution of the algorithm.*

If, instead, an item were allocated to agent 2 rather than agent 1, then agent 1 would EFX-envy agent 2. In that case, agent 1 would receive five low-value items, yielding a value of 5, while agent 2 would receive one high-value item and four low-value items. Agent 1's valuation for agent 2's bundle, even after removing one low-value item, would be 5.5, violating EFX.

We also observe that Algorithm 1 allocates at least one item in each round, ensuring that it runs in polynomial time (in  $n$  and  $m$ ), as a maximum-weight maximal matching can be computed in polynomial time (Kuhn 1955). This leads to the following corollary.

**Corollary 4.5.** *For any instance with personalized bivalued valuations, an EFX allocation can be computed in polynomial time. Furthermore, for any instance with factored personalized bivalued valuations, a PMMS allocation can also be computed in polynomial time.*

Finally, we note that Amanatidis et al. (2021) also showed that for non-personalized bivalued valuations, an EFX allocation can be achieved by maximizing the Nash welfare (i.e., the geometric mean of agents' utilities,  $(\prod_{i \in N} v_i(X_i))^{1/n}$ ). In sharp contrast, in the extended version, we show that this result does not extend to personalized bivalued valuations, as demonstrated by the following proposition.

**Proposition 4.6.** *There exists an instance with personalized bivalued valuations where none of the allocations maximizing Nash welfare satisfies EFX.*

## 5 Binary-Valued Valuations

In this section, we consider the class of binary-valued valuations, where all possible bundles are classified as either desirable or non-desirable; see Definition 2.7 for a formal definition. Importantly, we do not assume that these valuations are monotone or normalized, we only assume that they satisfy the MMS-feasibility condition (Definition 2.6). We now proceed to state our main theorem.

**Theorem 3** (Binary-Valued Valuations). *For any instance with binary-valued MMS-feasible valuations, there exists a PMMS allocation.*

The proof of Theorem 3 is based on the notion of the cut-and-choose graph, which is formally defined in the following definition. This notion is inspired by the envy graph introduced by Lipton et al. (2004).

**Definition 5.1** (Cut-and-Choose Graph). *For a fixed allocation  $(X_1, \dots, X_n)$  and a fixed agent  $s \in N$ , we define the cut-and-choose graph as follows. The set of vertices is the set of agents  $N$ . For every agent  $i \in N$ , we add an edge to precisely one agent  $\pi(i) \in N$ . If  $v_i(X_s) \geq \mu_i(X_s \cup X_j)$  for all  $j \in N$ , then we set  $\pi(i) = s$ . Otherwise, we set  $\pi(i) = j$  for some  $j \in N$  such that  $v_i(X_s) < \mu_i(X_s \cup X_j)$ .*

We use the following key properties of the cut-and-choose graph.

**Lemma 5.2.** *Fix an allocation  $(X_1, \dots, X_n)$  and an agent  $s \in N$ , and consider the cut-and-choose graph. Then, for every agent  $i \in N$ , the following holds:*

---

### Algorithm 2: Cut-and-Choose-Graph Procedure

---

**Input** : A binary-valued MMS-feasible  $\langle N, M, \mathcal{V} \rangle$   
**Output**: A PMMS allocation  $X = \langle X_1, \dots, X_n \rangle$

- 1 Let  $X = \langle X_1, \dots, X_n \rangle$  be an arbitrary allocation
- 2 **while**  $X$  is not PMMS **do**
- 3     Choose  $s \in [n]$  s.t.  $\exists j \in N$   $v_s(X_s) < \mu_s(X_s \cup X_j)$
- 4     Let  $\pi$  be the cut-and-choose graph for  $X$  and  $s$
- 5     Set  $i_0 \leftarrow s$
- 6     Set  $k \leftarrow 0$
- 7     **while**  $\pi(i_k) \notin \{i_0, \dots, i_k\}$  **do**
- 8         Set  $i_{k+1} \leftarrow \pi(i_k)$
- 9         Set  $k \leftarrow k + 1$
- 10    **if**  $\pi(i_k) = i_0$  **then**
- 11        Set  $X_i \leftarrow X_{\pi(i)}$  for all  $i \in \{i_0, \dots, i_k\}$
- 12    **else**
- 13        Let  $w$  be such that  $\pi(i_k) = i_w$
- 14        Let  $(A, B)$  be a maximin partition of  $X_{i_0} \cup X_{i_w}$  by  $v_{i_k}$
- 15        **if**  $v_{i_{w-1}}(A) < v_{i_{w-1}}(B)$  **then**
- 16            Swap  $A$  and  $B$
- 17        Set  $X_i \leftarrow X_{\pi(i)}$  for all  $i \in \{i_0, \dots, i_{w-2}\} \cup \{i_w, \dots, i_{k-1}\}$
- 18        Set  $X_{i_{w-1}} \leftarrow A$
- 19        Set  $X_{i_k} \leftarrow B$
- 20 **return**  $X = \langle X_1, \dots, X_n \rangle$

---

- (i) If  $\pi(i) = s$ , then either  $v_i(X_s) = 1$ , or  $\mu_i(X_s \cup X_j) = 0$  for all  $j \in N$ .
- (ii) If  $\pi(i) = j$  for some  $j \neq s$ , then  $v_i(X_j) = 1$  and  $\mu_i(X_s \cup X_j) = 1$ .

We use Lemma 5.2 to prove Theorem 3. The proof relies on Algorithm 2, and Figure 1 illustrates the operations performed by the algorithm.


We present a simplified version of the proof of Theorem 3, under the assumption that the if-condition in Line 10 is not satisfied. For the full proof, see the extended version.

*Simplified Proof of Theorem 3.* By the while-loop condition in Line 2, it is clear that the output of Algorithm 2 satisfies the PMMS condition. Hence, to prove Theorem 3, it suffices to show that Algorithm 2 terminates for any input.


Consider the tuple  $\langle W, E \rangle$ , where  $W = \sum_{i \in N} v_i(X_i)$  and  $E$  is the number of agents for whom the PMMS condition holds. We will show that in each iteration, either  $W$  strictly increases, or  $W$  remains the same and  $E$  strictly increases. Since  $W \in \{0, 1, \dots, n\}$  and  $E \in \{0, 1, \dots, n\}$ , the algorithm must terminate after at most  $n^2$  iterations.

To prove this claim, fix an iteration of the while-loop. Let  $X$  be the allocation at the beginning of the iteration and let  $Y$  be the allocation at the end.

In the simplified version of the proof, we suppose the if-condition in Line 10 is not satisfied. By Lemma 5.2(ii), we have  $v_i(Y_i) = v_i(X_{\pi(i)}) = 1$  for all  $i \in \{i_0, \dots, i_k\} \setminus \{i_{w-1}, i_k\}$ . Also, from the same lemma, we know that  $\mu_{i_{w-1}}(X_{i_w} \cup X_{i_0}) = 1$  and  $\mu_{i_k}(X_{i_w} \cup X_{i_0}) = 1$ .



(i) Case 1: Cycle.



(ii) Case 2: Lollipop.

Figure 1: An illustration of the operations performed by the Cut-and-Choose-Graph Procedure (Algorithm 2). The nodes represent agents, and the arrows depict the cut-and-choose graph. Subfigure (i) illustrates Case (1), where the if-condition in Line 10 holds, and the new allocation is obtained by assigning each agent  $i$  in the cycle the bundle of the agent to whom  $i$  points. Subfigure (ii) illustrates Case (2), where the if-condition in Line 10 does not hold. In this case, the new allocation is obtained by giving all agents in the cycle—except for agents  $i_{w-1}$  and  $i_k$ —the bundles they are pointing to, and having agents  $i_{w-1}$  and  $i_k$  divide the two bundles  $X_s$  and  $X_{i_w}$  between themselves using the cut-and-choose method.

Since  $(A, B)$  is set to be a maximin partition of  $X_{i_w} \cup X_{i_0}$  by  $v_{i_k}$  in Line 14, we have  $v_{i_k}(B) = 1$  since  $\mu_{i_k}(X_{i_w} \cup X_{i_0}) = 1$ . Finally, by the swap operation in Line 16 and the MMS-feasibility of  $v_{i_{w-1}}$ , we have

$$\begin{aligned} v_{i_{w-1}}(Y_{i_{w-1}}) &= v_{i_{w-1}}(A) = \max\{v_{i_{w-1}}(A), v_{i_{w-1}}(B)\} \\ &\geq \mu_{i_{w-1}}(A \cup B) = 1. \end{aligned}$$

Therefore,  $W$  strictly increases in this case. This proves the claim, and consequently, completes the simplified proof of the theorem.  $\square$

As noted in the proof of Theorem 3, the tuple  $\langle W, E \rangle$  increases lexicographically with every iteration of Algorithm 2. Since there are only  $n^2$  possible values for this tuple, Algorithm 2 runs in polynomial time (in  $n$  and  $m$ ). This leads to the following corollary.

**Corollary 5.3.** *For instances with binary-valued valuations, PMMS allocations can be computed in polynomial time.*

## 6 Pair-Demand Valuations

In this section, we analyze the natural class of pair-demand valuations, where each agent desires at most two items; see Definition 2.8. This class generalizes unit-demand valuations, in which every agent wants only one item. While EFX and PMMS are straightforward to guarantee under unit-demand valuations, they become non-trivial in the pair-demand setting. Our main result is the following theorem.

**Theorem 4 (Pair-Demand Valuations).** *For any instance with pair-demand valuations, there exists a PMMS allocation.*

We remark that we do not establish the existence of EFX allocations for pair-demand valuations, which we leave as an open problem. In particular, PMMS does not imply EFX, as valuations may be degenerate. Achieving EFX may require allocating only a single item to one agent and more than two items to another. For example, consider an instance with two identical agents, one high-value item, and three low-value items. Here, one agent must receive the high-value item, while the other receives all three low-value items, even though the third low-value item no longer increases the second agent’s value and could instead benefit the first agent.

---

### Algorithm 3: Reversed Round-Robin Algorithm

---

**Input** : A pair-demand  $\langle N, M, \mathcal{V} \rangle$  with  $|M| \geq 2|N|$   
**Output**: A PMMS allocation  $X = \langle X_1, \dots, X_n \rangle$

- 1 Set  $X_i \leftarrow \emptyset$  for  $i \in \{1, \dots, n\}$
- 2 Set  $P \leftarrow M$
- 3 **for**  $i \in N$  in increasing order of indices **do**
- 4     Let  $g_i \in \arg \max_{g \in P} v_i(\{g\})$
- 5     Add  $g_i$  to  $X_i$
- 6     Remove  $g_i$  from  $P$
- 7 **for**  $i \in N$  in decreasing order of indices **do**
- 8     Let  $h_i \in \arg \max_{h \in P} v_i(\{h\})$
- 9     Add  $h_i$  to  $X_i$
- 10    Remove  $h_i$  from  $P$
- 11 Allocate all items left in  $P$  to any player
- 12 **return**  $X = \langle X_1, \dots, X_n \rangle$

---

As we show below, it is possible to obtain a PMMS allocation while giving each agent exactly two items, provided there are at least  $2n$  items in total.

Our proof is based on a two-stage round-robin algorithm (Algorithm 3) in which agents select items according to a fixed order in the first stage, and then in the reversed order in the second stage. The idea of reversing the order in round-robin algorithms has been explored in the literature; for instance, in ABBA picking sequences (Brams and Taylor 2000), the Double Round-Robin algorithm (Aziz et al. 2022), and as part of the Draft-and-Eliminate algorithm (Amanatidis, Markakis, and Ntokos 2020). We defer the proof to the extended version.

Since Algorithm 3 runs in polynomial time (in  $n$  and  $m$ ), we obtain the following corollary.

**Corollary 6.1.** *For instances with pair-demand valuations, PMMS allocations can be computed in polynomial time.*

## 7 Conclusion

Our work has provided several new insights into the existence of EFX and PMMS allocations. In the extended version, we additionally highlight several open problems that arise from our findings.

## Acknowledgments

This project has been partially funded by NCN grant number 2020/39/B/ST6/01641, by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 866132), by an Amazon Research Award, by the Israel Science Foundation Breakthrough Program (grant No. 2600/24), and by a grant from TAU Center for AI and Data Science (TAD), and by the NSF-BSF (grant number 2020788).

## References

- Akrami, H.; Alon, N.; Chaudhury, B. R.; Garg, J.; Mehlhorn, K.; and Mehta, R. 2025. EFX: A Simpler Approach and an (Almost) Optimal Guarantee via Rainbow Cycle Number. *Oper. Res.*, 73(2): 738–751.
- Akrami, H.; Chaudhury, B. R.; Hoefer, M.; Mehlhorn, K.; Schmalhofer, M.; Shahkarami, G.; Varricchio, G.; Vermande, Q.; and van Wijland, E. 2022. Maximizing Nash Social Welfare in 2-Value Instances. In *AAAI*, 4760–4767. AAAI Press.
- Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2021. Maximum Nash welfare and other stories about EFX. *Theor. Comput. Sci.*, 863: 69–85.
- Amanatidis, G.; Lolos, A.; Markakis, E.; and Turmel, V. 2025. Online Fair Division for Personalized 2-Value Instances. *CoRR*, abs/2505.22174.
- Amanatidis, G.; Markakis, E.; and Ntokos, A. 2020. Multiple birds with one stone: Beating 1/2 for EFX and GMMS via envy cycle elimination. *Theor. Comput. Sci.*, 841: 94–109.
- Aziz, H.; Caragiannis, I.; Igarashi, A.; and Walsh, T. 2022. Fair allocation of indivisible goods and chores. *Auton. Agents Multi Agent Syst.*, 36(1): 3.
- Aziz, H.; Lindsay, J.; Ritossa, A.; and Suzuki, M. 2023. Fair Allocation of Two Types of Chores. In *AAMAS*, 143–151. ACM.
- Berger, B.; Eden, A.; and Feldman, M. 2020. On the Power and Limits of Dynamic Pricing in Combinatorial Markets. In *WINE*, volume 12495 of *Lecture Notes in Computer Science*, 206–219. Springer.
- Bogomolnaia, A.; and Moulin, H. 2004. Random Matching Under Dichotomous Preferences. *Econometrica*, 72(1): 257–279.
- Bogomolnaia, A.; Moulin, H.; and Stong, R. 2005. Collective choice under dichotomous preferences. *J. Econ. Theory*, 122(2): 165–184.
- Brams, S.; and Taylor, A. D. 2000. *The Win-Win Solution: Guaranteeing Fair Shares to Everybody*. W. W. Norton & Company.
- Bu, X.; Li, Z.; Liu, S.; Song, J.; and Tao, B. 2023. Fair Division with Allocator’s Preference. In *WINE*, volume 14413 of *Lecture Notes in Computer Science*, 77–94. Springer.
- Caragiannis, I.; Gravin, N.; and Huang, X. 2019. Envy-Freeness Up to Any Item with High Nash Welfare: The Virtue of Donating Items. In *EC*, 527–545. ACM.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Economics and Comput.*, 7(3): 12:1–12:32.
- Chaudhury, B. R.; Garg, J.; and Mehlhorn, K. 2024. EFX Exists for Three Agents. *J. ACM*, 71(1): 4:1–4:27.
- Christoforidis, V.; and Santorinaios, C. 2024. On the Pursuit of EFX for Chores: Non-existence and Approximations. In *IJCAI*, 2713–2721. ijcai.org.
- Deligkas, A.; Melissourgos, T.; and Spirakis, P. G. 2021. Walrasian Equilibria in Markets with Small Demands. In *AAMAS*, 413–419. ACM.
- Ebadian, S.; Peters, D.; and Shah, N. 2022. How to Fairly Allocate Easy and Difficult Chores. In *AAMAS*, 372–380. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS).
- Foley, D. 1967. Resource Allocation and the Public Sector. *Yale Economic Essays*, 7(1): 45–98.
- Jin, J.; and Tao, B. 2025. On Pareto-Optimal and Fair Allocations with Personalized Bi-Valued Utilities. *CoRR*, abs/2507.18251.
- Kuhn, H. W. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2): 83–97.
- Kurokawa, D.; Procaccia, A. D.; and Shah, N. 2018. Leximin Allocations in the Real World. *ACM Trans. Economics and Comput.*, 6(3-4): 11:1–11:24.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *EC*, 125–131. ACM.
- Steinhaus, H. 1948. The problem of fair division. *Econometrica*, 16(1): 101–104.
- Stromquist, W. 1980. How to Cut a Cake Fairly. *The American Mathematical Monthly*, 87(8): 640–644.
- Zhang, H.; and Conitzer, V. 2020. Learning the Valuations of a k-demand Agent. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, 11066–11075. PMLR.