

Comparing Approximate Relaxations of Envy-Freeness

Georgios Amanatidis¹, Georgios Birmpas² and Evangelos Markakis²

¹ Centrum Wiskunde & Informatica (CWI), Amsterdam, the Netherlands

² Department of Informatics, Athens University of Economics and Business, Athens, Greece
 amanatid@cwi.nl, gebirbas@aueb.gr, markakis@aueb.gr

Abstract

In fair division problems with indivisible goods it is well known that one cannot have any guarantees for the classic fairness notions of envy-freeness and proportionality. As a result, several relaxations have been introduced, most of which in quite recent works. We focus on four such notions, namely envy-freeness up to one good (EF₁), envy-freeness up to any good (EFX), maximin share fairness (MMS), and pairwise maximin share fairness (PMMS). Since obtaining these relaxations also turns out to be problematic in several scenarios, approximate versions of them have been considered. In this work, we investigate further the connections between the four notions mentioned above and their approximate versions. We establish several tight, or almost tight, results concerning the approximation quality that any of these notions guarantees for the others, providing an almost complete picture of this landscape. Some of our findings reveal interesting and surprising consequences regarding the power of these notions, e.g., PMMS and EFX provide the same worst-case guarantee for MMS, despite PMMS being a strictly stronger notion than EFX. We believe such implications provide further insight on the quality of approximately fair solutions.

1 Introduction

In this work, we revisit fairness notions for allocating indivisible goods. The objective in fair division is to allocate a set of resources to a set of agents in a way that leaves everyone satisfied, according to their own preferences. Over the years, the field has grown along various directions, with a substantial literature by now, and with several applications. We refer the reader to the surveys [Bouveret *et al.*, 2016; Procaccia, 2016; Markakis, 2017] for more recent results, and to the classic textbooks [Brams and Taylor, 1996; Robertson and Webb, 1998; Moulin, 2003] for an overview of the area.

To model such allocation problems, one needs to specify the preferences of the agents, and the fairness criterion under consideration. For the preferences, we stick to the usual assumption of associating each agent with an *additive* valuation function on the set of resources. As for fairness criteria,

one of the classic desirable notions that have been proposed is *envy-freeness*, meaning that no agent has a higher value for the bundle of another agent than for her own [Foley, 1967; Varian, 1974]. Unfortunately, for problems with indivisible goods, this turns out to be a very strong requirement. Existence of envy-free allocations cannot be guaranteed, and the relevant algorithmic and approximability questions are also computationally hard [Lipton *et al.*, 2004].

Given these concerns, recent works have considered relaxations of envy-freeness that seem more appropriate for settings with indivisible items, see, e.g., [Budish, 2011; Caragiannis *et al.*, 2016]. Our work focuses on four such notions, namely envy-freeness up to one good (EF₁), envy-freeness up to any good (EFX), maximin share fairness (MMS), and pairwise maximin share fairness (PMMS). All these capture different ways of allowing envy in an allocation, albeit in a restricted way. For instance, EF₁ requires that no agent envies another agent's bundle after removing from it her most valued item.

These relaxations are still no panacea, especially since existence results have either been elusive or simply negative. So far, we only know that EF₁ allocations always exist, whereas this is not true for MMS allocations. As a result, this has naturally led to approximate versions of these notions, accompanied by some positive algorithmic results (see related work). What has not been well charted yet, however, is the relation between these four notions and their approximate counterparts, especially concerning the approximation quality that any of these notions guarantees for the others. For example, does an MMS or an approximate MMS allocation (for which we do know efficient algorithms) provide an approximation guarantee in terms of the EFX or the PMMS criteria? As it turns out (Prop. 4.5, Cor. 4.9), it does not. As another example, while we know that PMMS implies EFX, it is not clear if an approximate PMMS allocation is also approximately EFX. In fact (Prop. 3.7 and 4.8), it is the other way around! Such results can be conceptually helpful, as they deepen our understanding of the similarities and differences between fairness criteria. Furthermore, this insight allows us to either translate an approximation algorithm for one notion into an approximation algorithm for another, or to establish that such an approach cannot yield approximability or existential results. Given the growing interest in these relaxations, we find it important to further explore these interconnections.

Contribution. We investigate how the four notions mentioned above and their approximate versions are related. For each one of them, we examine the approximation guarantee that it provides in terms of the other three notions. Our results provide an almost complete mapping of the landscape, and in most cases our approximation guarantees are tight. Some of our results provide interesting and surprising consequences. Indicatively, some highlights of our results include:

- PMMS and EFX allocations both provide the same constant approximation with respect to the MMS criterion.
- Although PMMS implies EFX, an approximate PMMS allocation may be arbitrarily bad in terms of approximate EFX. On the contrary, an approximate EFX allocation does provide some guarantee with respect to PMMS.
- While EFX is a much stronger concept than EF1, they both provide comparable constant approximations for the PMMS criterion and this degrades smoothly for approximate EFX and EF1 allocations.
- Although exact PMMS and MMS are defined in a similar manner, the former implies EFX, EF1, and a $4/7$ -approximation in terms of MMS, while the latter provides no guarantee with respect to the other notions.
- Our results also suggest a simple efficient algorithm for computing a $1/2$ -approximate PMMS allocation (the current best),¹ and improvements in certain special cases.

Some of the implications between the different notions are depicted in Figure 1.

Related Work. Regarding the several relaxations of envy-freeness, the notion of EF1 originates in the work of Lipton et al. [2004], where both existence and algorithmic results are provided. The concept of MMS was introduced by Budish [2011], building on concepts of Moulin [1990], and later defined as considered here by Bouveret and Lemaître [2016], who studied a hierarchy of exact fairness concepts. Kurokawa, Procaccia and Wang [2018; 2016] showed that MMS allocations do not always exist even for three agents. From the computational point of view, $2/3$ -approximation algorithms for MMS allocations are known [Kurokawa et al., 2018; Amanatidis et al., 2017b; Barman and Krishnamurthy, 2017] and recently this approximation factor has been improved to $3/4$ [Seddighin et al., 2018]. Some variants of the maximin share criterion have also been considered [Suksompong, 2018]. As for the notions of EFX and PMMS, they were introduced in the work of Caragiannis et al. [2016], which provided some initial results on how these are related to each other as well as to MMS and to EF1. In particular, they established that a PMMS allocation is also an EFX, EF1 and $1/2$ -MMS allocation. Barman and Krishnamurthy [2017] showed that when the agents agree on the ordering of the items according to their value, then EFX allocations do exist and a specific EFX allocation can be produced, that is also a $2/3$ -MMS allocation. The existence and computation of exact and approximate EFX allocations under additive and general valua-

¹Independently, the recent work of [Barman et al., 2018a] (Theorem 1), which concerns a stronger concept, also implies an efficient algorithm for computing $1/2$ -PMMS allocations.




Figure 1: An indicative chart of implications with envy-freeness as a starting point. All implications shown are tight or almost tight.

tions was studied by Plaut and Roughgarden [2018], establishing the currently best $1/2$ -approximation. Finally, in a recent work, Barman et al. [2018a] introduce two new fairness notions that are strongly related to the notions studied here, namely *groupwise maximin share fairness* (GMMS) and *envy-freeness up to one less-preferred good* (EFL). Studying how their approximate versions fit into the landscape explored here is an interesting direction for future work.

2 Preliminaries

We assume we have a set of n agents, $N = \{1, 2, \dots, n\}$ and a set M of m indivisible goods. Following the usual assumptions in the majority of the fair division literature, each agent is associated with a monotone, *additive* valuation function. Hence, for every $S \subseteq M$, $v_i(S) = \sum_{g \in S} v_i(\{g\})$. For simplicity, we will use $v_i(g)$ instead of $v_i(\{g\})$ for $g \in M$. Monotonicity here implies that $v_i(g) \geq 0$ for every $i \in N, g \in M$.

We are interested in solutions that allocate the whole set of goods M to the agents. An allocation of M to the n agents is therefore a partition, $\mathcal{A} = (A_1, \dots, A_n)$, where $A_i \cap A_j = \emptyset$ and $\bigcup_i A_i = M$. By $\Pi_n(M)$ we denote the set of all partitions of a set M into n bundles.

2.1 Fairness Concepts

Our work focuses on relaxations of the classic notion of envy-freeness, initially suggested by Gamow and Stern [1958], and more formally by Foley [1967] and Varian [1974].

Definition 2.1. An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is *envy-free* (EF), if for every $i, j \in N$, $v_i(A_i) \geq v_i(A_j)$.

It is well known that envy-freeness is a very strict requirement in the presence of indivisible goods. This gives rise to considering relaxations of envy-freeness, with the hope of obtaining more positive results.

We begin with two such relaxations, and their approximate versions, where an agent may envy some other agent, but only by an amount dependent on the value of a single item in the other agent’s bundle. Formally:

Definition 2.2. An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is an

- α -EF1 allocation (α -envy-free up to one good), if for every pair of agents $i, j \in N$, with $A_j \neq \emptyset$, there exists an item $g \in A_j$, such that $v_i(A_i) \geq \alpha \cdot v_i(A_j \setminus \{g\})$.
- α -EFX allocation (α -envy-free up to any good), if $v_i(A_i) \geq \alpha \cdot v_i(A_j \setminus \{g\})$ holds for every pair $i, j \in N$, with $A_j \neq \emptyset$, and for every $g \in A_j$ with $v_i(g) > 0$.²

²The requirement that $v_i(g) > 0$ in the definition of α -EFX has

Note that for $\alpha = 1$ in the above definitions, we obtain precisely the notions of envy-freeness up to one good (EF1) [Budish, 2011] and envy-freeness up to any good (EFX) [Caragiannis *et al.*, 2016]. The difference between these two notions is simply the quantifier regarding the item that eliminates envy when removed from an agent’s bundle. Clearly, EF implies EFX, which in turn implies EF1.

We now move on to a different notion, also proposed by Budish [2011]. Motivated by the question of what we can guarantee in the worst case to the agents, the rationale is to think of a generalization of the well-known cut-and-choose protocol to multiple agents: Suppose that agent i is asked to partition the goods into n bundles and then the rest of the agents choose a bundle before i . In the worst case, agent i will be left with her least valuable bundle. Hence, a risk-averse agent would choose a partition that maximizes the minimum value of a bundle in the partition. This gives rise to the following definition.

Definition 2.3. *Given n agents, and a subset $S \subseteq M$ of items, the n -maximin share of agent i with respect to S is:*

$$\mu_i(n, S) = \max_{\mathcal{A} \in \Pi_n(S)} \min_{A_j \in \mathcal{A}} v_i(A_j).$$

From the definition, it directly follows that $n \cdot \mu_i(n, S) \leq v_i(S)$. When $S = M$, this quantity is called the *maximin share* of agent i . We say that $\mathcal{T} \in \Pi_n(M)$ is an n -maximin share defining partition for agent i , if $\min_{T_j \in \mathcal{T}} v_i(T_j) = \mu_i(n, M)$. When it is clear from context what n and M are, we will simply write μ_i instead of $\mu_i(n, M)$. The solution concept we are interested in, asks for a partition that gives each agent her (approximate) maximin share.

Definition 2.4. *An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is called an α -MMS (α -maximin share) allocation if $v_i(A_i) \geq \alpha \cdot \mu_i$, for every $i \in N$.*

The last notion we define is related but not directly comparable to MMS. The idea is that instead of imagining an agent i partitioning the items into n bundles, we think of i as partitioning the combined bundle of herself and another agent into two bundles and receiving the one she values less.

Definition 2.5. *An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is called an α -PMMS (α -pairwise maximin share) allocation if for every pair of agents $i, j \in N$,*

$$v_i(A_i) \geq \alpha \cdot \max_{(B_1, B_2)} \min\{v_i(B_1), v_i(B_2)\},$$

where $(B_1, B_2) \in \Pi_2(A_i \cup A_j)$.

In Definitions 2.4 and 2.5, when $\alpha = 1$, we refer to the corresponding allocations as MMS and PMMS allocations respectively. It has been already observed that the notion of PMMS is stronger than EFX [Caragiannis *et al.*, 2016].

Example 1. To illustrate the relevant definitions, let us consider an instance with 3 agents and 5 items:

	a	b	c	d	e
Agent 1	3	1	1	1	4
Agent 2	4	3	3	1	4
Agent 3	3	2	1	3	4

been dropped by Plaut and Roughgarden [2018] but several of their results hold under the assumption of all values being strictly positive.

If $M = \{a, b, c, d, e\}$ is the set of items, one can see that $\mu_1(3, M) = 3$, $\mu_2(3, M) = 4$, $\mu_3(3, M) = 4$. For example, looking at agent 1, there exists a partition so that the worst bundle is worth a value of 3, and there is no partition where the worst bundle is better. Similarly, agent 2 can produce a partition where the worst bundle has a value of 4 for her, and there is no other partition that can guarantee a better worst-case value.

Let us examine the allocation $\mathcal{A} = (\{e\}, \{b, c\}, \{a, d\})$. Clearly, this is an EF allocation, and hence it is also an EFX, EF1, MMS, and PMMS allocation. Consider now the allocation $\mathcal{B} = (\{a\}, \{b, e\}, \{c, d\})$. This is no longer EF and it is also neither EFX nor PMMS. However, it is an EF1 as well as an MMS allocation. Clearly, each agent i receives a value of at least $\mu_i(3, M)$. To see why it is EF1, note that agents 1 and 3 envy agent 2 but removing item e from the bundle of agent 2 eliminates the envy from either agent. We can also see that \mathcal{B} is a $\frac{3}{4}$ -EFX allocation. To verify this, we can look at agent 1, and compare the value of her bundle to the bundle of agent 2 when we remove either item b or e . Finally, it is not hard to check that \mathcal{B} is also a $\frac{3}{4}$ -PMMS allocation.

3 EFX and EF1 Allocations

We begin our exposition with the more direct relaxations of envy-freeness, EF1 and EFX. Within this section, we always start with either an α -EF1 or α -EFX allocation, for some $\alpha \in (0, 1]$, and explore the implications and fairness guarantees that can be derived. We also pay particular attention to the case of exact EFX or EF1 allocations, i.e., for $\alpha = 1$. Due to lack of space, several of our proofs, are deferred to the full version of the paper.

We already know that EFX is stronger than EF1. Our first warm-up proposition states that this also holds in an approximation preserving sense. We complement this by the fact that EF1 allocations cannot provide any approximation to EFX.

Proposition 3.1. *For $n \geq 2$, any α -EFX allocation is also an α -EF1 allocation for any $\alpha \in (0, 1]$. On the other hand, an EF1 allocation is not necessarily an α -EFX allocation for any $\alpha \in (0, 1]$.*

Before proving guarantees in terms of MMS and PMMS, we state a useful technical lemma which generalizes both (the $k = 1$ case of) Lemma 1 of Bouveret and Lemaître [2016] and Lemma 3.4 of Amanatidis *et al.* [2017b].

Lemma 3.2. *Suppose $\mathcal{T} \in \Pi_n(M)$ is an n -maximin share defining partition for agent i , i.e., $\min_{T_j \in \mathcal{T}} v_i(T_j) = \mu_i(n, M)$. Then, for any set of goods S , such that there exists some j with $S \subseteq T_j$, it holds that*

$$\mu_i(n-1, M \setminus S) \geq \mu_i(n, M).$$

In particular, for any item g , $\mu_i(n-1, M \setminus \{g\}) \geq \mu_i(n, M)$.

Next we study EFX allocations in terms of the MMS guarantees they can provide, starting with the case of a small number of agents.

Proposition 3.3. *For $n \in \{2, 3\}$, an EFX allocation is also a $\frac{2}{3}$ -MMS allocation. Moreover, this guarantee is tight.*

Proof. We prove the statement for $n = 2$. The proof for the case of three agents is of similar flavor, albeit more complex, and is deferred to the full version.

Suppose that we have an EFX allocation $\mathcal{A} = (A_1, A_2)$. We show that for agent 1 we have $v_1(A_1) \geq \frac{2}{3}\mu_1(2, M)$. The case of agent 2 is symmetric. Since items of no value for agent 1 are completely irrelevant for her view of both EFX and MMS allocations, we can assume without loss of generality, that $v_1(g) > 0$ for all $g \in M$.

In the sequel, we write μ_1 instead of $\mu_1(n, M)$. If $A_2 = \emptyset$, then clearly $v_1(A_1) = v_1(M) \geq \mu_1$, so we may assume $|A_2| \geq 1$. If $|A_2| = \{g\}$, then $v_1(A_1) = v_1(M \setminus \{g\}) = \mu_1(1, M \setminus \{g\}) \geq \mu_1(2, M) = \mu_1$, where the inequality follows from Lemma 3.2. The remaining case is when $|A_2| \geq 2$. Suppose, towards a contradiction, that $v_1(A_1) < \frac{2}{3}\mu_1$. Since \mathcal{A} is an EFX allocation, we have that

$$v_1(A_2) \leq v_1(A_1) + v_1(g), \quad (1)$$

where $g \in \arg \min_{h \in A_2} v_1(h)$. Since A_2 contains at least two items, we have

$$v_1(g) \leq \frac{1}{2}v_1(A_2). \quad (2)$$

Combining (1) and (2) we get $v_1(g) \leq v_1(A_1)$. Again by (1), this implies that $v_1(A_2) < \frac{4}{3}\mu_1$. But then, $v_1(M) = v_1(A_1) + v_1(A_2) < \frac{6}{3}\mu_1 = 2\mu_1$, a contradiction, since by definition, we know that $v_i(M) \geq n \cdot \mu_i$ for every $i \in N$.

Regarding tightness, consider the following simple example. Suppose that we have 2 agents and 4 items, a, b, c, d . The agents have identical values over the items, specifically $v_i(a) = v_i(b) = 2$ and $v_i(c) = v_i(d) = 1$, for $i \in \{1, 2\}$. We now look at the allocation $\mathcal{A} = (\{a, b\}, \{c, d\})$. It is easy to see that $\mu_1 = \mu_2 = 3$ and that \mathcal{A} an EFX allocation. However $v_2(\{c, d\}) = 2 = \frac{2}{3}\mu_2$. By adding an arbitrary number of copies of agent 2 and an equal number of items of value 3, this instance can be generalized to any number of agents. \square

Beyond three agents, the picture gets somewhat more complicated. The next result is one of the main highlights of our work. When moving from the case of $n \leq 3$, to a higher number of agents, the approximation guarantee achieved, in terms of MMS, degrades from $2/3$ to a quantity between $4/7 = 0.5714$ and 0.5914 . Surprisingly, the same happens to the MMS guarantee of a PMMS allocation, as we show in the next section.

It is interesting to note that recently Barman and Krishnamurthy [2017] obtained a simple $2/3$ -approximation algorithm for MMS by showing that when agents agree on the ordering of the values of the items, then there exists a particular EFX allocation that is also a $2/3$ -MMS allocation. As indicated by the proof of the following proposition, this is not true for all EFX allocations, even when the agents are identical.

Proposition 3.4. *For $n \geq 4$, any EFX allocation is also a $\frac{4}{7}$ -MMS allocation. On the other hand, an EFX allocation is not necessarily an α -MMS allocation for $\alpha > \frac{8}{13}$ and, as n grows large, for $\alpha > 0.5914$.*

Proof. Let $\mathcal{A} = (A_1, \dots, A_n)$ be an EFX allocation. Suppose, towards a contradiction, that \mathcal{A} is not a $\frac{4}{7}$ -MMS allocation, i.e., there exists some j so that $v_j(A_j) < \frac{4}{7}\mu_j(n, M)$.

Without loss of generality, we assume that agent 1 is such an agent, and write μ_1 instead of $\mu_1(n, M)$. Note that we may remove any agent, other than agent 1, that receives exactly one good, and still end up with an EFX allocation that is not a $\frac{4}{7}$ -MMS allocation. Indeed, if $|A_i| = 1$ for some $i > 1$, then $(A_1, \dots, A_{i-1}, A_{i+1}, \dots, A_n)$ is an EFX allocation of $M \setminus A_i$ to $N \setminus \{i\}$ and, by Lemma 3.2, $\mu_1(n-1, M \setminus A_i) \geq \mu_1(n, M)$. Thus, $v_1(A_1) < \frac{4}{7}\mu_1(n-1, M \setminus A_i)$. Therefore, again without loss of generality, we may assume that $|A_i| > 1$ for all $i > 1$ in the initial allocation \mathcal{A} .

Given \mathcal{A} , we say that a bundle A_j is *bad* if $j > 1$, $|A_j| = 2$, and $\min_{g \in A_j} v_1(g) > \frac{1}{2}v_1(A_1)$. An item is *bad* if it belongs to a bad bundle, while a bundle is *good* if it is not bad. Let B be the set of all bad items.

It is not hard to show that if A_i is good, then $v_1(A_1) \geq \frac{2}{3}v_1(A_i)$. When $i = 1$ this is straightforward; otherwise we consider two cases. First assume that $|A_i| = 2$. Then, by definition, we have $\min_{g \in A_i} v_1(g) \leq \frac{1}{2}v_1(A_1)$ and, since \mathcal{A} is EFX, we have $\max_{g \in A_i} v_1(g) \leq v_1(A_1)$. Thus, $v_1(A_i) \leq \frac{3}{2}v_1(A_1)$. Next, assume that $|A_i| \geq 3$. Then, $\min_{g \in A_i} v_1(g) \leq \frac{1}{3}v_1(A_i)$ and, since \mathcal{A} is EFX, we have $v_1(A_i) - \min_{g \in A_i} v_1(g) \leq v_1(A_1)$. Thus, we get $v_1(A_1) \geq v_1(A_i) - \frac{1}{3}v_1(A_i) = \frac{2}{3}v_1(A_i)$.

Now we are going to show that $v_1(A_1) \geq \frac{4}{7}\mu_1(n', M')$ for a reduced instance that we get by possibly removing some of the items of B , i.e., bad items. We do so in a way that ensures that $\mu_1(n', M') \geq \mu_1$, thus contradicting the choices of \mathcal{A} and A_1 . We consider an n -maximin share defining partition \mathcal{T} for agent 1, i.e., $\min_{T_i \in \mathcal{T}} v_1(T_i) = \mu_1$. If there is a bundle of \mathcal{T} containing two items of B , g_1, g_2 , then we remove those two items and reduce the number of agents by one. By Lemma 3.2, we have that $\mu_1(n-1, M \setminus \{g_1, g_2\}) \geq \mu_1$. We repeat as many times as necessary to get a reduced instance with $n' \leq n$ agents and a set of items $M' \subseteq M$ for which there is an n' -maximin share defining partition \mathcal{T}' for agent 1, such that no bundle contains more than 1 item of B . By repeatedly using Lemma 3.2, we have $\mu_1(n', M') \geq \mu_1$.

Let x be the number of items of B in the reduced instance (recall that these items are defined with respect to \mathcal{A} , hence they belong to a bundle of size 2 in \mathcal{A} and have value at least $\frac{1}{2}v_1(A_1)$ for agent 1). Clearly, x cannot be greater than n' , or some bundle of \mathcal{T}' would contain at least 2 bad items. Further, if $|B| = y$, i.e., the number of bad items in the original instance, then we know that the number of good bundles in \mathcal{A} was $n - \frac{y}{2}$, and that the number of agents was reduced $\frac{y-x}{2}$ times, i.e., $n' = n - \frac{y-x}{2}$. Hence, we may express the number of good bundles in the original instance in terms of n' and x only, as $n' - \frac{x}{2}$.

Recall that $\mu_1 \leq \mu_1(n', M')$ and, by the definition of maximin share, $\mu_1(n', M') \leq \frac{1}{n'}v_1(M')$. Thus

$$\mu_1 \leq \frac{1}{n'}v_1(M'). \quad (3)$$

In order to upper bound $v_1(M')$, notice that M' contains all the items of all the good bundles of \mathcal{A} plus x bad items. As discussed above, the value (with respect to agent 1) of each good bundle is upper bounded by $\frac{3}{2}v_1(A_1)$. On the other hand, \mathcal{A} being EFX implies that $\max_{g \in A_i} v_1(g) \leq v_1(A_1)$

for any bad bundle A_i . That is, any bad item's value is upper bounded by $v_1(A_1)$. So, we have

$$\begin{aligned} v_1(M') &\leq x \cdot v_1(A_1) + \left(n' - \frac{x}{2} - 1\right) \frac{3}{2} v_1(A_1) + v_1(A_1) \\ &= \left(\frac{3n'}{2} + \frac{x}{4} - \frac{1}{2}\right) v_1(A_1) \\ &\leq \left(\frac{3n'}{2} + \frac{n'}{4} - \frac{1}{2}\right) v_1(A_1) \\ &= \frac{7n' - 2}{4} v_1(A_1). \end{aligned} \quad (4)$$

We combine (3) and (4) to get $v_1(A_1) \geq \frac{4n'}{7n'-2} \mu_1 \geq \frac{4}{7} \mu_1$, which contradicts the choices of \mathcal{A} and A_1 .

This establishes the positive part of the theorem. To see that EFX does not imply anything stronger than $\frac{8}{13}$ -MMS, notice that for $n' = n = 4$ the above analysis can be tight. On one hand, $\frac{4n'}{7n'-2} = \frac{8}{13}$ in this case, while on the other hand the following instance indicates that this is the best one can guarantee for 4 agents. Suppose that we have 12 items with the following values for agent 1:

$$v_1(g_i) = \begin{cases} 1/8 & 1 \leq i \leq 4 \\ 1/2 & 5 \leq i \leq 8 \\ 1 & 9 \leq i \leq 12 \end{cases}$$

It is not hard to see that $\mu_1 = 1 + \frac{1}{2} + \frac{1}{8} = \frac{13}{8}$ in this instance. Now consider the allocation $\mathcal{A} = (\{g_1, \dots, g_5\}, \{g_6, g_7, g_8\}, \{g_9, g_{10}\}, \{g_{11}, g_{12}\})$. Assuming that agents 2, 3 and 4 are identical to agent 1, it is not hard to check that this is an EFX allocation. Yet, $v_1(A_1) = 1 = \frac{8}{13} \mu_1$. By adding an arbitrary number of copies of agent 4 and her bundle, this instance can be generalized to any number of agents.

The stronger bound for α as n grows large follows by Proposition 4.1 and Proposition 4.3. \square

By following a similar analysis as in the proof of Proposition 3.3, it can be shown that for any number of agents α -EFX implies at least $\frac{\alpha}{2}$ -MMS. The actual guarantee, moreover, cannot be much better than this.

Proposition 3.5. *For $n \geq 2$ and $\alpha \in (0, 1)$, any α -EFX allocation is also an $\frac{\alpha n}{\alpha + 2n - 2}$ -MMS allocation but not necessarily a β -MMS allocation, for any $\beta > \frac{2\alpha}{2+\alpha}$. For $n \geq 4$, the upper bound is improved to $\max\{\frac{\alpha}{1+\alpha}, \frac{8\alpha}{11+2\alpha}\}$.*

In contrast to EFX, α -EF1 allocations provide a much weaker approximation of maximin shares, namely a ratio of $O(\frac{1}{n})$ for constant α . The following proposition generalizes a result of Caragiannis et al. [2016].³

Proposition 3.6. *For $n \geq 2$, any α -EF1 allocation is also a $\frac{\alpha}{n-1+\alpha}$ -MMS allocation for any $\alpha \in (0, 1]$, and this is tight.*

Finally, we investigate the implications that can be derived for EFX and EF1 allocations, in terms of PMMS guarantees. Despite Proposition 3.1 suggesting that (approximate) EF1 is

much weaker than (approximate) EFX, the two notions give comparable guarantees with respect to PMMS. In particular, the guarantee implied by α -EFX can be at most $4/3$ times the guarantee implied by α -EF1.

Proposition 3.7. *For $n \geq 2$, any α -EFX allocation is also a $\frac{2\alpha}{2+\alpha}$ -PMMS allocation for any $\alpha \in (0, 1]$, and this is tight.*

Proposition 3.8. *For $n \geq 3$, any α -EF1 allocation is also an $\frac{\alpha}{1+\alpha}$ -PMMS allocation for any $\alpha \in (0, 1]$, and this is tight.*

Algorithmic Implications for PMMS. The last two propositions also have further consequences. First of all, it is known that for additive valuations, EF1 allocations can be computed efficiently by a simple round-robin algorithm [Lipton et al., 2004; Caragiannis et al., 2016]. Hence Proposition 3.8 yields a $1/2$ -PMMS allocation in polynomial time. Moreover, exact EFX allocations can be computed efficiently when agents have the same ordering on the values of the goods [Barman and Krishnamurthy, 2017; Plaut and Roughgarden, 2018], implying a $2/3$ -approximation by Proposition 3.7. These are summarized below.

Corollary 3.9. *For $n \geq 3$,*

- (1) *the round-robin algorithm, where each agent picks in her turn her favorite available item, produces a $\frac{1}{2}$ -PMMS allocation.*
- (2) *we can compute in polynomial time a $\frac{2}{3}$ -PMMS allocation when all agents agree on the ordering of the goods with respect to their value.*

It is an interesting open problem to compute a better than $1/2$ -PMMS allocation for additive valuations. So far, the existence of 0.618 -PMMS allocations has been established but without an efficient algorithm [Caragiannis et al., 2016].

4 PMMS and MMS Allocations

We now explore analogous questions with Section 3, but starting now with (approximate) MMS or PMMS allocations. We begin with the guarantees that PMMS allocations imply for MMS and vice versa. In order to proceed, we will make use of the following observation, that PMMS implies EFX.

Proposition 4.1 ([Caragiannis et al., 2016]). *For $n \geq 2$, any PMMS allocation is also an EFX allocation.*

For two agents, it is clear by their definition that the notions of MMS and PMMS are identical. For more agents, by Propositions 3.3, 3.4 and 4.1, we have the following corollary.

Corollary 4.2. *For $n = 3$, a PMMS allocation is also a $\frac{2}{3}$ -MMS allocation. Moreover, for $n \geq 4$, a PMMS allocation is also a $\frac{4}{7}$ -MMS allocation.*

Moreover, the guarantees of Corollary 4.2 are tight for a small number of agents and almost tight for bigger instances.

Proposition 4.3. *For $n \geq 3$, a PMMS allocation is not necessarily an α -MMS allocation for $\alpha > \frac{2}{3}$ and, as n grows large, for $\alpha \geq 0.5914$.*

³The bound by Caragiannis et al. [2016] is for $\alpha = 1$ and for allocations that are both EF1 and Pareto optimal. It follows by their proof, however, that it holds even when Pareto optimality is dropped.

Proof. Suppose that we have the following instance with 3 agents and 6 items. The items have the following values for agent 1:

$$v_1(g_i) = \begin{cases} 1/2 & 1 \leq i \leq 3 \\ 1 & 4 \leq i \leq 6 \end{cases}$$

Clearly, $\mu_1 = \frac{3}{2}$. Consider the allocation $\mathcal{A} = (\{g_4\}, \{g_1, g_2, g_3\}, \{g_5, g_6\})$. Assuming that agents 2 and 3 have large value for their bundles and negligible value for everything else, it is easy to check that \mathcal{A} is a PMMS allocation. However, $v_1(A_1) = 1 = \frac{2}{3}\mu_1$. By adding an arbitrary number of copies of agent 3 and her bundle, this instance can be generalized to any number of agents.

Next, we show a stronger bound as n grows large. The construction below achieves the desired bound for $n \geq 3 \cdot 7 \cdot 43 \cdot 1806 = 1,631,721$. However, there is a smooth transition from $2/3$ to that; e.g., already for $n \geq 21$ we get a bound of $3/5$, which worsens as n increases. We are going to consider the suggested allocation from the viewpoint of agent 1, while assuming that agents 2 through n have large value for their bundles and negligible value for everything else. Since there is a large number of items, we are not going to define $v_1(\cdot)$ explicitly, but implicitly through the different types of bundles the agents get. So consider the allocation \mathcal{A} where:

- agent 1 receives 1 item of value 1,
- $\lfloor \frac{n}{3} \rfloor$ agents receive 3 items of value $\frac{1}{2}$ each,
- $\lfloor \frac{n}{7} \rfloor$ agents receive 7 items of value $\frac{1}{6}$ each,
- $\lfloor \frac{n}{43} \rfloor$ agents receive 43 items of value $\frac{1}{42}$ each,
- $\lfloor \frac{n}{1807} \rfloor$ agents receive 1807 items of value $\frac{1}{1806}$ each,
- the remaining (at least $\frac{n-1}{2}$) agents receive 2 items of value 1 each.

It is easy to see that \mathcal{A} is EFX. What may not be obvious is that the number of agents receiving 2 items is at least $\frac{n-1}{2}$. However, it is a matter of simple calculations to check that $(\lfloor \frac{n}{3} \rfloor + \lfloor \frac{n}{7} \rfloor + \lfloor \frac{n}{43} \rfloor + \lfloor \frac{n}{1807} \rfloor) + 1 \leq \frac{n+1}{2}$ for $n \geq 3$. We show now how to get the bound for $n = 1,631,721$ (in which case the agents receiving 2 items are exactly $\frac{n-1}{2}$). Then, this instance can be generalized to any number of agents by just adding more agents who receive 2 items of value 1 each.

Calculating now μ_1 is straightforward: in an n -maximin share defining partition for agent 1, each agent would receive exactly one item of each type. That is, $\mu_1 = 1 + \frac{1}{2} + \frac{1}{6} + \frac{1}{42} + \frac{1}{1806}$, and $v_1(A_1) = 1 < 0.5914 \cdot \mu_1$. \square

We continue with the worst-case guarantee we can get for MMS by an α -PMMS allocation, with $\alpha < 1$. Notice that now the guarantee degrades with n .

Proposition 4.4. *For $n \geq 3$ and $\alpha \in (0, 1)$, any α -PMMS allocation is also an $\frac{\alpha}{2(n-1)-\alpha(n-2)}$ -MMS allocation but not necessarily a β -MMS allocation, for any $\beta > \frac{\alpha}{n-1-\alpha(n-2)}$.*

The next result exhibits a sharp contrast between PMMS and MMS. Although PMMS allocations (exact or approximate) imply some MMS guarantee, even exact MMS allocations do not imply any approximation with respect to PMMS.

Proposition 4.5. *For $n \geq 3$, an α -MMS allocation is not necessarily a β -PMMS allocation for any $\alpha, \beta \in (0, 1]$.*

We conclude this section by discussing the guarantees of MMS and PMMS with respect to EF1 and EFX. Even for EF1, we see that MMS and PMMS differ significantly in what they can achieve.

Proposition 4.6. *For $n \geq 2$, an α -PMMS allocation is also an $\frac{\alpha}{2-\alpha}$ -EF1 allocation for any $\alpha \in (0, 1)$, and this is tight.*

Proposition 4.7. *For $n \geq 3$, an α -MMS allocation is not necessarily a β -EF1 allocation for any $\alpha, \beta \in (0, 1]$.*

When one focuses on guarantees in terms of EFX, there are only negative results. In fact, it is rather surprising that even though PMMS implies EFX, an α -PMMS allocation with $\alpha < 1$ does not imply any approximation in terms of EFX.

Proposition 4.8. *For $n \geq 2$, an α -PMMS allocation is not necessarily a β -EFX allocation for any $\alpha, \beta \in (0, 1]$.*

Proof. Consider an instance with $n \geq 2$ identical agents and $n + 1$ items. Let $\alpha \in (0, 1)$ and $V \gg \max\{1, \frac{1}{\alpha} - 1\}$. For every agent i we have

$$v_i(g_j) = \begin{cases} V & j = 1 \\ \frac{1}{\alpha} - 1 & j = 2 \\ 1 & 3 \leq j \leq n + 1 \end{cases}$$

We examine the allocation $\mathcal{A} = (A_1, \dots, A_n) = (\{g_1, g_2\}, \{g_3\}, \{g_4\}, \dots, \{g_{n+1}\})$. For any $i \neq 1$, we focus on agents i and 1 from i 's viewpoint. It is easy to see that $\mu_i(2, A_1 \cup A_i) = \frac{1}{\alpha} - 1 + 1 = \frac{1}{\alpha}$. Since $v_i(A_i) = 1$, we get $v_i(A_i) = \alpha \mu_i(2, A_1 \cup A_i)$. As it is straightforward to see that for any other pair of agents that there is no envy, we have that \mathcal{A} is an α -PMMS allocation. On the other hand, every agent $i \in N \setminus \{1\}$ envies agent 1 up to any item by a factor $\beta = \frac{1}{V}$ which can become arbitrarily small for sufficiently large V . \square

By Propositions 4.8, 4.7, and 3.1, we also have the same result for approximate MMS allocations.

Corollary 4.9. *For $n \geq 2$, an α -MMS allocation is not necessarily a β -EFX allocation for any $\alpha, \beta \in (0, 1]$.*

5 Discussion

The main purpose of this work is to provide a deeper understanding of the connections between the exact and the approximate versions of four prominent fairness notions used in discrete models of fair division. In most cases we give tight, or almost tight, results, providing therefore a close to complete picture on the status of these questions. Some of our findings also strike as counter-intuitive, given what was known for the exact versions of these notions.

A direct implication of our results is the non-existence of truthful allocation mechanisms for two agents that achieve any constant approximation of PMMS, EFX, or even EF1. This follows from the corresponding negative result of Amanatidis et al. [2017a] for MMS and Propositions 3.5, 3.6, and 4.4. We defer a more detailed discussion on this to the full version of the paper.

Aside from the questions raised here, there is an abundance of interesting open problems for future research. Deciding the existence of exact EFX or exact PMMS allocations seem to be

two of the most puzzling such problems. Furthermore, obtaining (algorithmically or existentially) allocations with better approximation ratios is another class of equally intriguing problems. So far, we know there exist allocations with ratios of 0.75, 0.618, and 0.5, for MMS, PMMS, and EFX respectively, out of which, the result for PMMS is existential; our Corollary 3.9 only yields a 0.5-approximation in polynomial-time. Finally, the combination of fairness with other desired objectives, like Pareto optimality, creates further algorithmic challenges, even for the seemingly easier notion of EF₁ [Barman *et al.*, 2018b].

Interestingly enough, our results suggest that improving the current state of the art in the approximation of one of these notions does not imply an immediate improvement on the best approximation ratio for the others, with the notable exception of α -EFX. An algorithmic result with $\alpha > 2/3$ or an existential result with $\alpha > 0.895$ for α -EFX would imply an improved algorithmic or existential result for β -PMMS.

References

- [Amanatidis *et al.*, 2017a] G. Amanatidis, G. Birmpas, G. Christodoulou, and E. Markakis. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 545–562, 2017.
- [Amanatidis *et al.*, 2017b] G. Amanatidis, E. Markakis, A. Nikzad, and A. Saberi. Approximation algorithms for computing maximin share allocations. *ACM Trans. Algorithms*, 13(4):52:1–52:28, 2017. *A preliminary conference version appeared in IICALP 2015.*
- [Barman and Krishnamurthy, 2017] S. Barman and S. K. Krishnamurthy. Approximation algorithms for maximin fair division. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 647–664, 2017.
- [Barman *et al.*, 2018a] S. Barman, A. Biswas, S. K. Krishnamurthy, and Y. Narahari. Groupwise maximin fair allocation of indivisible goods. *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, 2018.
- [Barman *et al.*, 2018b] S. Barman, S. K. Krishnamurthy, and R. Vaish. Finding fair and efficient allocations. In *19th ACM Conference on Economics and Computation (EC 2018)*, 2018. To appear.
- [Bouveret and Lemaître, 2016] S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290, 2016. *A preliminary conference version appeared in AAMAS 2014.*
- [Bouveret *et al.*, 2016] S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair allocation of indivisible goods. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12. Cambridge University Press, 2016.
- [Brams and Taylor, 1996] S. J. Brams and A. D. Taylor. *Fair Division: from Cake Cutting to Dispute Resolution*. Cambridge University press, 1996.
- [Budish, 2011] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2016] I. Caragiannis, D. Kurokawa, H. C. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 305–322, 2016.
- [Foley, 1967] D. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.
- [Gamow and Stern, 1958] G. Gamow and M. Stern. *Puzzle-Math*. Viking press, 1958.
- [Kurokawa *et al.*, 2016] D. Kurokawa, A. D. Procaccia, and J. Wang. When can the maximin share guarantee be guaranteed? In *30th AAAI Conference on Artificial Intelligence, AAAI 2016*, pages 523–529, 2016.
- [Kurokawa *et al.*, 2018] D. Kurokawa, A. D. Procaccia, and J. Wang. Fair enough: Guaranteeing approximate maximin shares. *J. ACM*, 65(2):8:1–8:27, 2018.
- [Lipton *et al.*, 2004] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *ACM Conference on Electronic Commerce (EC)*, pages 125–131, 2004.
- [Markakis, 2017] E. Markakis. Approximation algorithms and hardness results for fair division with indivisible goods. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 12. AI Access, 2017.
- [Moulin, 1990] H. Moulin. Uniform externalities: Two axioms for fair allocation. *Journal of Public Economics*, 43(3):305–326, 1990.
- [Moulin, 2003] H. Moulin. *Fair division and collective welfare*. MIT Press, 2003.
- [Plaut and Roughgarden, 2018] B. Plaut and T. Roughgarden. Almost envy-freeness with general valuations. In *29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 2584–2603. SIAM, 2018.
- [Procaccia, 2016] A. D. Procaccia. Cake cutting algorithms. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13. Cambridge University Press, 2016.
- [Robertson and Webb, 1998] J. M. Robertson and W. A. Webb. *Cake Cutting Algorithms: be fair if you can*. AK Peters, 1998.
- [Seddighin *et al.*, 2018] M. Seddighin, M. Ghodsi, M. T. Hajiaghayi, S. Seddighin, and H. Yami. Fair allocation of indivisible goods: Improvement and generalization. In *19th ACM Conference on Economics and Computation (EC 2018)*, 2018. To appear.
- [Suksompong, 2018] W. Suksompong. Approximate maximin shares for groups of agents. *Mathematical Social Sciences*, 2018. To appear.
- [Varian, 1974] H. Varian. Equity, envy and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.