

Multiple Birds with One Stone: Beating $1/2$ for EFX and GMMS via Envy Cycle Elimination

Georgios Amanatidis,¹ Evangelos Markakis,² Apostolos Ntokos²

¹Department of Computer, Control and Management Engineering
Sapienza University of Rome
amanatidis@diag.uniroma1.it

²Department of Informatics
Athens University of Economics and Business
{markakis, apostntok}@gmail.com

Abstract

Several relaxations of envy-freeness, tailored to fair division in settings with indivisible goods, have been introduced within the last decade. Due to the lack of general existence results for most of these concepts, great attention has been paid to establishing approximation guarantees. In this work, we propose a simple algorithm that is *universally fair* in the sense that it returns allocations that have good approximation guarantees with respect to four such fairness notions at once. In particular, this is the first algorithm achieving a $(\phi - 1)$ -approximation of *envy-freeness up to any good* (EFX) and a $\frac{2}{\phi+2}$ -approximation of *groupwise maximin share fairness* (GMMS), where ϕ is the golden ratio. The best known approximation factor, in polynomial time, for either one of these fairness notions prior to this work was $1/2$. Moreover, the returned allocation achieves *envy-freeness up to one good* (EF1) and a $2/3$ -approximation of *pairwise maximin share fairness* (PMMS). While EFX is our primary focus, we also exhibit how to fine-tune our algorithm and improve further the guarantees for GMMS or PMMS.

Finally, we show that GMMS—and thus PMMS and EFX—allocations always exist when the number of goods does not exceed the number of agents by more than two.

1 Introduction

The mathematical study of *fair division* has a long and intriguing history, starting with the formal introduction of the cake-cutting problem by Banach, Knaster and Steinhaus (Steinhaus 1948). Ever since, we have seen the emergence of several fairness criteria, such as the classic notion of *envy-freeness*, that has a dominant role in the literature, see e.g., (Brandt et al. 2016) and references therein. On the other hand, the computational study of finding fair allocations when the resources are *indivisible items* is more recent. It is motivated by the realization that envy-freeness and other classic fairness notions are too demanding for the discrete setting. In particular, even with two agents and one item, it is impossible to produce an allocation with *any* reasonable worst-case approximation guarantee with respect to envy-freeness.

Within the last decade, these considerations have led to natural relaxations of envy-freeness, which are more suitable

for the context of indivisible goods. The most prominent examples, that are also the focus of our work, include the notions of envy-freeness up to one good (EF1) and up to any good (EFX), maximin share fairness (MMS), as well as pairwise and groupwise maximin share fairness (PMMS and GMMS respectively). These relatively new concepts breathed new life into the field of fair division, but they do not come without their issues. Most importantly, although they are generally easier to satisfy than envy-freeness, proving existence results has turned out to be a very challenging task (with the exception of EF1). For instance, it is an open problem to resolve whether EFX or PMMS allocations always exist, even for three agents with additive valuation functions. Surprisingly, existence remains unresolved even when the number of items is just slightly larger than the number of agents.

A reasonable approach is to focus on approximate versions of these relaxations. Indeed, this has led to a series of positive results, obtaining constant factor approximation algorithms for all the aforementioned relaxed criteria (see Related Work). However, improving on the currently known factors seems to be approaching a stagnation point. For example, soon after the introduction of EFX, a $1/2$ -approximation was established (Plaut and Roughgarden 2018), but there has been no progress beyond $1/2$, despite the active interest on this notion.

We should also stress that these notions capture quite different aspects of fairness. A good approximation of any one of EF1, EFX, MMS and PMMS does not necessarily imply particularly strong guarantees for any of the others (Amanatidis, Birmpas, and Markakis 2018). Hence, it becomes compelling to ask for allocations that attain good guarantees with respect to several fairness notions simultaneously. Such results are rather scarce in the literature, e.g., (Barman et al. 2018; Garg and McGlaughlin 2019) or are purely existential (Cargiannis et al. 2019).

Motivated by the lack of such *universally fair* algorithms, we look at the problem of computing allocations that (approximately) satisfy several fairness notions at the same time. Along the way, we aim to improve the state-of-the-art for two of these notions, namely EFX and GMMS. Somewhat unexpectedly, to do so we rely on simple subroutines that have been repeatedly used in fair division before.

Contribution. Our main contribution is an algorithm that is

universally fair, in the sense that it achieves a better than $1/2$ -approximation for all the notions under consideration. The main results can be summarized in the following statement.

Main Theorem. *We can efficiently compute an allocation that is simultaneously*

- i) EFX up to a factor of 0.618,
- ii) GMMS up to a factor of 0.553 (thus, ditto for MMS),
- iii) EF1, and
- iv) PMMS up to a factor of 0.667.

We view parts i) and ii) of breaking the $1/2$ -approximation barrier for EFX and GMMS, as the highlights of this work. These desirable properties are attained by Algorithm 3 (Section 3). We also suggest variations with improved guarantees for one notion at the expense of the others. The factors achieved by Algorithm 3 and its variants, compared against the state of the art for each notion, are shown in Table 1.

	EFX	EF1	GMMS	PMMS
Best known (poly-time)	0.5	1	0.5	0.781
Algorithm 3	0.618	1	0.553	0.667
Variant in Thm. 15	0.6	1	0.571	0.667
Variant in Thm. 17	0.618	0.894	0.553	0.717

Table 1: Summary of our results and state of the art. Known results in the first row are due to Plaut and Roughgarden (2018), Lipton et al. (2004), Barman et al. (2018), and Kurokawa (2017), respectively.

At a technical level, our results are making use of two algorithms that are known to produce only EF1 allocations. The first one is a simple draft algorithm and the second one is the envy-cycle-elimination algorithm of (Lipton et al. 2004). Although these algorithms on their own do not possess any good approximations with respect to EFX or GMMS, our main insight is that by carefully combining parametric versions of these algorithms, we can obtain approximation guarantees for all the fairness criteria of interest here.

In Section 5, we return to the intriguing issue of existence. We show that GMMS—and thus PMMS and EFX—allocations always exist, and can be found efficiently, when the number of goods does not exceed the number of agents by more than two. While this is a simple case, it is still non-trivial to tackle and has remained unresolved. Quite surprisingly, the idea of envy cycle elimination again comes to the rescue, after we carefully alter a small part of the instance.

Related work. Envy-freeness was initially suggested by Gamow and Stern (1958), and more formally by Foley (1967) and Varian (1974). Regarding the relaxations of envy-freeness, EF1 was defined by Budish (2011), but it was also implicit in the work of Lipton et al. (2004). Budish also defined the notion of maximin shares, based on concepts by Moulin (1990). Later on, Caragiannis et al. (2019) introduced the notions of EFX and PMMS, and even more recently, Barman et al. (2018) proposed to study GMMS allocations. Further variants and generalizations of the criteria we present here have also been considered, see e.g., (Suksompong 2018).

EF1 allocations are known to be efficiently computable by the envy-cycle-elimination algorithm of Lipton et al. (2004). For all other notions, the focus has been on approximation algorithms since existence is either not guaranteed or is still an open problem. The most well studied notion is MMS with a series of positive results (Amanatidis et al. 2017b; Kurokawa, Procaccia, and Wang 2018; Barman and Murthy 2017; Garg, McGlaughlin, and Taki 2019), and best known approximation of $3/4$ (Ghodsi et al. 2018; Garg and Taki 2019). Exact and approximate EFX allocations with both additive and general valuations were studied by Plaut and Roughgarden (2018), achieving the currently best $1/2$ -approximation. Recently, a polynomial time algorithm with the same guarantee has been obtained by Chan et al. (2019). The same factor is also the best known for GMMS allocations in polynomial time, by Barman et al. (2018) via a variant of envy cycle elimination. Finally, the currently best approximation of 0.781 for PMMS is due to Kurokawa (2017), using an approach similar to ours. Connections between the approximate versions of these criteria have been investigated by Amanatidis et al. (2018), and we also refer the reader to the full version of our work (Amanatidis, Markakis, and Ntokos 2019) for a comparison of these implications with our results.

Some of these fairness criteria have also been studied in combination with other objectives, such as Pareto optimality (Barman, Krishnamurthy, and Vaish 2018), truthfulness (Amanatidis, Birmpas, and Markakis 2016; Amanatidis et al. 2017a), or maximizing the Nash welfare (Caragiannis et al. 2019; Caragiannis, Gravin, and Huang 2019).

Finally, in parallel and independently of our work, Chaudhury et al. (2020) also improve the $1/2$ factor for GMMS but not in polynomial time. The focus of their work is different, and involves algorithms for EFX allocations by discarding a relatively small number of items. As an implication of their main results, they follow closely the proof of Proposition 3.4 of Amanatidis, Birmpas, and Markakis (2018), albeit from a different starting point, and obtain a pseudo-polynomial time $4/7$ -approximation algorithm. This matches the guarantee of our Theorem 15 for GMMS but does not provide any good approximation for EFX.

2 Preliminaries

Let $N = \{1, 2, \dots, n\}$ be a set of n agents and M be a set of m indivisible items. Unless otherwise stated, we assume that each agent is associated with a monotone, *additive* valuation function, i.e., for $S \subseteq M$, $v_i(S) = \sum_{g \in S} v_i(\{g\})$. For simplicity, we write $v_i(g)$ instead of $v_i(\{g\})$, for $g \in M$. Monotonicity in this setting is equivalent to all items being *goods*, i.e., $v_i(g) \geq 0$ for every $i \in N, g \in M$. For the algorithms presented in this work, we assume that their input contains the valuation function of each involved agent, i.e., $v_i(g)$ is given to the algorithm for every agent i and good g .

We consider the most standard setting in fair division, where we want to allocate all the goods to the agents (*no free disposal*). 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 $\cup_i A_i = M$. By $\Pi_n(M)$ we denote the set of all partitions of a set M into n bundles. Although we allow for multiple

goods to have the exact same value for a specific agent, we assume a deterministic tie-breaking rule for the goods (e.g., break ties lexicographically). This way we may abuse the notation and write $g = \arg \max_{h \in M} v_i(h)$ instead of “let g be the lexicographically first element of $\arg \max_{h \in M} v_i(h)$ ”.

2.1 Fairness Concepts

All the fairness notions we work with are relaxations of the classic notion of envy-freeness.

Definition 1. An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is *envy-free* (EF), if for every $i, j \in N$, $v_i(A_i) \geq v_i(A_j)$.

As envy-freeness is too strong to ask for, when we deal with indivisible goods, several relaxed fairness notions have been introduced so as to obtain meaningful positive results. We start with two additive relaxations, and their approximate versions, where an agent may envy another agent, but only by an amount dependent on the value of a single good in the other agent’s bundle.

Definition 2. An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is an

- a) 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 a good $g \in A_j$, such that $v_i(A_i) \geq \alpha \cdot v_i(A_j \setminus \{g\})$.
- b) α -EFX allocation (α -envy-free up to any good), if for every pair $i, j \in N$, with $A_j \neq \emptyset$ and every good $g \in A_j$, it holds that $v_i(A_i) \geq \alpha \cdot v_i(A_j \setminus \{g\})$.¹

Of course, for $\alpha = 1$ 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. 2019). It is easy to see that EF implies EFX, which in turn implies EF1.

On a different direction, an interesting family of fairness criteria has been developed around the notion of *maximin shares*, also proposed by Budish (2011). The idea behind maximin shares is to capture the worst-case guarantees of generalizing the famous cut-and-choose protocol to multiple agents: Suppose agent i is asked to partition the goods into n bundles, while knowing that the other agents will choose a bundle before her. In the worst case, she will be left with her least valuable bundle. Assuming that agents are risk-averse, agent i would choose a partition that maximizes the minimum value of a bundle. This gives rise to the following definition.

Definition 3. Given n agents, and a subset $S \subseteq M$ of goods, 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 just 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 simply write μ_i instead of $\mu_i(n, M)$.

The most popular fairness notion based on maximin shares, referred to as *maximin share fairness*, asks for a partition that gives each agent her (approximate) maximin share.

¹ The original definition required the condition to hold for all $g \in A_j$ with $v_i(g) > 0$. This is often dropped, assuming that all values are positive (Plaut and Roughgarden 2018; Caragiannis, Gravin, and Huang 2019). For our work *neither* assumption is needed.

Definition 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 $i \in N$.

Variations of maximin share fairness have also been proposed. Here we focus on two notable examples. The first one, *pairwise maximin share fairness*, is related but not directly comparable to MMS and was introduced by Caragiannis et al. (2019). The idea is to demand an MMS-type guarantee but for any *pair* of agents. That is, we can think of an agent i as considering the combined bundle of herself and another agent and requesting to receive at least her maximin share of this bundle if split into two subsets.

Definition 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 $i, j \in N$, $v_i(A_i) \geq \alpha \cdot \mu_i(2, A_i \cup A_j)$.

Taking this one step further, we can demand an allocation to have an MMS-type guarantee for *any subset* of agents. This is referred to as *groupwise maximin share fairness*, introduced by Barman et al. (2018).

Definition 6. An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is called an α -GMMS (α -groupwise maximin share) allocation if for every subset of agents $N' \subseteq N$ and any agent $i \in N'$, $v_i(A_i) \geq \alpha \cdot \mu_i(|N'|, \cup_{j \in N'} A_j)$.

In Definitions 4, 5, and 6, when $\alpha = 1$, we refer to the corresponding allocations as MMS, PMMS, and GMMS allocations respectively. Clearly, the notion of GMMS is stronger than both MMS and PMMS. Further, it has been observed that EF is stronger than GMMS (Barman et al. 2018) and, when all values are positive, PMMS is stronger than EFX (Caragiannis et al. 2019). It should be noted that the approximate versions of these notions are related in non-straightforward ways (Amanatidis, Birmpas, and Markakis 2018).

We note that in the remainder of the paper, all missing proofs are deferred to the full version of our work (Amanatidis, Markakis, and Ntokos 2019).

2.2 Known EF1 Algorithms

Among the fairness notions defined above, EF1 is the only one for which we know that it can always be achieved. Furthermore, two simple algorithms are already known for computing such allocations in polynomial time. We state below a parametric version of these algorithms so that they can run for a limited number of steps or on a strict subset of the goods, as we are going to use them later as subroutines.

In order to define the *envy-cycle-elimination algorithm* (Algorithm 1) of Lipton et al. (2004), we first need to introduce the notion of an *envy graph*. Suppose we have a partial allocation $\mathcal{P} = (P_1, \dots, P_n)$, i.e., an allocation of a strict subset of M . We define the directed envy graph $G_{\mathcal{P}} = (N, E_{\mathcal{P}})$, where $(i, j) \in E_{\mathcal{P}}$ if and only if agent i currently envies agent j , i.e., $v_i(P_i) < v_i(P_j)$. Algorithm 1 builds an allocation one good at a time; in each step, an agent that no one envies receives the next available good. To ensure that such an agent always exists, the algorithm identifies cycles that are created in the envy graph and eliminates them by appropriately reallocating some of the current bundles.

Regarding tie-breaking in line 10 of the algorithm, we assume that agent i is the lexicographically first node of

Algorithm 1: Envy-Cycle-Elimination(N, \mathcal{P}, M')
where N : set of agents, \mathcal{P} : initial partial allocation, M' :
set of unallocated goods

```

1 Construct the envy graph  $G_{\mathcal{P}}$ 
2 for every  $g \in M'$  in lexicographic order do
3   while there is no node of in-degree 0 in  $G_{\mathcal{P}}$  do
4     Find a cycle  $j_1 \rightarrow j_2 \rightarrow \dots \rightarrow j_r \rightarrow j_1$  in  $G_{\mathcal{P}}$ 
5      $B = P_{j_1}$ 
6     for  $k = 1$  to  $r - 1$  do
7        $P_{j_k} = P_{j_{k+1}}$  /* shift the bundles */
8      $P_{j_r} = B$ 
9     Update  $G_{\mathcal{P}}$ 
10  Let  $i \in N$  be a node of in-degree 0
11   $P_i = P_i \cup \{g\}$ 
12  Update  $G_{\mathcal{P}}$ 
13 return  $\mathcal{P}$ 

```

$G_{\mathcal{P}}$ with in-degree 0. Below we summarize the main known properties of Algorithm 1 that we will utilize in our analysis.

Theorem 7 (Follows by (Lipton et al. 2004)). *Let \mathcal{P} be any EF1 partial allocation and $M' = M \setminus \cup_{i=1}^n P_i$. Then,*

- a) *at the end of each iteration of the for loop, the resulting partial allocation is EF1. Hence, the algorithm terminates with an EF1 allocation in polynomial time (actually, even for agents with general monotone valuation functions).*
- b) *Fix an agent i , and let A_i be the bundle assigned to i at the end of some iteration of the for loop. If A'_i is assigned to i at the end of a future iteration, then $v_i(A'_i) \geq v_i(A_i)$.*

The first property of Theorem 7 simply says that the EF1 property is maintained during the execution of the algorithm, given an initial EF1 allocation. The second property states that agents only get happier throughout the course of the algorithm, since they keep getting better and better bundles.

For additive valuation functions there is, in fact, an even simpler greedy algorithm, referred to in the literature as the *round-robin algorithm*, or the *draft algorithm* (Algorithm 2) that also outputs EF1 allocations, see e.g., (Markakis 2017). Given a fixed ordering of the agents, they simply pick their favorite unallocated good one by one, according to that ordering, until there are no goods left.

Algorithm 2: Round-Robin($N, \mathcal{P}, M', \ell, \tau$)
where N : set of agents, \mathcal{P} : partial allocation, M' : set of unallocated goods, ℓ : an ordering of N , τ : number of steps

```

1  $k = 1$ 
2 while  $M' \neq \emptyset$  and  $\tau > 0$  do
3    $g = \arg \max_{h \in M'} v_{\ell[k]}(h)$ 
4    $P_{\ell[k]} = P_{\ell[k]} \cup \{g\}$ 
5    $M' = M' \setminus \{g\}$ 
6    $k = k + 1 \pmod n$ 
7    $\tau = \tau - 1$ 
8 return  $(\mathcal{P}, M')$ 

```

Theorem 8. *Let ℓ be any ordering of N and $\mathcal{P}_0 = (\emptyset, \dots, \emptyset)$. Then Algorithm 2 with input $(N, \mathcal{P}_0, M, \ell, |M|)$ produces an EF1 allocation in polynomial time.*

3 A Simple Universally Fair Algorithm

As mentioned above, Algorithm 3 is built on Algorithms 1 and 2. In particular, it first runs a simple preprocessing step (Algorithm 4) that determines an appropriate ordering ℓ of the set of agents N . Then, it suffices to run only two rounds of the round-robin algorithm, once with respect to ℓ and once with respect to the reverse of ℓ (the second run is also restricted to a subset of the agents), and finally run the envy-cycle-elimination algorithm on the remaining instance. It should be noted here that the preprocessing step is mostly introduced to facilitate the presentation and the analysis of the algorithm. As it can be seen by its description, Algorithm 4 could be combined with the first run of the round-robin algorithm. Indeed, the final assignments for the h_i s in Algorithm 4 are exactly the goods that the agents receive in line 3 of Algorithm 3 (see also Lemma 9 in the next section).

We use ϕ to denote the *golden ratio*. Recall that $\phi = \frac{1+\sqrt{5}}{2} \approx 1.618$ and that $\phi - 1 = \phi^{-1} \approx 0.618$.

Before we move to the analysis of our algorithm, it is useful to build some more intuition on how things work. The

Algorithm 3: Draft-and-Eliminate(N, M)

```

1  $(\ell, n') = \text{Preprocessing}(N, M)$ 
2 Let  $\mathcal{A} = (A_1, \dots, A_n)$  with  $A_i = \emptyset$  for each  $i \in N$ 
3  $(\mathcal{A}, M') = \text{Round-Robin}(N, \mathcal{A}, M, \ell, n)$ 
4  $\ell^R = (\ell[n], \ell[n-1], \dots, \ell[1])$ 
5  $(\mathcal{A}, M') = \text{Round-Robin}(N, \mathcal{A}, M', \ell^R, n - n')$ 
6  $\mathcal{A} = \text{Envy-Cycle-Elimination}(N, \mathcal{A}, M')$ 
7 return  $\mathcal{A}$ 

```

Algorithm 4: Preprocessing(N, M)

```

1  $L = \emptyset$ ;  $A = N$ ;  $k = 1$ 
2 while  $A \neq \emptyset$  do
3   Let  $i$  be the lexicographically first agent of  $A$ 
4    $h_i = \arg \max_{g \in M} v_i(g)$ 
5    $t_i = m - |M| + 1$  /*  $i$ 's timestamp */
6   Let  $R = (N \setminus (A \cup L)) \cup \{i\}$ 
7    $j = \arg \max_{t \in R} v_i(h_t)$ 
8   if  $\phi \cdot v_i(h_i) < v_i(h_j)$  then
9      $h_i = h_j$ 
10     $L = L \cup \{i\}$ 
11     $\ell[k] = i$ 
12     $k = k + 1$ 
13     $A = (A \setminus \{i\}) \cup \{j\}$ 
14  else
15     $A = A \setminus \{i\}$ 
16     $M = M \setminus \{h_i\}$ 
17 for every  $i \in N \setminus L$  in order of increasing timestamp  $t_i$  do
18    $\ell[k] = i$ 
19    $k = k + 1$ 
20 return  $(\ell, |L|)$ 

```

preprocessing part essentially reorders N so that the first few agents (namely, the first $|L|$ agents) are *quite happy* with their pick in the first round of the round-robin subroutine. For the remaining agents, we make sure that they get a second good before we move to the envy-cycle-elimination algorithm. To do so in a “balanced” way, these agents pick goods in reverse order. The resulting partial allocation, where everyone receives one or two goods, turns out to have all the fairness properties we want to achieve at the end, e.g., it is $(\phi - 1)$ -EFX with respect to the currently allocated goods. Crucially, we show that starting from there and then applying the envy-cycle-elimination algorithm maintains these properties.

Coming back to the preprocessing part, the intuition is to simulate a first round of Algorithm 2 and correct any occurrences of extreme envy. In particular, if an agent envies someone that chose before her by a factor greater than ϕ , then she is moved to a position of high priority in the ordering that is created. The agents moved to the first positions during this process (i.e., agents in L) are guaranteed a good of high value in line 3 of Algorithm 3. To counterbalance their advantage, they are not allowed to pick a second good later in line 5.

To see that Algorithm 3 runs in polynomial time, given the properties we have seen for Algorithms 1 and 2, it suffices to check that the preprocessing step can be efficiently implemented. Indeed, the *if* branch of the while loop in Algorithm 4 may be executed at most n times, since agents are irrevocably added to L . Similarly, the *else* branch may be executed at most n times, as each time the set A becomes smaller and its size never increases in the other parts of the algorithm.

4 Fairness Guarantees of Algorithm 3

We begin our analysis with two useful lemmata about Algorithm 4. We stress that within Algorithm 4, every agent i is associated with a distinct good h_i , although nothing is allocated at this step. The first lemma establishes some comparisons regarding the goods associated with the agents. The second lemma states that this association actually coincides with the partial allocation produced in line 3 of Algorithm 3.

Recall that the set L , defined in Algorithm 4, contains the agents that get to pick first in line 3 of Algorithm 3 at the expense of not choosing a second good in line 5. In terms of Algorithm 3, $L = \{\ell[1], \ell[2], \dots, \ell[n']\}$. The partition of N into L and $N \setminus L$ is pivotal for distinguishing the different cases that are relevant in the analysis.

Lemma 9. *Algorithm 4, with input (N, M) , associates each agent i with a single good h_i , so that*

- a) $v_i(h_i) > \phi \cdot v_i(g)$, for any $i \in L$ and $g \in M \setminus \cup_{k=1}^n \{h_k\}$,
- b) $\phi \cdot v_i(h_i) \geq v_i(h_j)$, for any $i, j \in N \setminus L$.

In order for the above lemma to be of any use, we need a connection between the h_i s and the partial allocations that are produced in the first part of Algorithm 3 (lines 3-5). At a first glance, the issue is that the order in which the goods are assigned in Preprocessing(N, M) is somewhat different than the order in which the goods are allocated in Round-Robin($N, \mathcal{A}, M, \ell, n$). Next we establish this connection.

Lemma 10. *The partial allocation produced in line 3 of Algorithm 3 is $\mathcal{A} = (\{h_1\}, \{h_2\}, \dots, \{h_n\})$, where the h_i s are as in Lemma 9.*

Given the above, we are going to consistently use the h_i notation for the goods allocated in line 3 of Algorithm 3, throughout the remaining of this section. Further, for the agents who receive a second good in line 5 of Algorithm 3 we use h'_i to denote that second good of agent i .

As a warm-up we first obtain that Algorithm 3 maintains the fairness guarantee of its components, i.e., EF1 fairness.

Proposition 11. *Algorithm 3 returns an EF1 allocation.*

4.1 Envy-Freeness up to Any Good

Proving whether EFX allocations always exist or not seems very challenging (Plaut and Roughgarden 2018; Caragianis, Gravin, and Huang 2019). Even improving on the $1/2$ approximation factor of Plaut and Roughgarden (2018) has been one of the most intriguing recent open problems in fair division. In this sense, we view the following as one of the highlights of this work.

Theorem 12. *The allocation $\mathcal{A} = (A_1, \dots, A_n)$ returned by Algorithm 3 is a $(\phi - 1)$ -EFX allocation.*

Proof. Consider the allocation $\mathcal{A} = (A_1, \dots, A_n)$ returned by the algorithm, and fix two distinct agents $i, j \in N$. If $|A_j| = 1$, then clearly, $v_i(A_i) \geq \max_{g \in A_j} v_i(A_j \setminus \{g\}) = 0$. So, assume that $|A_j| \geq 2$ and let h be the last good added to A_j (either in line 5 by reverse round-robin or in line 6 by envy-cycle-elimination). Of course, at the time this happened, A_j may belong to an agent j' other than j . Finally, let $A_i^{\text{old}}, A_{j'}^{\text{old}}$ be the bundles of i and j' , respectively, right before h was allocated (i.e., h was added to $A_{j'}^{\text{old}}$). Note that A_i^{old} may not necessarily be a subset of A_i due to the possible swaps imposed by Algorithm 1, but part b) of Theorem 7 implies that $v_i(A_i) \geq v_i(A_i^{\text{old}})$. We consider four cases, depending on whether $i \in L$ and on the type of step during which h was added to $A_{j'}^{\text{old}}$.

Case 1 ($i \in L$ and h added in line 5). We have $A_i^{\text{old}} = \{h_i\}$, as well as $j' \in N \setminus L$ and $A_j = \{h_{j'}, h'_{j'}\}$. This immediately implies that $v_i(A_i^{\text{old}}) \geq \max\{v_i(h_{j'}), v_i(h'_{j'})\}$ and, thus, $v_i(A_i) \geq \max_{g \in A_j} v_i(A_j \setminus \{g\})$.²

Case 2 ($i \in L$ and h added in line 6). By the way that envy-cycle-elimination chooses who to give the next good to, (line 10 of Algorithm 1), we know that right before h was added, no one envied j' . In particular, $v_i(A_i^{\text{old}}) \geq v_i(A_{j'}^{\text{old}})$. We further have $v_i(A_i^{\text{old}}) \geq v_i(h_i) > \phi \cdot v_i(h)$, where the last inequality directly follows from part a) of Lemma 9. Putting everything together,

$$v_i(A_j) = v_i(A_{j'}^{\text{old}}) + v_i(h) \leq (1 + \phi^{-1})v_i(A_i^{\text{old}}) \leq \phi \cdot v_i(A_i),$$

or, equivalently, $v_i(A_i) \geq \phi^{-1} \cdot v_i(A_j) = (\phi - 1)v_i(A_j)$.

Case 3 ($i \notin L$ and h added in line 5). We have $i, j' \in N \setminus L$ and $A_j = \{h_{j'}, h'_{j'}\}$. If $\ell[i] < \ell[j']$, then we proceed in a way similar to Case 1. Indeed,

$$v_i(A_i) \geq v_i(A_i^{\text{old}}) \geq v_i(h_i) \geq \max\{v_i(h_{j'}), v_i(h'_{j'})\}$$

² Here we achieve the EFX objective exactly. Instead, in a similar argument as in Case 2, we could have used that $v_i(h_i) > v_i(h_{j'})$ and $v_i(h_i) > \phi \cdot v_i(h'_{j'})$ to get $v_i(A_i) \geq (\phi - 1)v_i(A_j)$.

$$= \max_{g \in A_j} v_i(A_j \setminus \{g\}).$$

So, assume that $\ell[i] > \ell[j']$. This, in particular, means that $v_i(h'_i) \geq v_i(h'_{j'})$. We have

$$\begin{aligned} v_i(A_i) &\geq v_i(A_i^{\text{old}}) \geq v_i(h_i) + v_i(h'_i) \geq \frac{1}{\phi} v_i(h_{j'}) + v_i(h'_{j'}) \\ &\geq \frac{1}{\phi} (v_i(h_{j'}) + v_i(h'_{j'})) = (\phi - 1)v_i(A_j), \end{aligned}$$

where the third inequality directly follows from part b) of Lemma 9.

Case 4 ($i \notin L$ and h added in line 6). Arguing like in Case 2, we have $v_i(A_i^{\text{old}}) \geq v_i(A_{j'})^{\text{old}}$. Moreover, by the way round-robin works, we know that $v_i(h_i) \geq v_i(h'_i) \geq v_i(h)$. In particular, $v_i(h) \leq \frac{1}{2}v_i(\{h_i, h'_i\}) \leq \frac{1}{2}v_i(A_i^{\text{old}})$. Putting things together, we have

$$v_i(A_j) = v_i(A_{j'}^{\text{old}}) + v_i(h) \leq \left(1 + \frac{1}{2}\right)v_i(A_i^{\text{old}}) \leq \phi \cdot v_i(A_i).$$

Equivalently, $v_i(A_i) \geq (\phi - 1)v_i(A_j)$. \square

It is not hard to see that our analysis is tight, i.e., there are instances (even with $n = 2$ and $m = 4$) for which the resulting allocation is not $(\phi - 1 + \varepsilon)$ -EFX for any $\varepsilon > 0$.

4.2 Groupwise Maximin Share Fairness

A result of Amanatidis et al. (2018) (Proposition 3.4) implies that every *exact* EFX allocation is also a $4/7$ -GMMS allocation. Of course, the allocation produced by Algorithm 3 is not exact EFX and, in general, an arbitrary $(\phi - 1)$ -EFX allocation need not even be a 0.404 -GMMS allocation (see full version). For the particular allocation returned by Algorithm 3, however, we can show that the GMMS guarantee is significantly better. Parts of our proof closely follow the proof of the aforementioned proposition of Amanatidis et al. (2018).

We are going to need the following simple lemma that allows to remove appropriately chosen subsets of goods, while reducing the number of agents, so that the maximin share of a specific agent does not decrease. In particular, the lemma implies that for any good g , $\mu_i(n - 1, M \setminus \{g\}) \geq \mu_i(n, M)$.

Lemma 13 (Amanatidis et al (2018)). *Suppose $\mathcal{T} \in \Pi_n(M)$ is an n -maximin share defining partition for agent i . Then, for any set of goods S , such that there exists some j with $S \subseteq T_j$, it holds that $\mu_i(n - 1, M \setminus S) \geq \mu_i(n, M)$.*

Theorem 14. *The allocation $\mathcal{A} = (A_1, \dots, A_n)$ returned by Algorithm 3 is a $\frac{2}{\phi+2}$ -GMMS allocation.*

Proof. Suppose that \mathcal{A} is not a $\frac{2}{\phi+2}$ -GMMS allocation, i.e., there exists a subset of agents $Q \subseteq N$ with $|Q| = q$, and some agent $j \in Q$, so that $v_j(A_j) < \frac{2}{\phi+2}\mu_j(q, R)$, where $R = \cup_{k \in Q} A_k$. That is, with respect to Q and R , the restriction of \mathcal{A} to Q is not a $\frac{2}{\phi+2}$ -MMS allocation. To facilitate the presentation, and without loss of generality, we may assume that $Q = [q]$ and that agent 1 is such a ‘‘dissatisfied’’ agent. We write μ_1 instead of $\mu_1(q, R)$.

We may remove any agent in Q , other than agent 1, that receives exactly one good, and still end up with a suballocation that is not a $\frac{2}{\phi+2}$ -GMMS allocation. Indeed, if $|A_i| = 1$

for some $i \in Q \setminus \{1\}$, then $(A_1, \dots, A_{i-1}, A_{i+1}, \dots, A_q)$ is an allocation of $R \setminus A_i$ to $Q \setminus \{i\}$ and, by Lemma 13, $\mu_1' = \mu_1(q - 1, R \setminus A_i) \geq \mu_1$. Thus, $v_1(A_1) < \frac{2}{\phi+2}\mu_1'$. Therefore, again without loss of generality, we may assume that $|A_i| \geq 2$ for all $i \in Q \setminus \{1\}$ in the initial allocation \mathcal{A} . At this point, we make the distinction on whether $1 \in L$ or not. **Case 1 ($1 \in L$).** As we see from (the footnote of) Case 1 and from Case 2 of the proof of Theorem 12, we always have $v_1(A_1) \geq (\phi - 1)v_1(A_i)$ (or equivalently $v_1(A_i) \leq \phi v_1(A_1)$) for all $i \in Q \setminus \{1\}$. Recall that, by the definition of maximin share, $\mu_1 \leq \frac{1}{q}v_1(R)$. Thus

$$q\mu_1 \leq v_1(R) = \sum_{k \in Q} v_1(A_k) \leq q\phi v_1(A_1).$$

That is, we get $v_1(A_1) \geq (\phi - 1)\mu_1 \geq \frac{2}{\phi+2}\mu_1$, which contradicts the choices of \mathcal{A} and A_1 .

Case 2 ($1 \notin L$). Consider some $i \in Q \setminus \{1\}$ and let h be the last good added to A_i . Following the notation introduced in the proof of Theorem 12, this bundle belonged to some agent i' and $A_{i'}^{\text{old}}$ denotes the bundle allocated to i' right before h was added. According to Cases 3 and 4 in the proof of Theorem 12, if h was added in line 5 of Algorithm 3 and $\ell[1] \geq \ell[i']$ or if it was added in line 6, then $v_1(A_i) \leq \phi v_1(A_1)$. We still need to deal with the subcase where h was added in line 5 but $\ell[1] < \ell[i']$. We call such an A_i *dubious*. For dubious bundles, by their definition, we directly have $|A_i| = 2$ and $v_1(A_1) \geq v_1(h_1) \geq \max_{g \in A_i} v_1(g)$. If a bundle A_i is not dubious, or if it is dubious but we have $v_1(A_i) \leq \frac{3}{2}v_1(A_1) < \phi v_1(A_1)$, we say that A_i is *convenient*. A (dubious) bundle is *inconvenient* if it is not convenient. A good is inconvenient if it belongs to an inconvenient bundle. Let B be the set of all inconvenient goods.

Now we are going to show that $v_1(A_1) \geq \frac{2}{\phi+2}\mu_1(q', R')$ for a reduced instance that we get by possibly removing some inconvenient goods. We do so in a way that ensures that $\mu_1(q', R') \geq \mu_1$, thus contradicting the choices of \mathcal{A} and A_1 . We consider a q -maximin share defining partition \mathcal{T} for agent 1 with respect to R , i.e., $\min_{T_i \in \mathcal{T}} v_1(T_i) = \mu_1$ and $\cup_{k \in Q} T_k = R$.³ If there is a bundle of \mathcal{T} containing two goods of B , g_1, g_2 , then we remove those two goods and reduce the number of agents by one. By Lemma 13, we have that $\mu_1(q - 1, R \setminus \{g_1, g_2\}) \geq \mu_1$. We repeat as many times as necessary to get a reduced instance with $q' \leq q$ agents and a set of goods $R' \subseteq R$ for which there is a q' -maximin share defining partition \mathcal{T}' for agent 1, such that no bundle contains more than *one* good from B . By repeatedly using Lemma 13, we get $\mu_1(q', R') \geq \mu_1$.

Let x be the number of goods from B in the reduced instance. Clearly, x cannot be greater than q' , or some bundle of \mathcal{T}' would contain at least 2 inconvenient goods. Further, if $|B| = y$, i.e., the number of inconvenient goods in the original instance, then we know that the number of convenient bundles in the restriction of \mathcal{A} on Q was $q - \frac{y}{2}$, and that the number of agents was reduced $\frac{y-x}{2}$ times, i.e., $q' = q - \frac{y-x}{2}$.

³ Note that while goods in B may have no apparent significance for the reduced instances and the allocations we talk about from this point on, we keep referring to them as *inconvenient*.

That is, we can express the number of convenient bundles in the original instance in terms of q' and x only, as $q' - \frac{x}{2}$.

In order to upper bound $v_1(R')$, notice that R' contains all the goods of all the convenient bundles plus x inconvenient goods. Recall that any good of a dubious bundle has value at most $v_1(A_1)$, and that if A_i is convenient then $v_1(A_i) \leq \phi v_1(A_1)$. So, we have

$$\begin{aligned} v_1(R') &\leq x v_1(A_1) + \left(q' - \frac{x}{2} - 1\right) \phi v_1(A_1) + v_1(A_1) \\ &= (\phi q' + (1 - \phi/2)x - (\phi - 1)) v_1(A_1) \\ &\leq (\phi + (1 - \phi/2)) q' v_1(A_1) = \frac{\phi + 2}{2} q' v_1(A_1). \end{aligned}$$

Combining this inequality with $\mu_1 \leq \mu_1(q', R')$ (by the construction of the reduced instance) and $\mu_1(q', R') \leq \frac{1}{q'} v_1(R')$ (by the definition of maximin share), we get $v_1(A_1) \geq \frac{2}{\phi + 2} \mu_1$, which contradicts the choices of \mathcal{A} and A_1 . \square

A natural question is why we do not achieve the factor of $4/7$ of the original proposition of Amanatidis et al. (2018) instead. A close inspection of the original proof reveals that we need a slightly stronger upper bound for the value of the convenient bundles, i.e., a factor of $3/2$ rather than ϕ that we have here. There is no easy way to fix this *in general* without other things breaking down badly in the analysis of Algorithm 3. The crucial observation, however, is that we only need the distinction of convenient and inconvenient bundles for agents in $N \setminus L$. By fine-tuning line 8 of Algorithm 4, we are able to improve the inequalities about the convenient bundles just for agents in $N \setminus L$ and obtain a $4/7$ -GMMS allocation, at the expense of some loss with respect to EFX.

Theorem 15. *Suppose we modified Algorithm 3 by changing ϕ in line 8 of Algorithm 4 to $3/2$. Then the resulting allocation is a $4/7$ -GMMS allocation. It is also a $3/5$ -EFX, a $2/3$ -PMMS, and an EF1 allocation.*

4.3 Pairwise Maximin Share Fairness

Any result for GMMS directly translates to a result for PMMS with the exact same guarantee. Note, however, that the proof of Theorem 14 suggests that the bad event with respect to GMMS is having many inconvenient bundles. When we only deal with two agents at a time, it is not hard to see that inconvenient bundles are not an issue. In fact, their existence would not be able to force the approximation ratio for PMMS below $2/3$. Indeed, following the cases in the proof of Theorem 12 it is easy to show that this is exactly the guarantee achieved.

Theorem 16. *The allocation $\mathcal{A} = (A_1, \dots, A_n)$ returned by Algorithm 3 is a $2/3$ -PMMS allocation.*

While the above factor is tight, it is possible to improve the PMMS guarantee by modifying slightly the envy graph. The high level idea—due to Kurokawa (2017)—is that an agent from $N \setminus L$ should only exchange her *initial* bundle of two goods for something significantly better.

Let $\mathcal{P} = (P_1, \dots, P_n)$ be the partial allocation produced in line 5 of Algorithm 3. For $\alpha > 1$, the α -modified envy graph $G_{\mathcal{P}}^{\alpha}$ is defined like the envy graph $G_{\mathcal{P}}$ but we drop any edge (i, j) where: $i \in N \setminus L$, and i still has her original bundle, and

$\alpha \cdot v_i(P_i) > v_i(P_j)$. That is, agents in $N \setminus L$ are represented in the envy graph as having an artificially amplified value (by a factor of α) specifically for their original bundles.

The following theorem indicates how far we can push the factor for PMMS, at the expense of EF1, while preserving the original guarantees with respect to EFX and GMMS.

Theorem 17. *Suppose we modified Algorithm 3 by using the $(\phi - \frac{1}{2})$ -adjusted envy graph in Algorithm 1. Then the resulting allocation is a $\frac{4\phi-2}{2\phi+3}$ -PMMS and a $\frac{2}{2\phi-1}$ -EF1 allocation. The guarantees of Theorems 12 and 14 are not affected.*

5 GMMS, PMMS, and EFX with a Few Goods

In this section we focus on the exact versions of the fairness notions under consideration. In particular, we show that GMMS allocations always exist when $m \leq n + 2$. This implies that PMMS and EFX allocations also exist for this case by the discussion in Section 2.⁴

The interesting case is when $m = n + 2$ and is tackled by Algorithm 5. When $m \leq n$ the problem is trivial, and the $m = n + 1$ case is straightforward as well. Adding one extra good, however, makes things significantly more complex. To point out how challenging these simple restricted cases can be, we note that for the much better studied notion of MMS fairness it is still open whether exact MMS allocations exist when $m = n + 5$ (Kurokawa, Procaccia, and Wang 2018).

Quite surprisingly, the envy-cycle-elimination algorithm again comes to the rescue. We first run the round-robin algorithm to allocate $n - 1$ goods to the first $n - 1$ agents. After this, we have 3 goods remaining. Allocating these goods to the last agent may destroy the properties we are after, so we need to be careful on how to handle these three goods. Instead, we (pretend to) pack them into two boxes; the big box (virtual good p) “contains” two goods and the small box (virtual good q) “contains” one. We tell each agent separately that the big box contains her favorite two out of the three items and give the big box to the last agent. Then we proceed using the envy-cycle-elimination algorithm. At the end, the owner of the big box gets her two favorite goods, while the owner of the small box gets the remaining good.

Theorem 18. *For instances with $m \leq n + 2$, a GMMS allocation always exists and can be efficiently computed.*

Algorithm 5: Draft-Pack-and-Eliminate(N, M)

- 1 Let $\ell = (1, 2, \dots, n)$ and $\mathcal{A} = (\emptyset, \dots, \emptyset)$
 - 2 $(\mathcal{A}, M') = \text{Round-Robin}(N, \mathcal{A}, M, \ell, n - 1)$
 - 3 Create two virtual goods p and q , such that for all $i \in N$:
 $v_i(q) = \min_{g \in M'} v_i(g)$ and $v_i(p) = v_i(M') - v_i(q)$
 - 4 Allocate p to agent n
 - 5 $\mathcal{A} = \text{Envy-Cycle-Elimination}(N, \mathcal{A}, \{q\})$
 - 6 Give the owner of p her two favorite goods from M' and the owner of q the remaining good
 - 7 **return** \mathcal{A}
-

⁴ Actually, the existence of EFX allocations is directly implied by the existence of PMMS allocations *only* when all values are positive. However, our result is more general.

Corollary 19. When $m \leq n + 2$, we can efficiently find PMMS and EFX allocations.

Acknowledgments

The first author was supported by NWO Gravitation project NETWORKS (No. 024.002.003), ERC Advanced Grant AM-DROMA (No. 788893), and MIUR PRIN project ALGADIMAR. This work has been partly supported by the COST Action CA16228 “European Network for Game Theory”.

References

- Amanatidis, G.; Birmpas, G.; Christodoulou, G.; and Markakis, E. 2017a. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, 545–562. ACM.
- Amanatidis, G.; Markakis, E.; Nikzad, A.; and Saberi, A. 2017b. Approximation algorithms for computing maximin share allocations. *ACM Trans. Algorithms* 13(4):52:1–52:28.
- Amanatidis, G.; Birmpas, G.; and Markakis, E. 2016. On truthful mechanisms for maximin share allocations. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, 31–37. IJCAI/AAAI Press.
- Amanatidis, G.; Birmpas, G.; and Markakis, V. 2018. Comparing approximate relaxations of envy-freeness. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden.*, 42–48. ijcai.org.
- Amanatidis, G.; Markakis, E.; and Ntokos, A. 2019. Multiple birds with one stone: Beating 1/2 for EFX and GMMS via envy cycle elimination. *CoRR* abs/1909.07650.
- Barman, S., and Murthy, S. K. K. 2017. Approximation algorithms for maximin fair division. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, 647–664. ACM.
- Barman, S.; Biswas, A.; Murthy, S. K. K.; and Narahari, Y. 2018. Groupwise maximin fair allocation of indivisible goods. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, LA, USA, February 2-7, 2018*, 917–924. AAAI Press.
- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, 557–574. ACM.
- Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6):1061–1103.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Trans. Economics and Comput.* 7(3):12:1–12:32.
- Caragiannis, I.; Gravin, N.; and Huang, X. 2019. Envy-freeness up to any item with high nash welfare: The virtue of donating items. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019.*, 527–545. ACM.
- Chan, H.; Chen, J.; Li, B.; and Wu, X. 2019. Maximin-aware allocations of indivisible goods. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, 137–143. ijcai.org.
- Chaudhury, B. R.; Kavitha, T.; Mehlhorn, K.; and Sgouritsa, A. 2020. A little charity guarantees almost envy-freeness. In *Proceedings of the Forty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*. SIAM. To appear.
- Foley, D. K. 1967. Resource allocation and the public sector. *Yale Economics Essays* 7:45–98.
- Gamow, G., and Stern, M. 1958. *Puzzle-Math*. Viking press.
- Garg, J., and McGlaughlin, P. 2019. Improving nash social welfare approximations. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, 294–300. ijcai.org.
- Garg, J., and Taki, S. 2019. An improved approximation algorithm for maximin shares. *CoRR* abs/1903.00029.
- Garg, J.; McGlaughlin, P.; and Taki, S. 2019. Approximating maximin share allocations. In *Fineman, J. T., and Mitzenmacher, M., eds., 2nd Symposium on Simplicity in Algorithms, SOSA@SODA 2019, January 8-9, 2019 - San Diego, CA, USA*, volume 69 of *OASICS*, 20:1–20:11. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.
- Ghods, M.; Hajiaghayi, M. T.; Seddighin, M.; Seddighin, S.; and Yami, H. 2018. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, 539–556. ACM.
- Kurokawa, D.; Procaccia, A. D.; and Wang, J. 2018. Fair enough: Guaranteeing approximate maximin shares. *J. ACM* 65(2):8:1–8:27.
- Kurokawa, D. 2017. *Fair Division in Game Theoretic Settings*. Ph.D. Dissertation, Carnegie Mellon University.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings 5th ACM Conference on Electronic Commerce (EC-2004), New York, NY, USA, May 17-20, 2004*, 125–131. ACM.
- Markakis, E. 2017. Approximation algorithms and hardness results for fair division with indivisible goods. In *Endriss, U., ed., Trends in Computational Social Choice*. AI Access. chapter 12.
- Moulin, H. 1990. Uniform externalities: Two axioms for fair allocation. *Journal of Public Economics* 43(3):305–326.
- Plaut, B., and Roughgarden, T. 2018. Almost envy-freeness with general valuations. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, 2584–2603. SIAM.
- Steinhaus, H. 1948. The problem of fair division. *Econometrica* 16:101–104.
- Suksompong, W. 2018. Approximate maximin shares for groups of agents. *Mathematical Social Sciences* 92:40–47.
- Varian, H. R. 1974. Equity, envy and efficiency. *Journal of Economic Theory* 9:63–91.