




Contents lists available at ScienceDirect

Artificial Intelligence

journal homepage: www.elsevier.com/locate/artint

Fair division of indivisible goods: Recent progress and open questions [☆]

Georgios Amanatidis ^a, Haris Aziz ^b, Georgios Birmpas ^c, Aris Filos-Ratsikas ^d,
Bo Li ^e, Hervé Moulin ^f, Alexandros A. Voudouris ^{a,*}, Xiaowei Wu ^g

^a University of Essex, United Kingdom of Great Britain and Northern Ireland

^b UNSW Sydney, Australia

^c Sapienza University of Rome, Italy

^d University of Edinburgh, United Kingdom of Great Britain and Northern Ireland

^e The Hong Kong Polytechnic University, Hong Kong

^f University of Glasgow, United Kingdom of Great Britain and Northern Ireland

^g University of Macau, Macau

ARTICLE INFO

Article history:

Received 15 August 2022

Received in revised form 18 April 2023

Accepted 12 June 2023

Available online 16 June 2023

Keywords:

Discrete fair division

Envy-freeness

Proportionality

EF1

EFX

MMS

ABSTRACT

Allocating resources to individuals in a *fair* manner has been a topic of interest since ancient times, with most of the early mathematical work on the problem focusing on resources that are infinitely divisible. Over the last decade, there has been a surge of papers studying computational questions regarding the *indivisible* case, for which exact fairness notions such as envy-freeness and proportionality are hard to satisfy. One main theme in the recent research agenda is to investigate the extent to which their relaxations, like maximin share fairness (MMS) and envy-freeness up to any good (EFX), can be achieved. In this survey, we present a comprehensive review of the recent progress made in the related literature by highlighting different ways to relax fairness notions, common algorithm design techniques, and the most interesting questions for future research.

© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Fair division is concerned with the fundamental task of *fairly* partitioning or allocating a set of resources to people with diverse and heterogeneous preferences over these resources. Some examples of fair division in the real world include Course Match, which is employed for course allocation at the Wharton School in the University of Pennsylvania, and the websites Spliddit (spliddit.org) and Fair-Outcomes (fairoutcomes.com), that provide implementations of fair allocation algorithms for various allocation problems, such as sharing rent among roommates, splitting taxi fares, and for assigning goods to a set of individuals (which is the focus of this survey). The associated theory originated in the works of Steinhaus [149], Banach, and Knaster (see [89]), and has been in the focus of economics, mathematics and computer science for the better part of the last century [138]. Most of the classic work on the problem has been devoted to the fair division of *infinitely divisible* resources, where “fair” here may have different interpretations, with two predominant ones being *proportionality* [149] and *envy-freeness* [100,154]; see also [143,132]. In a recent breakthrough, Aziz and Mackenzie [19] showed that an envy-free

[☆] Preliminary versions of this survey appeared as [12] and [28].

* Corresponding author.

E-mail address: alexandros.voudouris@essex.ac.uk (A.A. Voudouris).

allocation of divisible resources (which is also proportional in the standard additive case) can always be found in a finite number of steps.

Compared to the divisible setting, the fair division of *indivisible* resources, referred to as *discrete fair division*, turns out to be inherently more challenging. Indeed, no reasonable fair solution can be guaranteed in some cases; for example, when there is a single valuable item, no matter who gets it, the allocation is not fair to the others. A typical remedy to this situation is to employ randomization, and aim for fairness (such as envy-freeness) in expectation [121,55]. A fundamentally different approach to discrete fair division came via the introduction of appropriate *relaxations* of envy-freeness and proportionality, originating in the works of Lipton et al. [133], Budish [66], Caragiannis et al. [71], and Gourvès et al. [113], which are geared to escape such adverse examples. The main notions that were introduced in this literature are *envy-freeness up to one good* (EF1), *envy-freeness up to any good* (EFX) and *maximin share fairness* (MMS). Since then, work on the topic has flourished, centered around fundamental questions about the existence and the efficient computation of allocations satisfying these or other related fairness criteria.

We remark that the problem of fairly dividing indivisible goods has a very long history, and our aim in this survey is not to provide an extensive review of all of this literature. Instead, our focus is on the recent developments over the past decade, with an emphasis on the algorithmic aspects of the associated problems. For other results, there are several surveys highlighting different perspectives of the theory of fair resource allocation. Moulin [139] provides a review from an economics perspective, Aleksandrov and Walsh [3] and Suksompong [151] focus on online and constrained settings, respectively, whereas Walsh [155] and Aziz [17] provide short surveys targeted more towards a broader audience. Lang and Rothe [128], Bouveret et al. [63], and Markakis [137] consider general discrete fair division settings, and thus are more closely related to this work. Our survey differs from these in that we pay particular attention on algorithmic progress made in the last decade, focusing on scenarios where it is assumed that the agents are equipped with valuation functions (most often additive ones) on the sets of items. Over this period of time, discrete fair division has been at the epicenter of computational fair division, for several different fairness notions and a variety of different settings. In this survey, we highlight the main contributions of this literature, the most significant variants of the main setting, common algorithmic techniques, as well as some of the major open problems in the area. Although they lack most of this content, the surveys of Lang and Rothe [128], and Bouveret et al. [63] remain excellent starting points on questions of preference elicitation and compact representation.

Roadmap. The rest of the survey is organized as follows. We first introduce the problem of discrete fair division and the main fairness notions in Section 2. We then discuss results on EF1 in Section 3, EFX in Section 4, and MMS in Section 5. In Section 6 we introduce other notable fairness notions for the main setting. In Section 7, we talk about two more properties, efficiency and truthfulness, that are usually desirable together with fairness. Finally, Section 8 is dedicated to other settings, including limited information, general valuations, randomness, and more.

2. The setting

For the general discrete fair division problem we consider here, there is a set N of n agents and a set M of m goods that cannot be divided or shared. Each agent $i \in N$ is equipped with a *valuation function* $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$, which assigns a non-negative real number to each possible subset of items and is *normalized* and *monotone*, i.e., $v_i(\emptyset) = 0$ and $v_i(S) \leq v_i(T)$ for all $S \subseteq T \subseteq M$. In this survey, we mostly focus on the case where the valuation function of each agent i is also assumed to be *additive*, so that $v_i(S) = \sum_{g \in S} v_i(g)$ for any subset of items $S \subseteq M$, where $v_i(g)$ is used as a shorthand for $v_i(\{g\})$. Other types of valuation functions have also been studied and are briefly discussed but, unless otherwise specified, in what follows, we refer to the additive case. For the sake of readability, we avoid defining all different function classes mentioned here; we define submodular, XOS, and subadditive valuation functions in Section 8.2, and we refer the interested reader to the corresponding references for the others. A fair allocation instance is denoted by $I = (N, M, v)$ where $v = (v_1, \dots, v_n)$ is the vector of valuation functions and can be represented by a table with a row per agent and a column per good, such that cell (i, g) contains the value $v_i(g)$. An *allocation* is a tuple of subsets of M , $A = (A_1, \dots, A_n)$, such that each agent $i \in N$ receives the *bundle*, $A_i \subseteq M$, $A_i \cap A_j = \emptyset$ for every pair of agents $i, j \in N$, and $\bigcup_{i \in N} A_i = M$. If $\bigcup_{i \in N} A_i \subsetneq M$, the allocation is called *partial*.

2.1. Solution concepts

The objective is to compute a *fair* allocation, i.e., an allocation that satisfies a desired fairness criterion. As already mentioned, since the early fair division literature, there are two predominant fairness notions, namely *envy-freeness* and *proportionality*. An allocation is said to be *envy-free* if no agent believes that another agent was given a better bundle; note that envy-freeness depends on pairwise comparisons.

Definition 1 (Envy-freeness). An allocation A is *envy-free* (EF) if $v_i(A_i) \geq v_i(A_j)$ for every pair of agents $i, j \in N$.

On the other hand, an allocation is said to be *proportional* if each agent is guaranteed her *proportional share* in terms of the total value, independently of what others get.

Table 1

The values of the agents in the instance considered in Example 1. In particular, the number in the cell defined by row $i \in \{a_1, a_2, a_3\}$ and column $g \in \{g_1, g_2, g_3, g_4\}$ is the value $v_i(g)$ of agent i for good g .

	g_1	g_2	g_3	g_4
a_1	10	6	6	8
a_2	10	5	5	10
a_3	10	0	10	10

Table 2

The instance considered in Examples 2, 3, and 4.

	g_1	g_2	g_3	g_4	g_5
a_1	15	3	2	2	6
a_2	7	5	5	5	7
a_3	20	3	3	3	3

Definition 2 (Proportionality). An allocation A is *proportional* (PROP) if $v_i(A_i) \geq v_i(M)/n$ for every agent $i \in N$.

It is not hard to see that an EF allocation is also PROP, but the converse is not necessarily true.

Example 1. Consider an instance with three agents and four goods. The values of the agents for the goods are given in Table 1.

The allocation $A = \{A_1, A_2, A_3\}$ with $A_1 = \{g_2, g_3\}$, $A_2 = \{g_4\}$ and $A_3 = \{g_1\}$ is EF. In particular, agent a_1 does not envy any other agent as $v_1(A_1) = 12 > v_1(A_2) = 8$ and $v_1(A_1) = 12 > v_1(A_3) = 10$, agent a_2 does not envy as $v_2(A_2) = 10 = v_2(A_1) = v_2(A_3)$, and agent a_3 does not envy as $v_3(A_3) = 10 = v_3(A_1) = v_3(A_2)$.

On the other hand, the allocation $B = \{B_1, B_2, B_3\}$ with $B_1 = \{g_1\}$, $B_2 = \{g_2, g_3\}$ and $B_3 = \{g_4\}$ is not EF, but it is PROP. Indeed, agent a_1 envies agent 2: $v_1(B_1) = 10 < v_1(B_2) = 12$. However, $v_i(B_i) \geq v_i(M)/3 = 10$ for every agent a_i , and thus each agent obtains her proportional share. \square

As mentioned above, EF and PROP allocations do not always exist when allocating indivisible items. The simplest example is the case of two agents and a single good that is positively valued by both agents. Since only one of the agents receives the good, the other agent gets zero value, she envies the agent with the item, and does not achieve her proportional share. See also Example 2 for another instance that does not admit any EF or PROP allocation. Despite this impossibility, one could still be interested in finding EF or PROP allocations *when* they exist. Unfortunately, it turns out that the problem of even deciding whether an instance admits an EF or PROP allocation is NP-complete, which can be shown via a simple reduction from PARTITION [133]. These straightforward impossibility results have led to the definition of multiple relaxations of these two notions, tailored for discrete fair division.

2.2. Important relaxations

The first relaxation of envy-freeness is *envy-freeness up to one good* (EF1), implicitly introduced by Lipton et al. [133], but formally defined by Budish [66]. According to EF1 it is acceptable for an agent i to envy another agent j , as long as there exists a good in j 's bundle the hypothetical removal of which would eliminate i 's envy towards j .

Definition 3 (EF1). An allocation A is *envy-free up to one good* (EF1) if, for every pair of agents $i, j \in N$, it holds that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ for some $g \in A_j$.

Example 2. To demonstrate the notion of EF1 (as well as EFX and MMS later on), let us consider a simple example with three agents and five goods. The values of the agents for the goods are given in Table 2.

This instance does not admit any EF or PROP allocations. To see this, observe that in any PROP allocation, agent a_3 must get at least $\{g_1\}$ or $\{g_2, g_3, g_4, g_5\}$. In the latter case, at least one of a_1 and a_2 will get no goods, whereas in the former case a_1 must get at least three of the remaining four goods and a_2 must get at least two, which is not possible. On the other hand, note that the allocation $A_1 = \{g_3, g_4\}$, $A_2 = \{g_2, g_5\}$, $A_3 = \{g_1\}$ is EF1: a_2 and a_3 are not envious, and the envy of a_1 towards a_2 and a_3 can be eliminated by the hypothetical removals of g_5 from A_2 and g_1 from A_3 , respectively. \square

As we will see, EF1 is easy to achieve, even when the valuation functions are general monotone. However, in many cases it is a fairly weak fairness notion; an EF1 allocation is considered to be fair for an agent even when a very highly-valued good is hypothetically removed from another agent's bundle (e.g., a house or an expensive car). For example, consider agent a_1 's perspective of the allocation in Example 2, where the proposed EF1 solution requires the removal of rather valuable

goods for the agent. A very natural refinement of the notion is the stricter relaxation of *envy-freeness up to any good (EFX)* that was introduced in 2016 by Caragiannis et al. [71] in the conference version of their work but also somewhat earlier by Gourvès et al. [113] under the name *near envy-freeness*. An allocation is said to be EFX if the envy of an agent i towards another agent j can be eliminated by the hypothetical removal of *any* good in j 's bundle.

Definition 4 (EFX). An allocation A is *envy-free up to any good (EFX)* if, for every pair of agents $i, j \in N$, it holds that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ for any $g \in A_j$.

Example 3. Consider the instance of Example 2 again. The allocation $A_1 = \{g_3, g_4\}, A_2 = \{g_2, g_5\}, A_3 = \{g_1\}$ is not EFX since the envy of a_1 towards a_2 cannot be eliminated by removing g_2 (a_1 's least favorite good in A_2) from A_2 . Nevertheless, it is easy to modify this allocation to get $B_1 = \{g_4, g_5\}, B_2 = \{g_2, g_3\}, B_3 = \{g_1\}$ that is EFX. Indeed, the envy of a_1 towards a_3 can be eliminated by removing g_1 from B_3 , whereas the envy of a_2 towards a_1 can be eliminated by removing g_4 from B_1 ; in both cases, the hypothetical removal involves the envious agent's least valued good in the other agent's bundle. \square

In contrast to EF1, the existence of EFX allocations is a challenging open problem. In fact, [144] referred to this as "fair division's most enigmatic question". In the past few years, a sequence of works have partially or approximately answered this question; see Section 4.

Besides the two additive relaxations of envy-freeness discussed so far, an extensively studied fairness notion in discrete fair division is *maximin share fairness*, also introduced by [66]. The notion can be seen as a generalization of the rationale of the well-known cut-and-choose protocol, which is known to guarantee an envy-free partition of a divisible resource among two agents. Here, the goal is to give each agent i goods of value at least as much as her *maximin share* $\mu_i^n(M)$, which is the maximum value this agent could guarantee for herself by partitioning the set of goods M into n disjoint bundles and keeping the worst of them. As such, it is a relaxation of proportionality.

Definition 5 (MMS). Let $\mathcal{A}_n(M)$ be the collection of possible allocations of the goods in M to n agents. An allocation A is said to be *maximin share fair (MMS)* if for each agent $i \in N$,

$$v_i(A_i) \geq \mu_i^n(M) = \max_{B \in \mathcal{A}_n(M)} \min_{S \in B} v_i(S).$$

When M is clear from context, we may simply write μ_i^n for $\mu_i^n(M)$.

Example 4. Returning to the instance of Example 2, we can see that $\mu_1^3(M) = 6$, since it is not possible to partition the items into three sets with strictly more value, but 6 is guaranteed by the partition $\{g_1\}, \{g_2, g_3, g_4\}, \{g_5\}$. Similarly, $\mu_2^3(M) = 7$ and $\mu_3^3(M) = 6$. Therefore, $B_1 = \{g_4, g_5\}, B_2 = \{g_2, g_3\}, B_3 = \{g_1\}$, from Example 3, is an MMS allocation, but $A_1 = \{g_3, g_4\}, A_2 = \{g_2, g_5\}, A_3 = \{g_1\}$, from Example 2, is not as agent a_1 gets a bundle of value only $4 = 2/3 \cdot \mu_1^3(M)$. \square

While it is relatively easy to see that computing MMS allocations or even computing the maximin share of an agent is an NP-hard problem using a reduction from PARTITION, there is a PTAS for the latter task [156]. Since the introduction of the notion, the existence of MMS allocations was a very intriguing open problem. This was eventually answered in the negative by Kurokawa et al. [126,125], who proved that MMS allocations do not always exist when there are more than two agents. Still, it is possible to compute *approximate* MMS allocations; see Section 5.

3. Envy-freeness up to one good (EF1)

We start our discussion by presenting important results about EF1 allocations.

3.1. Computing EF1 allocations

There are several simple, polynomial-time algorithms for computing EF1 allocations. The simplest one is Round-Robin (Algorithm 1), which allocates the goods to the agents in multiple rounds. In each round, according to some arbitrarily fixed given ordering, each agent chooses her most valuable good among the available ones (that is, from the set of goods that have not been chosen by the agent's turn). It is not hard to see that Round-Robin leads to an EF1 allocation [71].

Theorem 1 ([71]). Round-Robin computes an EF1 allocation.

Proof. Consider two agents i and j , such that i chooses before j according to the given fixed ordering. As i has the chance to pick a good before j in every single round of the algorithm, i cannot envy j . Of course, agent j may envy agent i . Let g be the first good chosen by i . From that point on, we can consider the execution of the algorithm on the remaining goods as a fresh run where now j has the chance to pick a good before i in every round. So, j does not envy i 's bundle after the removal of good g from it. \square

Algorithm 1: Round-Robin.

```

1 Input: A fair allocation instance  $I = (N, M, v)$  with  $n$  agents and  $m$  goods.
2 Output: Allocation  $A = (A_1, \dots, A_n)$ .
3 for each agent  $i \in N$  do
4    $A_i \leftarrow \emptyset$ ;
5 end
6 for  $\ell = 1, \dots, m$  do
7   Let  $i \leftarrow \ell \bmod n$ ;
8   Let  $g^* \in \arg \max_{g \in M} v_i(g)$ ;
9    $A_i \leftarrow A_i \cup \{g^*\}$ ;
10   $M \leftarrow M \setminus \{g^*\}$ ;
11 end

```

In the basic definition of Round-Robin, the agents follow the same order in each round to select the goods. However, to compute an EF1 allocation, this is not necessary; as long as each agent selects her favorite good when it is her turn, the order of the agents in different rounds can be different. Essentially, Round-Robin is only a member of a larger class of algorithms that compute EF1 allocations by using *recursively balanced* sequences in different rounds, where the difference between the number of turns of any two agents is at most 1.

Another algorithm for computing an EF1 allocation is the Envy-Cycle Elimination (Algorithm 2), introduced by Lipton et al. [133]. In contrast to Round-Robin or any other sequential allocation algorithm, Envy-Cycle Elimination does not use a prefixed sequence for agents to select goods. Instead, it repeatedly chooses an agent that is in a disadvantage compared to other agents and gives an unallocated good. In the variant we present in Algorithm 2, whenever an agent gets a new good, this is her favorite available one. The algorithm maintains an *envy graph*, where the nodes correspond to agents and there is an edge from agent i to agent j if i is envious of j 's bundle. At each step of the algorithm, an unassigned good is allocated to some agent who is not envied by any other agent, i.e., an agent that corresponds to a node with in-degree 0 in the envy graph. If no such agent exists, the envy graph must contain a directed cycle, which can be eliminated by redistributing the current bundles among the agents that participate in the cycle. Formally, let $C = (i_1, \dots, i_d)$ be a directed cycle in the envy graph such that i_j envies i_{j+1} for each $j \in [d - 1]$, and i_d envies i_1 . The cycle can be resolved by exchanging the bundles of items along the cycle: each agent in the cycle gets the bundle of the agent she points to (Equation (1) in the description of Algorithm 2). Repeating this procedure, eventually, leads to a modified envy graph with at least one agent who is not envied by any other agent. The algorithm terminates when all items are allocated.

Algorithm 2: Envy-cycle elimination.

```

1 Input: A fair allocation instance  $I = (N, M, v)$  with  $n$  agents and  $m$  goods.
2 Output: Allocation  $A = (A_1, \dots, A_n)$ .
3 for each agent  $i \in N$  do
4    $A_i \leftarrow \emptyset$ ;
5 end
6 for  $\ell = 1, \dots, m$  do
7   while there does not exist an unenvied agent do
8     Find an envy-cycle  $C = (i_1, \dots, i_d)$  and resolve the cycle as follows:
          
$$A_{i_j}^C = \begin{cases} A_{i_{j+1}} & \text{for all } 1 \leq j \leq d - 1 \\ A_{i_1} & \text{for } j = d \end{cases} \quad (1)$$

           $A_i \leftarrow A_i^C$  for all  $i \in C$ ;
9   end
10  Let  $i$  be an unenvied agent;
11  Let  $g^* \in \arg \max_{g \in M} \{v_i(g)\}$ ;
12   $A_i \leftarrow A_i \cup \{g^*\}$ ;
13   $M \leftarrow M \setminus \{g^*\}$ ;
14 end

```

Lipton et al. [133] proved that Envy-Cycle Elimination runs in polynomial time and outputs an EF1 allocation. Indeed, in each round of the algorithm, an agent receives a good only if she is not envied by any other agent. Therefore, by removing the last good that an agent receives eliminates any possible envy of other agents towards her. Note that to ensure EF1, we can assign any good to agent i (instead of the item g^* with maximum value) in line 11. However, as we will see later, by assigning the most valuable good, we can ensure other nice properties regarding EFX and MMS.

3.2. EF1 and Pareto optimal (PO) allocations

Besides computing a fair allocation, another natural requirement is that of efficiency. Caragiannis et al. [71] identified an interesting inherent connection between EF1, Pareto optimality, and the notion of *maximum Nash welfare* (MNW).

Definition 6. An allocation A is said to be Pareto optimal (PO) if there is no allocation B such that $v_i(B_i) \geq v_i(A_i)$ for all $i \in N$ and $v_j(B_j) > v_j(A_j)$ for some $j \in N$. Equivalently, we will say that such an allocation is not *Pareto dominated* by any other allocation.

Definition 7 (MNW allocation). An allocation A is said to be a *maximum Nash welfare (MNW) allocation* if (a) it maximizes the number of agents with positive value, and (b) it maximizes the *Nash welfare*, defined as the product of values $\prod_i v_i(A_i)$ for the agents with positive value.¹

Theorem 2 (Caragiannis et al. [71]). Every MNW allocation is EF1 and PO.

Proof. It is straightforward to show that every MNW allocation A is PO. If there were a different allocation B with $v_i(B) \geq v_i(A)$ for all $i \in N$ and $v_j(B_j) > v_j(A_j)$ for some $j \in N$ it would have strictly larger Nash welfare compared to that of A , which would contradict that A is an MNW allocation.

To show that A is EF1, suppose to the contrary that there exist agents i and j such that $v_i(A_i) < v_i(A_j \setminus \{g\})$ for every $g \in A_j$. We will show that there exists another allocation A' with Nash welfare that is strictly larger than that of A . Since agent i envies agent j , there exists a $g^* = \arg \min_{g \in A_j, v_i(g) > 0} \left\{ \frac{v_j(g)}{v_i(g)} \right\}$. Using this, we define the allocation A' with $A'_\ell = A_\ell$ for each $\ell \in N \setminus \{i, j\}$, $A'_i = A_i \cup \{g^*\}$ and $A'_j = A_j \setminus \{g^*\}$. To show that the Nash welfare of A' is strictly larger, it suffices to show that

$$v_i(A'_i) \cdot v_j(A'_j) > v_i(A_i) \cdot v_j(A_j).$$

By the definition of g^* we have $\frac{v_j(g^*)}{v_i(g^*)} \leq \frac{v_j(A_j)}{v_i(A_j)}$, which implies

$$\frac{v_j(g^*)}{v_j(A_j)} \leq \frac{v_i(g^*)}{v_i(A_j)} < \frac{v_i(g^*)}{v_i(A_i) + v_i(g^*)}, \quad (2)$$

where the last inequality follows from the fact that agent i envies agent j even after the removal of any item in A_j . Therefore, we have

$$\begin{aligned} v_i(A'_i) \cdot v_j(A'_j) &= (v_i(A_i) + v_i(g^*)) \cdot (v_j(A_j) - v_j(g^*)) \\ &= \left(1 + \frac{v_i(g^*)}{v_i(A_i)}\right) \cdot \left(1 - \frac{v_j(g^*)}{v_j(A_j)}\right) \cdot v_i(A_i) \cdot v_j(A_j) \\ &> \frac{v_i(A_i) + v_i(g^*)}{v_i(A_i)} \cdot \left(1 - \frac{v_i(g^*)}{v_i(A_i) + v_i(g^*)}\right) \cdot v_i(A_i) \cdot v_j(A_j) \\ &= v_i(A_i) \cdot v_j(A_j), \end{aligned}$$

where the first inequality follows from Inequality (2). Hence, the Nash welfare of A' is strictly larger than that of A , a contradiction. \square

This result shows that there exist allocations that combine fairness with (Pareto) efficiency. Yuen and Suksompong [157] recently showed that Nash Welfare is the only welfarist function of agent valuations whose maximization leads to allocations that are EF1 and PO.

However, computing MNW allocations is not an easy task, as it is known that they are generally hard to even approximate in polynomial time; see the relevant discussion in Section 7.1. Barman et al. [42] recently made progress by computing such allocations in pseudo-polynomial time, but it still remains unknown whether this can be done in polynomial time.

We now mention our first open question.

Open Problem 1. Can an EF1 and PO allocation be computed in polynomial time?

4. Envy-freeness up to any good (EFX)

In contrast to EF1, where the existence is guaranteed via simple polynomial-time algorithms, the existence of EFX allocations is a challenging open problem. In the past few years, a sequence of works has positively answered this question for important special cases, centered mainly around two axes: a small number of agents or restricted agent valuations. In this

¹ Defining the Nash welfare as the product of values is a simplification. Typically, the Nash welfare of an allocation A is defined as the geometric mean of the values of the agents, that is, $(\prod_{i \in N} v_i(A_i))^{1/n}$. Note that exactly maximizing the geometric mean of values is equivalent to exactly maximizing the product of values, but this is not true when discussing about approximation.

section, we first review the existence of EFX allocations for these special cases, and then discuss results about approximate versions of EFX.

4.1. Existence and computation of EFX allocations for special cases

We first discuss a few cases for which EFX allocations are known to exist.

4.1.1. Identical valuations

The first positive existence result was shown by Plaut and Roughgarden [142] for when the agents have identical valuations, that is, $v_i(\cdot) = v(\cdot)$ for every agent i . In particular, Plaut and Roughgarden showed that allocations that satisfy particular properties about the minimum value are EFX.

A *leximin* allocation is an allocation that first maximizes the minimum value among all bundles, then maximizes the second minimum value, and so on. While a *leximin* allocation is not EFX, Plaut and Roughgarden [142] showed that a slight refinement of it is in fact EFX. The fix is that after maximizing the minimum value, the allocation must maximize the size of the bundle with minimum value, before maximizing the second minimum value, and so on. Such an allocation is called *leximin++*.

It is not hard to see that any *leximin++* allocation A must be EFX. Let $v(A_1) \leq \dots \leq v(A_n)$ and suppose towards the contrary that A is not EFX; thus, there exist $i < j$ such that $v(X_i) < v(X_j \setminus \{g\})$ for any $g \in X_j$. Then, either the i -th minimum value can be strictly increased (if $v(g) > 0$) or the size of this bundle can be increased (if $v(g) = 0$), which contradicts that A is *leximin++*. Note that a *leximin++* allocation is EFX even when the common valuation function is not additive.

4.1.2. Ordered valuations

Beyond identical valuations, Plaut and Roughgarden [142] also showed that the Envy-Cycle Elimination algorithm returns an EFX allocation for the slightly more general class of *ordered instances*, where the agents have identical orderings, but possibly distinct cardinal valuations. In particular, in such instances, there is an ordering of the goods g_1, \dots, g_m such that $v_i(g_1) \geq \dots \geq v_i(g_m)$ for every agent $i \in N$. The underlying intuition is that an agent can choose a good only when she is not envied by the others, and the newly added good is the least-valued (among the available ones) for every agent, which means that the removal of this good can eliminate any envy of the other agents towards the choosing agent.

Theorem 3 ([142]). *The Envy-Cycle Elimination Algorithm (Algorithm 2) computes an EFX allocation for every ordered instance.*

Proof. Consider any agent j and let g be the last good agent j is allocated. For any other agent i , we have $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ because when good g is allocated to agent j , j is not envied by any other agent. Since the instance is ordered, for every good $g' \in A_j \setminus \{g\}$ (that is allocated before g), we have $v_i(g') \geq v_i(g)$. Note that while some items in A_j might be swapped in the envy cycle elimination procedure, we can still ensure $v_i(g') \geq v_i(g)$ because the instance is ordered and items with larger value are allocated first. Hence, for every $g' \in A_j$,

$$v_i(A_i) \geq v_i(A_j \setminus \{g\}) = v_i(A_j) - v_i(g) \geq v_i(A_j) - v_i(g') = v_i(A_j \setminus \{g'\}).$$

In other words, the allocation is EFX. \square

4.1.3. EFX for two and three agents

When there are only two agents, we can compute an EFX allocation using a *divide-and-choose* approach that exploits the existence of EFX allocations for agents with identical valuations.

Theorem 4 ([142]). *EFX allocations always exist for instances with two agents.*

Proof. We fix any agent i and compute an EFX allocation (consisting of two bundles), pretending that both agents have the same valuation function v_i . For additive functions we can use the Envy-Cycle Elimination algorithm, as we have shown in Theorem 3; for non-additive functions we can use the *leximin++* allocation. The other agent is allocated her favorite among the two bundles, and the remaining bundle is allocated to agent i . Clearly, since the other agent gets her favorite bundle, she does not envy. In addition, since the allocation used for the definition of the two bundles is EFX to agent i , no matter which bundle is left for agent i , she must be EFX towards the other agent. \square

Goldberg et al. [111] recently showed that it is possible to efficiently compute EFX allocations when there are only two agents for some more general valuation functions such as gross substitute or budget-additive, but the problem becomes PLS-complete for submodular functions.

For instances with three agents, in a breakthrough paper, Chaudhury et al. [76] presented a procedure that computes an EFX allocation in pseudo-polynomial time. The original proof of Chaudhury et al. is quite involved and consists of careful and complex case analysis; a slightly simpler proof was recently given by Akrami et al. [1]. Computing EFX allocations for three agents in polynomial time is still an open problem.

4.1.4. Bi-valued valuations

For bi-valued instances, in which there are only two distinct possible values that an agent may have for the goods, Amanatidis et al. [10] showed that EFX allocations exist and can be efficiently computed for any number of agents. Later, Garg and Murhekar [101] showed that this is possible even in conjunction with PO. In fact, Amanatidis et al. [10] also demonstrated an interesting connection between EFX and MNW allocations for bi-valued instances, by showing that every MNW allocation implies EFX. A similar result was later shown by [30] for general valuations with binary (i.e., 0 or 1) marginals.

4.1.5. Discussion

It is instructive to mention that the original definition of EFX due to Caragiannis et al. [71] was a bit weaker and required the removed good to be positively-valued. In the case of binary valuations, a special case of bi-valued instances, the distinction between the two versions of EFX does make a difference, but in more general settings the existence and computation of the weak version can be reduced to the existence and computation of the stronger version; see Amanatidis et al. [10] for a related discussion. It is easy to see that EF implies EFX, which implies weak EFX, which in turn implies EF1.

The aforementioned results are positive first steps towards showing the existence of EFX allocations, but a general positive (or negative) result remains elusive. This brings us to our second open problem, which is one of the most important open problems in discrete fair division.

Open Problem 2. Do EFX allocations exist for instances with $n \geq 4$ agents and unrestricted additive valuations?

4.2. Relaxations of EFX

Instead of focusing on exact EFX allocations, a growing line of work has taken a different approach by aiming to compute allocations that are *approximately* EFX, for different notions of approximation. The first such notion is in terms of multiplicative approximations to the values obtained by the agents.

Definition 8 (α -EFX). Let $\alpha \in (0, 1]$. An allocation A is α -EFX if, for every pair of agents $i, j \in N$, $v_i(A_i) \geq \alpha \cdot v_i(A_j \setminus \{g\})$ for any $g \in A_j$.

Plaut and Roughgarden [142] were the first to pursue this, showing that 1/2-EFX allocations always exist, even for sub-additive valuation functions; later, Chan et al. [75] showed that computing such allocations can be done in polynomial time. Indeed, it can be shown that for agents with additive valuation functions, the Envy-Cycle Elimination algorithm (Algorithm 2) computes a 1/2-EFX allocation, as long as we break ties in favor of empty bundles when picking an unenvied agent.

Theorem 5 ([75]). *The Envy-Cycle Elimination algorithm computes a 1/2-EFX allocation.*

Proof. Consider any two agents i and j . We must show that $v_i(A_i) \geq 1/2 \cdot v_i(A_j \setminus \{g\})$ for any $g \in A_j$. The statement holds trivially if $|A_j| = 1$. So in the following, we assume that $|A_j| \geq 2$.

Let g^* be the last good allocated to agent j . Since we break ties in favor of empty bundles when picking an unenvied agent, it must be the case that $A_i \neq \emptyset$, as otherwise g^* would not be allocated to j , but to i . Since $A_i \neq \emptyset$, there must be a round before g^* is allocated during which agent i gets to pick an item. Such a round exists because empty bundles will not be involved in the envy cycle elimination procedure. Moreover, while items in A_j might be swapped a few times during the envy cycle elimination, the value of i for A_j is non-decreasing. Therefore, $v_i(A_i) \geq v_i(g^*)$. By the design of the algorithm, we have that $v_i(A_i) \geq v_i(A_j \setminus \{g^*\})$. Combining the two lower bounds, we have

$$2 \cdot v_i(A_i) \geq v_i(A_j) \geq v_i(A_j \setminus \{g\}), \quad \forall g \in A_j,$$

and thus the allocation is 1/2-EFX. \square

The approximation ratio for the additive case was further improved by Amanatidis et al. [9] to $\phi - 1 \approx 0.618$ by combining Round-Robin and Envy-Cycle Elimination with some appropriate pre-processing. To this end, we have the following open question:

Open Problem 3. What is the best possible α for which α -EFX allocations exist?

A positive answer to Open Problem 2 would establish that $\alpha = 1$ in Open Problem 3, but a negative answer would make the latter open problem very meaningful in its own right. Additionally, as is the case for all of these notions, the next natural question is whether existence can be paired with polynomial-time algorithms for finding such allocations, or whether some kind of computational hardness can be proven.

Another recent approach is that of relaxing the requirement to allocate *all* available goods. Clearly, if done without any constraints, this makes the problem trivial: simply leaving all goods unallocated results in an envy-free allocation. However, the objective here is to only leave “a few” goods unallocated (e.g., donate them to charity instead), or remove some goods without affecting the maximum possible Nash social welfare by “too much”. On this front, Caragiannis et al. [70] showed that it is possible to compute an EFX allocation of a subset of the goods, the Nash welfare of which is at least half of the maximum Nash welfare on the original set. Chaudhury et al. [80] presented an algorithm that computes a partial EFX allocation, but the number of unallocated goods is at most $n - 1$, and no agent prefers the set of these goods to her own bundle. The latter result was recently improved by Berger et al. [50] who showed that the unallocated goods can be decreased to $n - 2$ in general, and to just one for the case of four agents. Finally, Chaudhury et al. [78] showed that a $(1 - \varepsilon)$ -EFX allocation with a sublinear number of unallocated goods and high Nash welfare can be computed in polynomial time for every constant $\varepsilon \in (0, 0.5)$. This motivates the next open problem.

Open Problem 4. Is it possible to achieve an exact EFX allocation by donating a sublinear number of goods?

5. Maximin Share Fairness (MMS)

As we have mentioned, although the fairness notion of MMS significantly weakens the requirement of PROP, MMS allocations are still not guaranteed to exist. Therefore, the majority of the literature has focused on approximate MMS allocations.

Definition 9 (α -MMS). Let $\alpha \in (0, 1]$. An allocation A is α -MMS if $v_i(A_i) \geq \alpha \cdot \mu_i^n(M)$ for all $i \in N$.

Before introducing the state-of-the-art results, we first state some properties about MMS allocations that are helpful for designing approximation algorithms. The first property is the monotonicity of the MMS values; specifically, if we arbitrarily remove one agent and one good from the instance, the MMS value of each of the remaining agents in the remaining instance does not decrease. This property was first formally stated by Bouveret and Lemaître [62] (as the $k = 1$ case of Lemma 1 therein) and since then several variants have appeared in the literature. Here, we state the version of Amanatidis et al. [7].

Lemma 1 ([62,7]). For any agent i and any good g , $\mu_i^{n-1}(M \setminus \{g\}) \geq \mu_i^n(M)$.

Proof. Let (A_1, \dots, A_n) be an n -partition of M that achieves $\mu_i^n(M)$, i.e., $v_i(A_j) \geq \mu_i^n(M)$ for all $j \in [n]$. Without loss of generality, suppose $g \in A_n$. We arbitrarily reallocate the goods in $A_n \setminus \{g\}$ (if any) to the sets A_1, \dots, A_{n-1} , and denote the resulting allocation by (B_1, \dots, B_{n-1}) ; clearly, this is a partition of $M \setminus \{g\}$ into $n - 1$ bundles. Since $v_i(B_j) \geq v_i(A_j) \geq \mu_i^n(M)$ for all $j \in [n - 1]$, it follows that $\mu_i^{n-1}(M \setminus \{g\}) \geq \mu_i^n(M)$. \square

Lemma 1 essentially suggests that if there is a large good g that satisfies the desired approximation of MMS value for some agent i , we can simply assign g to i and finalize the bundle $A_i = \{g\}$ without hurting the remaining agents in the reduced instance.

5.1. Computing 1/2-MMS allocations

Here, we consider three simple algorithms for computing 1/2-MMS allocations in polynomial time. Better approximations can be obtained by variants and extensions of the techniques employed by these algorithms; see Section 5.2.

The monotonicity of the MMS value (Lemma 1) allows us to focus only on the case when the agents have sufficiently small values for the goods, in particular, when $v_i(g) \leq \mu_i^n/2$ for every $i \in N$ and $g \in M$. Specifically, as long as there is an agent i and a good g such that $v_i(g) > \mu_i^n/2$, we can allocate g to i so that i obtains half of her MMS value, and never consider i or g again. If we can guarantee 1/2-MMS for the remaining agents in the reduced instance, the same guarantee must also hold in the original instance by Lemma 1. One way to achieve this is by using Round-Robin (see Algorithm 1) [7].

Lemma 2. If $v_i(g) \leq 1/2 \cdot \mu_i^n$ for every $i \in N$ and $g \in M$, the Round-Robin algorithm returns an allocation that is 1/2-MMS.

Proof. Let $A = (A_1, \dots, A_n)$ be the allocation returned by the Round-Robin algorithm. If $m < n$, then $\mu_i^n = 0$ for every $i \in N$ and thus any allocation is MMS. So, assume that $m \geq n$. By the design of the algorithm, for each agent i , the item selected in any round must be at least as good as the items allocated to any agent in subsequent rounds. For any $i \in N$, let g_i be the good selected by agent i in the first round. Now let i be an arbitrary agent. Then, $v_i(A_i) \geq v_i(A_j \setminus \{g_j\})$ for any $j \in N$. By summing these inequalities over $j \in N$, we obtain

$$n \cdot v_i(A_i) \geq v_i(M) - \sum_{j \in N} v_i(g_j) \geq n \cdot \mu_i^n - \frac{n}{2} \cdot \mu_i^n = \frac{n}{2} \cdot \mu_i^n,$$

where the second inequality follows since $v_i(M) \geq n \cdot \mu_i^n$ and $v_i(g_j) \leq \mu_i^n/2$ for all $j \in N$. Therefore, for any agent i , we have that $v_i(A_i) \geq \mu_i^n/2$, that is, she obtains at least half of her MMS value and the allocation is 1/2-MMS. \square

It is also not hard to see that the Envy Cycle Elimination algorithm (Algorithm 2) returns an allocation that is 1/2-MMS (in addition to being EF1 and 1/2-EFX).

Theorem 6. *The Envy-Cycle Elimination algorithm computes a 1/2-MMS allocation.*

Proof. Consider any agent $i \in N$. As we have shown in the proof of Theorem 5, for every agent $j \neq i$, we have either $|A_j| = 1$ or $v_i(A_j) \leq 2 \cdot v_i(A_i)$. Let $N_1 = \{j \in N \setminus \{i\} : |A_j| = 1\}$ be the set of agents other than i that are allocated a single item, and $n_1 = |N_1|$. By Lemma 1, we have $\mu_i^{n-n_1}(M \setminus (\cup_{j \in N_1} A_j)) \geq \mu_i^n(M)$. In other words, the MMS value of agent i in the instance with set of agents $N_2 = N \setminus N_1$ and set of items $M \setminus (\cup_{j \in N_1} A_j)$ is at least as large as $\mu_i^n(M)$. Since $v_i(A_j) \leq 2 \cdot v_i(A_i)$ for every $j \in N_2$, we have that $v_i(\cup_{j \in N_2} A_j) \leq 2 \cdot |N_2| \cdot v_i(A_i)$. Reordering the inequality gives

$$v_i(A_i) \geq \frac{1}{2} \cdot \frac{v_i(\cup_{j \in N_2} A_j)}{|N_2|} \geq \frac{1}{2} \cdot \mu_i^{n-n_1}(M \setminus (\cup_{j \in N_1} A_j)) \geq \frac{1}{2} \cdot \mu_i^n(M),$$

which implies that the allocation is 1/2-MMS for every agent $i \in N$. \square

The third algorithm for computing a 1/2-MMS allocation relies on a commonly used technique known as *bag-filling* [109]. Without loss of generality, and for simplicity, we assume that the MMS values of all agents are normalized to 1: $\mu_i^n = 1$ for every $i \in N$. As in the case of the first algorithm, due to Lemma 1, we can focus on the case where $v_i(g) \leq 1/2 = \mu_i^n/2$ for every $i \in N$ and $g \in M$. We start with an empty bag and keep adding goods one by one into the bag until there is an agent with value at least 1/2 for the goods in the bag. Then, we allocate all the goods in the bag to this agent. This process is repeated for the remaining agents and goods. When a single agent remains, she is given all the available goods. See Algorithm 3.

Algorithm 3: Bag-filling algorithm.

```

1 Input: A fair allocation instance  $I = (N, M, v)$  with  $|N| \geq 2$  and  $\mu_i^n(M) = 1$  for all  $i \in N$ 
2 Output: Allocation  $A = (A_1, \dots, A_n)$ 
3 for each agent  $i \in N$  do
4    $A_i \leftarrow \emptyset$ ;
5 end
6  $\text{BAG} \leftarrow \emptyset$ ;
7 while  $|N| > 1$  and there exist  $i \in N$  and  $g \in M$  such that  $v_i(g) > 1/2$  do
8    $A_i \leftarrow \{g\}$ ;
9    $N \leftarrow N \setminus \{i\}$ ;
10   $M \leftarrow M \setminus \{g\}$ ;
11 end
12 while  $|N| > 1$  do
13   while  $v_i(\text{BAG}) < 1/2$  for every  $i \in N$  do
14     Select an arbitrary good  $g \in M$ ;
15     Set  $\text{BAG} \leftarrow \text{BAG} \cup \{g\}$ ;
16     Set  $M \leftarrow M \setminus \{g\}$ ;
17   end
18   Let  $i^* \in N$  be any agent with  $v_{i^*}(\text{BAG}) \geq 1/2$ ;
19    $A_{i^*} \leftarrow \text{BAG}$ ;
20    $N \leftarrow N \setminus \{i^*\}$ ;
21    $\text{BAG} \leftarrow \emptyset$ ;
22 end
23 Let  $i$  be the remaining agent in  $N$ ;
24  $A_i \leftarrow M$ ;

```

Theorem 7. *The Bag-Filling Algorithm (Algorithm 3) returns a 1/2-MMS allocation.*

Proof. First, recall that any agent has a total value of at least n for the all the goods due to the normalization. In each round of the algorithm, the bag is allocated to an agent that has value at least 1/2 for the goods in the bag, and thus this agent obtains at least half of her MMS value. The value of each of the other agents for the given bag is smaller than her MMS value of 1: before the last good is added to the bag their value is smaller than 1/2 and each good (including the last one added to the bag) has value at most 1/2. Consequently, when an agent takes away a bag in some round, each of the remaining $n' < n$ agents has total value at least n' for the remaining goods. The argument can be repeated until there is a single agent. Then she gets all the remaining goods of value at least 1, i.e., her whole MMS value. \square

5.2. Better approximations

Kurokawa et al. [126] were the first to go beyond $1/2$ by showing how to find $2/3$ -MMS allocations, albeit not in polynomial time. Amanatidis et al. [7] matched this guarantee in polynomial time, as did Barman and Krishnamurthy [38] with a much simpler algorithm, which relies on an important property that was first proved by Bouveret and Lemaître [62] for exact MMS allocations and then generalized by Barman and Krishnamurthy [38] for approximate MMS allocations. This property is very important and suggests that the hardest case is when the agents have the same order of preference over the goods; in other words, there exists a worst-case instance that is ordered.

Given any instance $I = (N, M, v)$ which may not be ordered, we can construct a corresponding ordered instance $I' = (N, M, v')$. Based on an ordering \succ of the goods, for each agent $i \in N$, we find a permutation $\sigma_i : M \rightarrow M$ such that $v_i(\sigma_i(g)) \geq v_i(\sigma_i(g'))$ for every $g \succ g'$. Then, we define a new valuation function v'_i with $v'_i(g) = v_i(\sigma_i(g))$ for every $g \in M$; hence, the new value of i for good g that is ranked ℓ -th in \succ equals the ℓ -th highest value of i in the original valuation. Given any allocation A' for the ordered instance I' , we can obtain an allocation A for the original instance I such that every agent has value for A_i at least as much as for A'_i . Hence, if I' admits an MMS allocation then I admits an MMS allocation as well, and the same holds for approximate MMS allocations.

Lemma 3 ([62,38]). *Let $I = (N, M, v)$ be any instance, I' be its corresponding instance, and $A' = (A'_1, \dots, A'_n)$ be an allocation for I' such that $v'_i(A'_i) \geq \alpha_i$ for every $i \in N$. Then, we can construct in polynomial time an allocation $A = (A_1, \dots, A_n)$ for I such that $v_i(A_i) \geq \alpha_i$ for every $i \in N$.*

Barman and Krishnamurthy [38] showed that the Envy-Cycle Elimination algorithm actually returns a $2/3$ -MMS allocation for ordered instances, and thus a $2/3$ -MMS fair allocation can be found in polynomial time for every instance due to Lemma 3. The barrier of $2/3$ was broken by Ghodsi et al. [109] who designed an elaborate $(3/4 - \epsilon)$ -approximation algorithm. A simpler algorithm with a slightly improved approximation guarantee of $3/4 + 1/(12n)$ was then proposed by Garg and Taki [103]. On the negative side, Feige et al. [97] recently showed that it is impossible to achieve an approximation bound better than $39/40$, even for the case of three agents.

Open Problem 5. Is it possible to improve upon the bound of $3/4 + 1/(12n)$ for additive valuations? Is there a stronger inapproximability bound than $39/40$?

5.3. Restricted instances

As expected, by restricting the number of agents or the space of the valuation functions, one can get stronger results.

MMS for four or fewer agents. When there are only two agents, a simple cut-and-choose-type protocol always produces an MMS allocation. Specifically, the first agent partitions the set of goods as equally as possible (thus the worst set has value equal to her maximin share) and the second agent decides who gets each of these sets. As suggested above, the first step is computationally hard but producing a $(1 - \epsilon)$ -MMS allocation in polynomial time is still possible. Even though in the general case, the algorithm of Kurokawa et al. [126] guarantees a $2/3$ -approximation, for three or four agents it guarantees an improved $3/4$ -approximation. The approximation factor for three agents was then improved to $7/8$ [7], later to $8/9$ [112] and recently to $11/12$ [96], whereas for the case of four agents the factor was improved to $4/5$ [109].

MMS for restricted valuations. It follows by Definition 5 that MMS allocations exist for instances where all agents have identical valuation functions. As discussed above, Bouveret and Lemaître [62] showed that, unlike with EFX, the hardest instances for MMS (among all possible instances) are the ordered instances. They also suggested a simple construction of exact MMS allocations when the valuation functions are binary. This result can be generalized, as MMS allocations always exist and can be computed efficiently for ternary valuation functions [7] and for bi-valued valuation functions [95]. Ebadian et al. [90] showed that this is also the case when there are at most two values per agent (possibly not common across all agents) and for general instances where the value of each good is at least as much as the value of all less valuable goods combined. Another notable class for which exact MMS allocations are known to exist is that of matroid-rank valuations [40].

Open Problem 6. Are there other classes of structured valuations for which MMS is guaranteed to exist, e.g., when there are only a few (but more than two) possible values?

6. Further notable fairness notions

In this section, we discuss other notable fairness notions beyond EF, PROP, EF1, EFX and MMS.

6.1. EFL, EFR, and epistemic notions

The EFX notion was defined as a more realistic counterpart to EF1, however, as we have discussed in Section 4, it is still unknown if it can always be guaranteed. This has led to the definition of notions that lie in-between EF1 and EFX.

Barman et al. [41] defined the notion of *envy-freeness up to one less-preferred good* (EFL) according to which an agent i does not envy another agent j if either A_j contains at most one good that i values positively, or the envy of i can be eliminated by the hypothetical removal of a good $g \in A_j$ such that $v_i(A_i) \geq v_i(g)$. They showed that EFL allocations always exist and can be computed using a variant of the Envy-Cycle Elimination algorithm. Farhadi et al. [93] defined the notion of *envy-freeness up to a random good* (EFR) according to which the envy of an agent i towards another agent j can be eliminated in *expectation* after the hypothetical removal of a randomly chosen good from A_j . They showed that a 0.73-EFR allocation can be computed in polynomial time.

Open Problem 7. Do EFR allocations always exist?

Another recent line of research has focused on *epistemic* fairness notions motivated by cases where the agents do not have complete knowledge about the computed allocation [82,23,118,73]. For example, Aziz et al. [23] considered a setting where the agents are connected via a social network and do not know the bundles allocated to the agents who are not their neighbors. They introduced the notion of *epistemic EF* under which an agent is satisfied only if there is a (re)distribution of the goods she is not aware of (i.e., not given to her or to her neighbors) such that the resulting allocation is EF to her. Recently, Caragiannis et al. [73] proved the existence of epistemic EFX allocations for additive valuations when the underlying social network contains only isolated nodes.

6.2. PMMS and GMMS

Several variations of MMS have also been considered. Caragiannis et al. [71] defined the notion of *pairwise maximin share fairness* (PMMS) according to which, for every pair of agents i and j , i 's value for A_i must be at least as much as the maximum she could obtain by redistributing the set of goods in $A_i \cup A_j$ into two bundles and picking the worst of them. In other words, instead of requiring the maximin share guarantee to be achieved for the set of all agents, PMMS requires that it is achieved for any pair of agents.

Definition 10 (α -PMMS). Let $\alpha \in (0, 1]$. An allocation A is α -PMMS if $v_i(A_i) \geq \alpha \cdot \mu_i^2(A_i \cup A_j)$ for all $i, j \in N$. When $\alpha = 1$, the allocation is called PMMS.

Despite the apparent similarities in the definitions of PMMS and MMS, [71] showed that their exact versions are actually incomparable. The main open problem here is the following.

Open Problem 8. Do PMMS allocations always exist?

In fact, any PMMS allocation must be EFX, and thus showing the existence of PMMS allocations would also imply the existence of EFX allocations (Open Problem 2). For approximate PMMS, the best known result is 0.781 by Kurokawa [124].

An even stronger notion, which implies both MMS and PMMS, is that of *groupwise maximin share fairness* (GMMS) defined by Barman et al. [41], and which requires that the maximin share guarantee is simultaneously achieved for any possible subset of agents.

Definition 11 (α -GMMS). Let $\alpha \in (0, 1]$. An allocation A is α -GMMS if $v_i(A_i) \geq \alpha \cdot \text{GMMS}_i$ for all $i \in N$, where

$$\text{GMMS}_i = \max_{S \subseteq N: i \in S} \mu_i^{|S|} \left(\bigcup_{j \in S} A_j \right).$$

When $\alpha = 1$, the allocation is called GMMS.

Barman et al. [41] showed that GMMS allocations exist for some restricted settings, such as when the agents have binary or identical values. They also showed that any EFL allocation is $1/2$ -GMMS, and thus such an allocation can be computed efficiently. The currently best known approximation of GMMS is $4/7$ [9,80]. The implication relations between all the aforementioned notions have been used many times to show that particular algorithms have guarantees that hold for multiple notions at once. We refer the reader to the paper of Amanatidis et al. [8] for a discussion of the relations between (approximate versions of) these notions.

Open Problem 9. What is the best possible α for which α -GMMS allocations exist?

6.3. Prop1, PropX and PropM

A line of work has also focused on relaxations of proportionality that are similar in essence to EF1 and EFX. Conitzer et al. [85] defined the notion of *proportionality up to one good* (Prop1) according to which each agent could obtain her

Table 3
An instance that does not admit any PropX allocation.

	g_1	g_2	g_3	g_4	g_5
a_1	3	3	3	3	1
a_2	3	3	3	3	1
a_3	3	3	3	3	1

proportional share if given one extra good. [25] defined *proportionality up to any good* (PropX) that demands that each agent can obtain her proportional share when given the least positively-valued good among those allocated to other agents.

Definition 12 (Prop1). An allocation A is proportional up to one good (Prop1) if, for every agent $i \in N$, there exists a good $g \in M \setminus A_i$ such that $v_i(A_i \cup \{g\}) \geq v_i(M)/n$.

Definition 13 (PropX). An allocation A is proportional up to any good (PropX) if, for every agent $i \in N$ and any good $g \in M \setminus A_i$, we have $v_i(A_i \cup \{g\}) \geq v_i(M)/n$.

An allocation that is Prop1 and PO always exists [85] and can be computed in polynomial time [39]. PropX is rather demanding and it cannot always be guaranteed, as shown by Aziz et al. [25] via the following simple example (see Table 3). It can be verified that there must exist an agent i that receives at most one item. Moreover, there exists a good $g \notin A_i$ whose inclusion gives agent i a total value of 4, which is lower than $13/3$, the proportional share of agent i .

Recently, Baklanov et al. [35,36] introduced the notion of *proportionality up to the maximin good* (PropM) which provides a middle-ground between Prop1 and PropX. Informally, given an allocation A , the value of a maximin good for agent i is

$$d_i(A) = \max_{j \neq i} \min_{g \in A_j} v_i(g),$$

and A is called PropM if $v_i(A_i) + d_i(A) \geq v_i(M)/n$. Baklanov et al. [36] showed that a PropM allocation always exists and can be computed in polynomial time.

6.4. Equitability and its relaxations

Besides envy-freeness and proportionality, another important fairness notion is that of *equitability* according to which all agents must derive the same value from the bundles they are allocated.

Definition 14 (Equitability [64]). An allocation $A = (A_1, \dots, A_n)$ is *equitable* (EQ) if $v_i(A_i) = v_j(A_j)$ for every pair of agents $i, j \in N$.

It is not hard to observe that an EQ allocation may not exist, even when there is just one good and two agents. Similarly to EF1 and EFX, we can relax EQ to EQ1 and EQX.

Definition 15 (EQ1). An allocation $A = (A_1, \dots, A_n)$ is *equitable up to one good* (EQ1) if, for every pair of agents $i, j \in N$, it holds that $v_i(A_i) \geq v_j(A_j \setminus \{g\})$ for some $g \in A_j$.

Definition 16 (EQX). An allocation $A = (A_1, \dots, A_n)$ is *equitable up to any good* (EQX) if, for every pair of agents $i, j \in N$, it holds that $v_i(A_i) \geq v_j(A_j \setminus \{g\})$ for any $g \in A_j$.

Clearly an allocation that is EQX is also EQ1. Gourvès et al. [113] considered EQX allocations under the alias of *nearly jealousy-free*, and showed that such allocations (and thus also EQ1 allocations) always exist and can be computed in polynomial time.

7. Beyond fairness: efficiency and incentives

Here we consider other interesting directions like the relation between fairness and efficiency, or the possibility of achieving fairness when the agents are strategic.

7.1. Fair and Pareto optimal allocations

There is a significant line of work that considers the question of whether it is possible to simultaneously achieve fairness and *efficiency*. A common type of efficiency is that of PO (see Definition 6), which, as we already discussed, can be guaranteed in conjunction with some fairness notions, like EF1 and Prop1. In particular, Caragiannis et al. [71] showed that any

Table 4
An instance that does not admit any EQ1 and PO allocation [98].

	g_1	g_2	g_3	g_4	g_5	g_6
a_1	1	1	1	0	0	0
a_2	0	0	0	1	1	1
a_3	0	0	0	1	1	1

allocation that maximizes the Nash welfare is EF1 and PO (Theorem 2). However, since the computation of Nash welfare maximizing allocations is NP-hard, this result only shows the existence of EF1+PO allocations.

For the computation of allocations that are fair and PO, a commonly used technique in the literature is to exploit the connection between fair division and *market equilibria* in *Fisher markets* [66], which has been widely used in the design of approximation algorithms for maximizing the Nash welfare [83,84]. We briefly introduce this framework.

Similarly to a fair division instance, a Fisher market consists of a set of agents N , a set of goods M , and each agent $i \in N$ has a value $v_i(g)$ for every good $g \in M$. In contrast to discrete fair division, each agent i also has a *budget* $b_i \geq 0$, and the goods can be fractionally allocated. A *fractional allocation* is a tuple of vectors $X = (x_i)_{i \in N}$, where $x_i = (x_{i1}, \dots, x_{im})$ is the allocation of agent i ; in particular, $x_{ig} \in [0, 1]$ denotes the fraction of g that is allocated to i , such that $\sum_{i \in N} x_{ig} \leq 1$ for every $g \in M$. We will be focusing on Fisher markets with linear utilities, where, given a fractional allocation X , the value of agent i value is $v_i(x_i) = \sum_{g \in M} x_{ig} \cdot v_i(g)$. A *market outcome* is a tuple (X, p) , where X is a fractional allocation and $p = (p_1, \dots, p_m)$ defines the price p_j for item j . An outcome (X, p) is a *market equilibrium* if it satisfies the following conditions:

1. The market clears (all items are completely allocated): $\sum_{i \in N} x_{ig} = 1$ for every $g \in M$ with $p_g > 0$;
2. The budgets of all agents are exhausted: $\sum_{g \in M} p_g \cdot x_{ig} = b_i$ for every $i \in N$;
3. The agents only spend money on the goods with maximum value per unit of money spent (bang-per-buck): For any $g \in M$, $x_{ig} > 0$ implies $\frac{v_i(g)}{p_g} \geq \frac{v_i(g')}{p_{g'}}$ for every $g' \in M$.

An allocation is *fractionally Pareto optimal* (fPO) if it is not Pareto dominated by any fractional allocation. Clearly, an fPO allocation is also PO. It is known that a market equilibrium leads to an fPO allocation. Moreover, when all agents have the same budget (that is, we have a *competitive equilibrium from equal incomes* (CEEI)), the allocation is EF [154]. Therefore, if a market equilibrium gives an integral allocation, this allocation is EF and fPO. However, this rarely happens.

Many works have used Fisher markets as a proxy. For example, Barman et al. [42] and Barman and Krishnamurthy [39] compute allocations by perturbing the budgets or the values in the market so that they admit market equilibria with integral allocations. Based on this, Barman et al. [42] designed a pseudo-polynomial time algorithm to compute an allocation that is EF1 and PO, and Barman and Krishnamurthy [39] designed a strongly polynomial time algorithm to compute an allocation that is PROP1 and PO. Garg and Murhekar [102] improved the result of Barman et al. [42] by designing a pseudo-polynomial time algorithm to compute an EF1 and fPO allocation. Garg and Murhekar [101] showed that in bi-valued instances, an EFX and PO allocation can be computed in polynomial-time, but in instances with three distinct values, EFX and PO are incompatible. Freeman et al. [98] showed that, if all values are strictly positive, there always exists a EQX and PO allocation. On the negative side, using Example 5, they showed that there do not exist EQ1 and PO allocations when the values can be zero. Garg and Murhekar [102] strengthened this result by providing a pseudo-polynomial time algorithm to compute an EQ1 and PO allocation.

Example 5. Consider the instance given in Table 4. In any PO allocation, $\{g_1, g_2, g_3\}$ must be assigned to agent a_1 , and the remaining goods $\{g_4, g_5, g_6\}$ are shared between agents a_2 and a_3 . Consequently, one of a_2 and a_3 is allocated at most one good, which violates the requirement of EQ1 with agent a_1 . \square

Approximating the Nash welfare. Since MNW allocations are both EF1 and PO (see Section 3), it is worth mentioning at this point that the problem of efficiently computing approximate maximum Nash welfare allocations has been among the most active ones in computational social choice. Garg et al. [104] showed that approximating the Nash welfare within a ratio better than 1.069 is NP-hard, even for additive valuation functions with only four values; in fact, it is APX-hard to approximate the Nash welfare even when there only two values [2]. For additive valuations, the best known approximation is $e^{1/e} + \epsilon \approx 1.45$ [42]. Closing the gap is an interesting open question. It should be noted, however, that approximate Nash welfare allocations are not known to be even approximately EF1.

When all agents have submodular valuation functions (see Section 8.2 for a formal definition of submodular, XOS and subadditive functions), the best known approximation of Nash welfare is $6 + \epsilon$ [107] and the best known lower bound is 1.5819 [105]. Under the value oracle model, there is an asymptotically tight approximation ratio of $\Theta(n)$ for both XOS and subadditive functions [43,79]. With stronger oracle assumptions, a sublinear approximation ratio is possible for XOS functions [45]. Whether sublinear approximation ratios are possible for subadditive functions with stronger oracle assumption remains unknown.

Table 5
A lower bound instance for the price of EF1 for two agents.

	g_1	g_2	g_3
a_1	$1/3 - 2\varepsilon$	$1/3 + \varepsilon$	$1/3 + \varepsilon$
a_2	0	1/2	1/2

Egalitarian welfare. Another common welfare function that is considered as a fairness criterion itself is the *egalitarian welfare*, defined as the minimum utility among all agents (see, e.g., [53]). Aiming to maximize the egalitarian welfare is a known APX-hard problem known as the *Santa Claus* problem in the literature (e.g., see [37,94,129]). While egalitarian welfare maximizing allocations are clearly PO, they may not achieve any constant approximation of important relaxations of EF and Prop, such as EF1 and MMS.

7.2. Price of fairness

Another natural goal is to (approximately) maximize some objective function of the values of the agents, such as the *social welfare*, i.e., the total value of the agents for the goods they receive. Formally, the social welfare of an allocation A is $SW(A) = \sum_{i \in N} v_i(A_i)$. Bertsimas et al. [52] and Caragiannis et al. [69] defined the *Price of Fairness*, a measure which, similarly to the approximation ratio in worst-case analysis, measures the deterioration of the objective due to the fairness requirement (which may refer to any fairness notion). Given an instance I , we denote by $OPT(I)$ the maximum social welfare over all allocations of I . Let F be a fairness criterion (such as EF1 and MMS) and $F(I)$ be the set of all allocations satisfying F .

Definition 17 (Price of Fairness (PoF)). The price of fairness with respect to fairness criterion F is

$$PoF(F) = \sup_I \min_{A \in F(I)} \frac{OPT(I)}{SW(A)}.$$

If no fair allocation can achieve non-zero welfare, the price of fairness is infinite (unbounded). The price of fairness with respect to fairness criterion F is also called the *price of F* . Bei et al. [48] were the first to consider the price of fairness with respect to relaxations of EF. For instances with two agents, using Example 6, they showed that the price of EF1 is at least $8/7$.

Example 6. Consider the instance given in Table 5, where $\varepsilon > 0$ is arbitrarily small.

The optimal social welfare in this instance is $4/3 - 2\varepsilon$, achieved by allocating g_1 to a_1 and other two goods to a_2 . However, in any EF1 allocation the last two goods cannot both be given to a_2 . Hence, the social welfare of an EF1 allocation is at most $(1/3 - 2\varepsilon) + (1/3 + \varepsilon) + 1/2 = 1/6 - \varepsilon$. Taking $\varepsilon \rightarrow 0$, the price of EF1 is at least $8/7$. □

An upper bound on the price of EF1 follows by a variant of the Adjusted-Winner algorithm [64]. The idea is to sort the goods according to the ratios between the values that they yield to the two agents:

$$\frac{v_1(g_1)}{v_2(g_1)} \geq \frac{v_1(g_2)}{v_2(g_2)} \geq \dots \geq \frac{v_1(g_m)}{v_2(g_m)}.$$

Intuitively, the first goods according to this ordering are more valuable to the first agent while the last ones are more valuable to the second agent. Bei et al. [48] proved that, if a minimal set of consecutive goods from left is allocated to the first agent such that the allocation is EF1 for her and the remaining goods are given to agent 2, then the allocation is EF1 with social welfare that is within a $\sqrt{3}/2$ -fraction of the optimal social welfare.

For arbitrary number of agents, Bei et al. [48] showed that a lower bound of $\Omega(\sqrt{n})$ on the price of EF1; Barman et al. [44] managed to give a tight upper bound of $O(\sqrt{n})$. Furthermore, Barman et al. [44] showed tight bounds for other fairness notions, in particular, for $1/2$ -MMS and Prop1. Halpern and Shah [115] showed tight bounds on the price of EF1 and of approximate MMS under the constraint of having only ordinal information about the agent values, a typical assumption made in the context of *distortion* in social choice [14].

7.3. Fair division with strategic agents

Most of the papers mentioned so far, studied the discrete fair division problem from an algorithmic perspective under the assumption that the agents are non-strategic. In the *strategic* version of the problem, an agent may intentionally misreport how she values the goods in order to end up with a better bundle. This introduces an additional layer of difficulty, as the goal is to produce fair allocations *according to the true values* of the agents, while their declarations might be far from being the truth. This version of the problem has been considered mostly from a mechanism design *without monetary transfers* perspective, in which the utility of an agent is defined as her (true) value for her bundle.

A first direction was the design of *truthful* mechanisms that are also fair. A mechanism is truthful if no agent has an incentive to lie, i.e., no matter what values are reported by the other agents, reporting her true values always brings an agent value no smaller than what she obtains when reporting false ones. Caragiannis et al. [68] showed that no truthful mechanism for two agents and two goods can always output allocations of minimum envy. Amanatidis et al. [5] revisited the problem for the case of two agents and any number of goods, and showed that no truthful mechanism can always output α -MMS allocations, for $\alpha > 2/m$. A characterization of truthful mechanisms for two agents, showing that truthfulness and fairness are incompatible (in the sense that there is no truthful mechanism with bounded fairness guarantees under *any* meaningful fairness notion) was provided by [6]. This impossibility, however, does not apply to restricted cases. For binary valuations, maximizing the Nash welfare results in a polynomial-time truthful mechanism that outputs EF1 and PO allocations, see [116,56,10]. Babaioff et al. [30] showed an analogous result with respect to MMS. In fact, Babaioff et al. also showed that for the submodular analog of binary valuations, there is a truthful mechanism that always outputs EFX allocations.

The aforementioned impossibility results led to a different direction, where the focus was shifted to the stable states of non-truthful mechanisms. In particular, [11] studied mechanisms that always have *pure Nash equilibria*, and showed that every allocation that corresponds to an equilibrium of Round-Robin is EF1 with respect to the (unknown) true values of the agents. Qualitatively similar results can be obtained for approximate pure Nash equilibria, even for agents with submodular valuation functions [13]. Bouveret and Lang [61] and Aziz et al. [21] studied the strategic setting for general sequential allocation algorithms.

Open Problem 10. Are there mechanisms that always have pure Nash equilibrium allocations with stronger guarantees than EF1?

Another direction to escape the impossibility is to relax the requirements of truthfulness to *not obviously manipulability* [145]. A mechanism is not obviously manipulable (NOM) if no agent can increase her best- and worst-case value by lying. Fortunately, Round-Robin algorithm is NOM, and thus EF1 and NOM are compatible. Furthermore, Psomas and Verma [145] showed that we can achieve EF1, PO and NOM simultaneously.

8. Different settings

In this last section, we briefly discuss further meaningful discrete fair division settings.

8.1. Limited information

There is also an increasing interest in settings with partial information, and particularly when only ordinal information is available. Even though the agents have cardinal values for the goods, the algorithm may only have access to each agent's ranking over the goods (induced by the values in non-increasing order). Given only the ordinal preferences of the agents, we cannot run algorithms such as Envy-Cycle Elimination which need to calculate and compare the values of the agents. Instead, sequential-picking algorithms become more powerful as they do not directly rely on the values.

It is not hard to obtain an EF1 allocation as the Round-Robin algorithm only needs to know which of the available goods every agent values the most in each round. However, it becomes harder to compute approximate MMS allocations. Let $H_n = \Theta(\ln n)$ be the n -th harmonic number. Halpern and Shah [115] showed that with only ordinal information, it is impossible to achieve better than $1/H_n$ -MMS, while in previous work, Amanatidis et al. [5] showed that $1/2H_n$ -MMS allocations can be computed. Hosseini et al. [119] showed the existence of PO and MMS or EFX allocations when agents have lexicographic preferences.

Another interesting question is to investigate the query complexity of unknown valuations. In this model, the algorithm can access the valuations by making queries to an oracle. Oh et al. [141] proved that $\Theta(\log m)$ queries suffice to define an algorithm that returns EF1 allocations. In general, it is an important research direction to explore how much information about the valuations of the agents is sufficient to design algorithms with strong fairness guarantees.

8.2. General valuations

Besides additive valuations, there are classes of valuation functions (such as submodular, fractional subadditive (XOS), and subadditive [140]) which capture more complex and combinatorial preferences. A valuation function v_i is

- *submodular* if $v_i(S) + v_i(T) \geq v_i(S \cup T) + v_i(S \cap T)$ for all S and T ;
- *subadditive* if $v_i(S) + v_i(T) \geq v_i(S \cup T)$ for all S and T ;
- *XOS* if there is a finite number of additive functions a_1, \dots, a_k such that $v_i(S) = \max_{\ell \in [k]} a_\ell(S)$.

It is known that any additive function is also submodular, any submodular function is XOS, and any XOS function is subadditive; all these implications are strict.

Table 6
A hard instance with general valuations.

	$S_1 = \{g_1, g_2\}$	$S_2 = \{g_3, g_4\}$	$S_3 = \{g_1, g_3\}$	$S_4 = \{g_2, g_4\}$	$S \neq S_1, S_2, S_3, S_4$
a_1	1	1	0	0	0
a_2	0	0	1	1	0

Table 7
Best known approximations of MMS for general valuations based on the works of Garg and Taki [103], Feige et al. [97], Ghodsi et al. [110], and Seddighin and Seddighin [146].

	Lower bound	Upper bound
Additive	$3/4 + 1/(12n)$	39/40
Submodular	1/3	3/4
XOS	1/4.6	1/2
Subadditive	$\Omega(1/(\log n \log \log n))$	1/2

The Envy-Cycle Elimination algorithm returns an EF1 allocation even when the valuation functions of the agents are as general as possible [133]. However, for MMS allocations, if the functions are not restricted, no approximation guarantee can be achieved as shown in Example 7.

Example 7.

Consider the instance given in Table 6 [62,110]. It is not hard to check that the MMS values of the two agents are $\mu_1^2 = \mu_2^2 = 1$. But, in any allocation there is at least one agent with value 0, leading to an unbounded approximation ratio with respect to MMS. □

Therefore, most of the literature has focused on more structured valuations. Barman and Krishnamurthy [38] first proved that the Round-Robin algorithm returns a 0.21-approximate MMS fair allocation for submodular valuations. The guarantee was improved to 1/3 by Ghodsi et al. [110], who complemented it with an upper bound of 3/4. For XOS valuations, Ghodsi et al. [110] proved the existence of 1/5-MMS allocations, and designed a polynomial-time algorithm to compute 1/8-MMS allocations. The existence result was improved to 1/4.6 by Seddighin and Seddighin [146]. On the negative side, Ghodsi et al. [110] showed that no algorithm can do better than 1/2-approximation for all XOS valuations. For subadditive valuations, Seddighin and Seddighin [146] proved the existence of a $\Omega(1/(\log n \log \log n))$ -MMS allocation, which exponentially improves the $\Omega(1/\log m)$ guarantee that was shown by Ghodsi et al. [110]. We summarize the best known results for approximate MMS allocations in Table 7.

Regarding EFX allocations, Plaut and Roughgarden [142] proved that 1/2-EFX allocations always exist for general valuations, and Chan et al. [75] showed that computing such allocations can be done in polynomial time. Recently, Akrami et al. [1] proved that, for instances with three agents, if one of them has an additive valuation and the others have general valuations, an exact EFX allocation always exists.

8.3. Arbitrary entitlements

So far, we assumed that all agents have equal entitlements over the goods. However, there are settings where the fairness of an allocation must be considered with respect to asymmetric entitlements; for example, in many inheritance scenarios, closer relatives have higher entitlements than distant ones. Formally, in the *weighted* or *asymmetric* setting, each agent i has an entitlement $0 < s_i < 1$ such that $\sum_{i \in N} s_i = 1$. The traditional setting with symmetric agents is a special case when $s_i = 1/n$ for every i . To capture fairness in the presence of arbitrary entitlements, one can generalize existing notions to their weighted counterparts, like *weighted* EF1 and *weighted* MMS.

Definition 18 (*Weighted EF and EF1*). An allocation A is *weighted envy-free* (WEF) if $v_i(A_i)/s_i \geq v_i(A_j)/s_j$ for every pair of agents $i, j \in N$. It is *weighted envy-free up to one item* (WEF1) if $v_i(A_i)/s_i \geq v_i(A_j \setminus \{g\})/s_j$ for every pair of agents $i, j \in N$ and some $g \in A_j$.

The concept of weighted fairness is based on ideas that were previously presented in, for example, [149,64]. Chakraborty et al. [74] proved that WEF1 allocations always exist and can be computed by a weight-dependent sequential picking algorithm that generalizes the Round-Robin algorithm.

Adapting MMS to the weighted setting is not as straightforward as WEF1. The underlying idea is to define a *fairness ratio* for each agent. Ideally, when the items are divisible, the fairest partition for agent i is such that $v_i(A_j) = s_j \cdot v_i(M)$ for all j (i.e., it is weighted proportional). However, since the items are not divisible, it is inevitable to induce some degree of unfairness and the best fairness ratio for agent i is

$$F_i = \max_{A \in \mathcal{A}_n(M)} \min_{j \in N} \left\{ \frac{v_i(A_j)}{s_j \cdot v_i(M)} \right\}.$$

Accordingly, the weighted MMS can be defined as follows.

Definition 19 (Weighted MMS [92]). An allocation A is *weighted MMS* (WMMS) if $v_i(A_i) \geq \mu_i^n = F_i \cdot s_i \cdot v_i(M)$ for every agent $i \in N$. For any $\alpha \in [0, 1]$, an allocation A is α -WMMS if $v_i(A_i) \geq \alpha \cdot \mu_i^n$ for every agent $i \in N$.

It can be verified that $F_i \leq 1$ and $\mu_i^n \leq s_i \cdot v_i(M)$. Moreover, when the agents are symmetric, WMMS coincides with MMS. Farhadi et al. [92] proved that with arbitrary entitlements, the best approximation ratio of WMMS is $1/n$, which is somewhat surprisingly guaranteed by the Round-Robin algorithm.

Novel fairness notions which highlight different perspectives of the weighted setting, such as ℓ -out-of- d share and AnyPrice share (APS), were proposed and studied by Babaioff et al. [32,31]. In particular, APS is defined as the maximum value an agent can guarantee to herself if she has a budget equal to her entitlement and the goods are adversarially priced (with prices that sum up to 1). In the unweighted setting, the APS value is at least as much as that of MMS, and sometimes it is strictly larger. Babaioff et al. [31] showed how to efficiently compute an allocation where every agent gets value no less than $3/5$ of her APS. Aziz et al. [25] considered weighted PROP1 and presented an algorithm for computing a weighted PROP1 and PO allocation for indivisible items.

8.4. Group fairness

In the model we discussed in the main part of the survey, we care about individual fairness. However, there are practical applications where it makes more sense for the allocation to be group-wisely fair. Depending on whether the groups of agents pre-exist or not, there are two lines of research on group fairness. When the groups of agents do not pre-exist, we aim to compute allocations that are fair for any group of agents. For divisible items, Berliant et al. [51] defined the notion of *group envy-freeness* (GEF) by restricting envy-freeness to hold for groups of agents with equal size; clearly, this is strictly stronger than the requirement for EF to hold for pairs of agents. Conitzer et al. [86] adapted and extended GEF for indivisible items, leading to the notion of *group-fairness* (GF), which takes into account groups of different size.

Definition 20 (GF). An allocation A is *group-fair* (GF) if, for any non-empty sets of agents $S, T \subseteq N$ and every partition $(B_i)_{i \in S}$ of $\bigcup_{j \in T} A_j$, $((|S|/|T|) \cdot v_i(B_i))_{i \in S}$ does not Pareto dominate $(v_i(A_i))_{i \in S}$.

From the above definition, we can see that the two extreme cases where $|S| = |T| = 1$ and $|S| = |T| = n$ imply the traditional requirements of EF and PO. Conitzer et al. [86] and Aziz and Rey [20] introduced different variants of “up to one” relaxations of GEF/GF and proved existential and hardness results.

When the agents are previously partitioned into several groups (e.g., each group might correspond to a family), then we only need to satisfy each given group of agents. Several models capturing scenarios along these lines have been considered in the literature. Suksompong [150] focused on a setting where each agent derives full value from all the goods allocated to the group she belongs to, and showed bounds on the best possible approximation of MMS. Kyropoulou et al. [127] considered EF1 and EFX allocations in the same setting, as well as in settings with dynamic group formation; some of these results were later improved by Manurangsi and Suksompong [136] using ideas from discrepancy theory. Segal-Halevi and Suksompong [148] focused on the case of democratic fairness, where the goal is to compute allocations that are considered fair (e.g., satisfying EF1) by a large fraction of the agents in each group.

8.5. Randomness in fair division

Best of both worlds. We mentioned in the introduction that randomization can be used as a tool to circumvent the impossibilities in the task of allocating indivisible resources. The guarantee that one can obtain this way however is that of *ex-ante* EF (EF with respect to the expected values). In fact, in the absence of any other requirement, this is rather trivial; simply allocate the goods to the agents uniformly at random. On the other hand, this sort of approach cannot help with achieving *ex-post* fairness, i.e., fairness in each resulting allocation, for any of the deterministic notions of fairness mentioned in this survey. Therefore, it seems natural to ask for a randomized allocation algorithm that also has good *ex-post* fairness guarantees.

As we have seen, Round-Robin algorithms return an EF1 allocation, but the agents have an advantage over those who are ordered after them. To eliminate this unfair advantage, we can consider a randomized version where the ordering of the agents is picked at random. Freeman et al. [99] showed that this randomized Round-Robin algorithm guarantees *ex-ante* PROP but not *ex-ante* EF.

The Probabilistic-Serial (PS) algorithm of Bogomolnaia and Moulin [55] is a randomized algorithm for allocating indivisible items in an *ex-ante* EF manner. Agents eat their most preferred good at a uniform rate and, once consumed, move on to the next available good. The algorithm terminates when all goods have been consumed, and the probability share of an

agent for a good is the fraction of the good eaten by the agent. Freeman et al. [99] presented a variant of this procedure, the recursive probabilistic serial algorithm, that maintains the ex-ante EF property, and moreover every realized allocation is ex-post EF1. Later, Aziz [16] showed that the random allocation generated by the original probabilistic serial algorithm is equivalent to lottery over a set of EF1 allocations. Aziz and Brandl [18] presented an eating algorithm that is suitable for any type of feasibility constraint and allocation problem with ordinal preferences.

Finally, in a somewhat different direction, Caragiannis et al. [72] studied *interim* envy-freeness, a notion which lies between ex-ante and ex-post envy-freeness. They showed positive and impossibility results for many settings, including one-sided matching, where each agent can be given only one good.

Stochastic settings. Although it is easy to construct instances where an envy-free allocation does not exist, there is a line of research showing that when the values are randomly drawn from some probability distributions (instead of being chosen adversarially), an envy-free allocation is likely to exist as long as the number of items is sufficiently large compared to the number of agents.

In particular, Dickerson et al. [88] showed that when the values of the agents are independently drawn from an identical distribution, the social welfare maximizing allocation (that is, allocating each good to the agent who has the highest value for it) is envy-free with high probability when $m = \Omega(n \log n)$. Manurangsi and Suksompong [134,135] further improved the lower bound on the existence of envy-free allocations, and showed that the Round-Robin algorithm produces an envy-free allocation for a slightly lower m . Recently, Bai and Gözl [33] extended this result to the case of asymmetric distributions.

In a different direction, Bai et al. [34] showed that given a worst-case instance that does not admit any envy-free allocation, randomly perturbing the values of each agent leads to the existence of envy-free allocations with high probability.

8.6. Online fair division

Most models we have discussed are static, as all goods, agents, and their valuation functions do not change over time. Online fair division considers settings where the agents or the goods arrive in an online manner. We briefly mention some of the most related works; we refer the readers to Aleksandrov et al. [4], Aleksandrov and Walsh [3] for a more detailed discussion. In the most common model, T items arrive online; when an item arrives, the values of the agents for it are realized and based on these we need to decide immediately and irrevocably how to allocate it, typically without knowing future events (including how many items will arrive later and the values of the agents for these items).

Let Envy_T be the cumulative envy until time T , i.e., the maximum difference between any agent's value for goods allocated to another agent and to herself in the first T rounds. Then, it is desired that the envy can vanish as time goes on, that is, $\text{Envy}_T/T \rightarrow 0$ when T goes to infinity. Benade et al. [49] proved that by allocating the new good to an agent chosen uniformly at random, $\text{Envy}_T = \tilde{O}(\sqrt{T/n})$, and thus we can indeed eliminate the envy over time. Benade et al. [49] also showed that there exists a deterministic polynomial-time algorithm with the same envy bound as the random allocation algorithm, and that this bound is asymptotically optimal against an adaptive adversary (who can see the random bits of the algorithm through the first T goods and decide the value v_{iT}), meaning that the allocation is far from being EF1. To circumvent this obstacle, He et al. [117] relaxed the requirement of "irrevocable decision" by allowing the algorithm to reallocate previously allocated items. Clearly, if we reallocate all goods at each time (i.e., with $\Theta(T^2)$ reallocations), we can surely achieve EF1, but this makes the online setting meaningless. He et al. [117] showed that with two agents we actually only need $\Theta(T)$ reallocations; for more than two agents, $O(T^{3/2})$ reallocations are sufficient and $\Omega(T)$ are necessary.

The stochastic setting of online fair allocation is also well-studied. Actually, the algorithms by Dickerson et al. [88], Kurokawa et al. [125] and Bai and Gözl [33] that we have introduced in Section 8.5 for allocating goods with random values do not need to access the order of goods and also work for the online setting. Thus, we can achieve envy-freeness with high probability by allocating each arriving good to the agent with the highest value. Zeng and Psomas [158] further showed that, with some modification, the returned allocation is either EF1, or envy-free with high probability. Zeng and Psomas [158] also investigated the trade-off between fairness and efficiency when the adversary has different levels of power.

It remains an open question whether there exist competitive online algorithms for the computation of (approximately) MMS or Prop1/PropX allocations. For the case of identical valuation functions, approximate MMS allocations correspond to maximizing the minimum load on the job scheduling problem, for which optimal competitive ratios have been proved by Azar and Epstein [15] and Tan and Wu [153]. The alternative model that considers a fixed set of resources and agents who arrive or depart over time has not been considered for indivisible resources, partially because it is very challenging to achieve positive results [122].

8.7. Subsidies

As we saw in Section 7.3, even in a game-theoretic setting no monetary transfers are allowed in fair division problems. Indeed, arbitrary payments would significantly alter the flavor of these problems and often go against their motivation. A recent line of work, however, considers the question of whether it is possible to pay the agents just a small amount of money (subsidy) on top of a given allocation in order to make it envy-free (when the subsidies are also taken into consideration). Allocations for which this can be done are called *envy-freeable*. Formally, a pair (A, p) consisting of an allocation $A = (A_1, \dots, A_n)$ and a payment profile $p = (p_1, \dots, p_n)$, where $p_i \geq 0$ for every i , is envy-free if $v_i(A_i) + p_i \geq v_i(A_j) + p_j$. An allocation A is *envy-freeable* if there are payments $p = (p_1, \dots, p_n)$ such that (A, p) is envy-free.

Halpern and Shah [114] first noted that not all allocations are envy-freeable by considering an example with two agents that have values 150 and 100 for a single good. If the good is allocated to the second agent, then we have to compensate the first agent at least 150, however, this would make the second agent envious. Halpern and Shah [114] then gave a characterization of all envy-freeable allocations. Specifically, they showed that allocation A is envy-freeable if and only if A maximizes the utilitarian welfare across all reassignments of its bundles to agents.

A natural question to ask in this setting is to quantify the minimum total amount of subsidy $\sum_i p_i$ required to find an envy-free allocation. Towards answering this question, Halpern and Shah [114] showed that, if a fixed envy-freeable allocation is given, the minimum total subsidy can be computed in polynomial time. However, finding the envy-freeable allocation with overall minimum total subsidy is an NP-hard problem; this follows since deciding the existence of an envy-free allocation without any subsidy is, as we already discussed, NP-hard [60]. Caragiannis and Ioannidis [67] provided (additive) approximation guarantees and hardness results for computing an envy-freeable allocation that minimizes the total amount of subsidies.

To bound the minimum subsidy required in the worst case over allocations, Halpern and Shah [114] proved that any envy-freeable allocation requires no more than $(n - 1)m \cdot v^*$ total subsidy, where v^* is the maximum value of an agent for a good. If we are able to choose the allocations, the $(n - 1)$ factor cannot be removed. Consider an instance with a single good that all agents value as v^* ; then, a total payment of at least $(n - 1) \cdot v^*$ is required to eliminate the envy of all agents, in any allocation. Halpern and Shah [114] conjectured that it is always possible to find an allocation with no more than $(n - 1) \cdot v^*$ total subsidy, which was later indeed proved by Brustle et al. [65]. In addition, Brustle et al. proved that it suffices to subsidize each agent by at most v^* .

8.8. Mixtures of indivisible and divisible items

There has also been recent work on models that involve both indivisible and divisible goods; note that limited subsidy can be thought of as a divisible good. Bei et al. [46] proposed a new fairness notion called *envy-freeness for mixed goods (EFM)*, which is a generalization of both EF and EF1 to the mixed goods setting. The key idea behind EFM is that (1) an agent i 's envy towards another agent j vanishes, if an indivisible good is removed from consideration when j only gets indivisible goods, or (2) agent i does not envy agent j when j gets some divisible goods. Bei et al. [46] proved that an EFM allocation always exists for any number of agents with additive valuations. Bei et al. [47] examined the same setting and explored approximations of MMS. Bhaskar et al. [54] showed that for a mixed resources model consisting of indivisible items and a divisible, undesirable heterogeneous resource, an EFM allocation always exists.

8.9. Chores and mixed Manna

Beyond discrete fair division of goods, there is a significant line of work that considers similar questions when items can be seen as *chores* (which are negatively valued by the agents), or *mixed manna* (a mixture of both goods and chores). Here, we give a brief overview of these settings.

The definitions of EF, PROP and MMS remain the same as in Definitions 1, 2, 5, and 9. As in the case of goods, the existence of EF or PROP allocations is rarely guaranteed. For additive valuations, Aziz et al. [22] proved that the Round-Robin algorithm returns a 2-MMS allocation. This result was later improved to 4/3 and 11/9 by Barman and Krishnamurthy [38] and Huang and Lu [120], respectively. On the negative side, Feige et al. [97] proved that no algorithm can ensure an approximation factor better than 44/43. If the algorithm has access only to the ordinal preferences of the agents, the best approximation of MMS that can be guaranteed is between 7/5 and 5/3 [29]. When the valuations are submodular, Li et al. [131] proved that no algorithm guarantees a better than n -MMS allocation. Aziz et al. [24] considered the case when the agents have asymmetric entitlements and extended the notion of weighted MMS to chores, but the tight approximation ratio is still unknown.

Relaxations of EF and PROP need to be changed a bit to capture the fact that the agents have negative values for chores. Essentially, instead of removing goods from bundles of other agents to eliminate the envy of an agent, we need to remove chores from the bundle of the agent herself.

Definition 21 (EF1 and EFX for chores). For chores, an allocation A is α -EF1 if $v_i(A_i \setminus \{g\}) \geq \alpha \cdot v_i(A_j)$ for any pair of agents i, j and some $g \in A_i$; A is α -EFX if the inequality holds for any $g \in A_i$.

For additive valuations, an EF1 allocation can be easily computed by the Round-Robin algorithm. However, as noted by Bhaskar et al. [54], the allocation returned by the Envy-Cycle Elimination algorithm may not be EF1 if the cycles are resolved arbitrarily. Instead, the top-trading technique (in which each agent only points to the agent she envies the most in the envy graph) must be used to preserve EF1. In contrast to the case of goods where EF1 and PO can be satisfied simultaneously, the compatibility of EF1 and PO is still unknown for chores. Ebadian et al. [90], Garg et al. [108] and Chaudhury et al. [81] proved that, for bi-valued instances, EF1+PO allocations always exist and can be found efficiently; this is the only known result so far. The existence of EFX allocations for chores also remains open, with the only positive result being the computation of $O(n^2)$ -EFX allocations due to Zhou and Wu [159].

Definition 22. For chores, an allocation A is Prop1 if $v_i(A_i \setminus \{g\}) \geq \alpha \cdot v_i(M)/n$ for any pair of agents i, j and some $g \in A_i$; A is PropX if the inequality holds for any $g \in A_i$.

The existence of Prop1 allocations is straightforward as EF1 implies Prop1. But the good news for Prop1 is that it can always be satisfied together with PO, even if the set of items is a mixture of goods and chores [25]. In contrast to goods, where PropX allocations may not exist, for chores, a PropX allocations always exist and can be computed efficiently, even when the agents have asymmetric entitlements [130]. However, it is still unknown whether PropX and PO are compatible or not. The relationships among various fairness notions for chores are discussed by Sun et al. [152].

The more general case of mixtures of goods and chores has recently also been studied [57,58,25,27]. This model is particularly interesting as it includes the setting with non-monotone valuations. Aziz et al. [27] proved that a double Round-Robin algorithm computes an EF1 allocation for any number of agents, and a generalized adjusted winner algorithm finds an EF1+PO allocation for two agents. Recently, Aziz et al. [25] and Kulkarni et al. [123] designed algorithms for computing Prop1+PO or approximately MMS+PO allocations, respectively. It is an intriguing future research direction to study the fair allocation problem under other non-monotonic valuations.

Finally, the fair allocation of divisible chores is also studied in the literature, such as the computation of envy-free allocations [87] and competitive equilibria [59,77]. As such settings are out of the scope of our survey, we refer the reader to these works and references therein.

8.10. Other settings

There are several other variants that we do not discuss in the current survey. Depending on the application, some allocations may not be feasible due to various restrictions, such as connectivity, cardinality, separation, or budget constraints. Such models have recently attracted the attention of the community. Rather than referring to specific works, we point the reader to the survey of [151] which discusses this part of the literature in detail. For another example, externalities have been studied in fair division by Seddighin et al. [147] and Aziz et al. [26], where the agents not only have value for what they get but also for what others get. Yet another example is related to public resource allocation where the items can be shared by the agents [85,91,106].

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgements

This work is partially supported by the ERC Advanced Grant 788893 AMDROMA “Algorithmic and Mechanism Design Research in Online Markets”, the MIUR PRIN project ALGADIMAR “Algorithms, Games, Digital Markets”, the NWO Veni project No. VI.Veni.192.153, NSFC No. 62102333, HKSAR RGC No. PolyU 25211321, PolyU Start-up No. P0034420, FDCT (File no. 0014/2022/AFJ, 0085/2022/A, 0143/2020/A3, SKL-IOTSC-2021-2023).

References

- [1] Hannaneh Akrami, Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta, EFX allocations: simplifications and improvements, CoRR, arXiv: 2205.07638 [abs], 2022.
- [2] Hannaneh Akrami, Bhaskar Ray Chaudhury, Martin Hoefer, Kurt Mehlhorn, Marco Schmalhofer, Golnoosh Shahkarami, Giovanna Varricchio, Quentin Vermande, Ernest van Wijnland, Maximizing Nash social welfare in 2-value instances, in: Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI), 2022, pp. 4760–4767.
- [3] Martin Aleksandrov, Toby Walsh, Online fair division: a survey, in: Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI), 2020, pp. 13557–13562.
- [4] Martin Aleksandrov, Haris Aziz, Serge Gaspers, Toby Walsh, Online fair division: analysing a food bank problem, in: Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI), 2015, pp. 2540–2546.
- [5] Georgios Amanatidis, Georgios Birmpas, Evangelos Markakis, On truthful mechanisms for maximin share allocations, in: Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI), 2016, pp. 31–37.
- [6] Georgios Amanatidis, Georgios Birmpas, George Christodoulou, Evangelos Markakis, Truthful allocation mechanisms without payments: characterization and implications on fairness, in: Proceedings of the 2017 ACM Conference on Economics and Computation (EC), 2017, pp. 545–562.
- [7] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, Amin Saberi, Approximation algorithms for computing maximin share allocations, ACM Trans. Algorithms 13 (4) (2017) 52:1–52:28.
- [8] Georgios Amanatidis, Georgios Birmpas, Vangelis Markakis, Comparing approximate relaxations of envy-freeness, in: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI), 2018, pp. 42–48.

- [9] Georgios Amanatidis, Evangelos Markakis, Apostolos Ntokos, Multiple birds with one stone: beating $1/2$ for EFX and GMMS via envy cycle elimination, *Theor. Comput. Sci.* 841 (2020) 94–109.
- [10] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, Alexandros A. Voudouris, Maximum Nash welfare and other stories about EFX, *Theor. Comput. Sci.* 863 (2021) 69–85.
- [11] Georgios Amanatidis, Georgios Birmpas, Federico Fusco, Philip Lazos, Stefano Leonardi, Rebecca Reiffenhäuser, Allocating indivisible goods to strategic agents: pure Nash equilibria and fairness, in: *Proceedings of the 17th International Conference on Web and Internet Economics (WINE)*, 2021, pp. 149–166.
- [12] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros A. Voudouris, Fair division of indivisible goods: a survey, in: *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 2022, pp. 5385–5393.
- [13] Georgios Amanatidis, Georgios Birmpas, Philip Lazos, Stefano Leonardi, Rebecca Reiffenhäuser, Round-Robin beyond additive agents: existence and fairness of approximate equilibria, *CoRR*, arXiv:2301.13652 [abs], 2023.
- [14] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, Alexandros A. Voudouris, Distortion in social choice problems: the first 15 years and beyond, in: *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 2021, pp. 4294–4301.
- [15] Yossi Azar, Leah Epstein, On-line machine covering, in: *Proceedings of the 5th Annual European Symposium on Algorithms (ESA)*, 1997, pp. 23–36.
- [16] Haris Aziz, Simultaneously achieving ex-ante and ex-post fairness, in: *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*, 2020, pp. 341–355.
- [17] Haris Aziz, Developments in multi-agent fair allocation, in: *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020, pp. 13563–13568.
- [18] Haris Aziz, Florian Brandl, The vigilant eating rule: a general approach for probabilistic economic design with constraints, *CoRR*, arXiv:2008.08991 [abs], 2020.
- [19] Haris Aziz, Simon Mackenzie, A discrete and bounded envy-free cake cutting protocol for any number of agents, in: *Proceedings of the 57th Annual Symposium on Foundations of Computer Science (FOCS)*, 2016, pp. 416–427.
- [20] Haris Aziz, Simon Rey, Almost group envy-free allocation of indivisible goods and chores, in: *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 2020, pp. 39–45.
- [21] Haris Aziz, Sylvain Bouveret, Jérôme Lang, Simon Mackenzie, Complexity of manipulating sequential allocation, in: *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017, pp. 328–334.
- [22] Haris Aziz, Gerhard Raucherer, Guido Schryen, Toby Walsh, Algorithms for max-min share fair allocation of indivisible chores, in: *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017, pp. 335–341.
- [23] Haris Aziz, Sylvain Bouveret, Ioannis Caragiannis, Ira Giagkousi, Jérôme Lang, Knowledge, fairness, and social constraints, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018, pp. 4638–4645.
- [24] Haris Aziz, Hau Chan, Bo Li, Weighted maxmin fair share allocation of indivisible chores, in: *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019, pp. 46–52.
- [25] Haris Aziz, Hervé Moulin, Fedor Sandomirskiy, A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation, *Oper. Res. Lett.* 48 (5) (2020) 573–578.
- [26] Haris Aziz, Warut Suksompong, Zhaohong Sun, Toby Walsh, Fairness concepts for indivisible items with externalities, *CoRR*, arXiv:2110.09066 [abs], 2021.
- [27] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, Toby Walsh, Fair allocation of indivisible goods and chores, *Auton. Agents Multi-Agent Syst.* 36 (1) (2022) 3.
- [28] Haris Aziz, Bo Li Hervé Moulin, Xiaowei Wu, Algorithmic fair allocation of indivisible items: a survey and new questions, *SIGecom Exch.* 20 (2022).
- [29] Haris Aziz, Bo Li, Xiaowei Wu, Approximate and strategyproof maximin share allocation of chores with ordinal preferences, *Math. Program.* (2022) 1–27.
- [30] Moshe Babaioff Tomer Ezra, Uriel Feige, Fair and truthful mechanisms for dichotomous valuations, in: *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 2021, pp. 5119–5126.
- [31] Moshe Babaioff Tomer Ezra, Uriel Feige, Fair-share allocations for agents with arbitrary entitlements, in: *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 2021, p. 127.
- [32] Moshe Babaioff, Noam Nisan, Inbal Talgam-Cohen, Competitive equilibrium with indivisible goods and generic budgets, *Math. Oper. Res.* 46 (1) (2021) 382–403.
- [33] Yushi Bai, Paul Gözl, Envy-free and Pareto-optimal allocations for agents with asymmetric random valuations, in: *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 2022, pp. 53–59.
- [34] Yushi Bai, Uriel Feige, Paul Gözl, Ariel D. Procaccia, Fair allocations for smoothed utilities, in: *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, 2022, pp. 436–465.
- [35] Artem Baklanov, Pranav Garimidi, Vasilis Gkatzelis, Daniel Schoepflin, Achieving proportionality up to the maximin item with indivisible goods, in: *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*, 2021, pp. 5143–5150.
- [36] Artem Baklanov, Pranav Garimidi, Vasilis Gkatzelis, Daniel Schoepflin, PROPm allocations of indivisible goods to multiple agents, in: *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, 2021, pp. 24–30.
- [37] Nikhil Bansal, Maxim Sviridenko, The santa claus problem, in: *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, 2006, pp. 31–40.
- [38] Siddharth Barman, Sanath K. Krishnamurthy, Approximation algorithms for maximin fair division, *ACM Trans. Econ. Comput.* 8 (1) (2020) 5:1–5:28.
- [39] Siddharth Barman, Sanath Kumar Krishnamurthy, On the proximity of markets with integral equilibria, in: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019, pp. 1748–1755.
- [40] Siddharth Barman, Paritosh Verma, Existence and computation of maximin fair allocations under matroid-rank valuations, in: *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, 2021, pp. 169–177.
- [41] Siddharth Barman, Arpita Biswas, Sanath Kumar Krishna Murthy, Yadati Narahari, Groupwise maximin fair allocation of indivisible goods, in: *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 2018, pp. 917–924.
- [42] Siddharth Barman, Sanath K. Krishnamurthy, Rohit Vaish, Finding fair and efficient allocations, in: *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, 2018, pp. 557–574.
- [43] Siddharth Barman, Umang Bhaskar, Anand Krishna, Ranjani G. Sundaram, Tight approximation algorithms for p-mean welfare under subadditive valuations, in: *ESA*, in: *LIPICs*, vol. 173, Schloss Dagstuhl - Leibniz-Zentrum Für Informatik, 2020, pp. 11:1–11:17.
- [44] Siddharth Barman, Umang Bhaskar, Nisarg Shah, Optimal bounds on the price of fairness for indivisible goods, in: *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*, 2020, pp. 356–369.
- [45] Siddharth Barman, Anand Krishna, Pooja Kulkarni, Shivika Narang, Sublinear approximation algorithm for Nash social welfare with XOS valuations, *CoRR*, arXiv:2110.00767 [abs], 2021.
- [46] Xiaohui Bei, Zihao Li, Jinyan Liu, Shengxin Liu, Xinhang Lu, Fair division of mixed divisible and indivisible goods, *Artif. Intell.* 293 (2021) 103436.
- [47] Xiaohui Bei, Shengxin Liu, Xinhang Lu, Hongao Wang, Maximin fairness with mixed divisible and indivisible goods, *Auton. Agents Multi-Agent Syst.* 35 (2) (2021) 34.

- [48] Xiaohui Bei, Xinhang Lu, Pasin Manurangsi, Warut Suksompong, The price of fairness for indivisible goods, *Theory Comput. Syst.* 65 (7) (2021) 1069–1093.
- [49] Gerdus Benade, Aleksandr M. Kazachkov, Ariel D. Procaccia, Christos-Alexandros Psomas, How to make envy vanish over time, in: *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, 2018, pp. 593–610.
- [50] Ben Berger, Avi Cohen, Michal Feldman, Amos Fiat, Almost full EFX exists for four agents, in: *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 2022, pp. 4826–4833.
- [51] Marcus Berliant, William Thomson, Karl Dunz, On the fair division of a heterogeneous commodity, *J. Math. Econ.* 21 (3) (1992) 201–216.
- [52] Dimitris Bertsimas, Vivek F. Farias, Nikolaos Trichakis, The price of fairness, *Oper. Res.* 59 (1) (2011) 17–31.
- [53] Ivona Bezáková, Varsha Dani, Allocating indivisible goods, *SIGecom Exch.* 5 (3) (2005) 11–18, <https://doi.org/10.1145/1120680.1120683>.
- [54] Umang Bhaskar, A.R. Sricharan, Rohit Vaish, On approximate envy-freeness for indivisible chores and mixed resources, in: *Proceedings of the 2021 International Conference on Approximation, Randomization, and Combinatorial Optimization (APPROX/RANDOM)*, 2021, pp. 1:1–1:23.
- [55] Anna Bogomolnaia, Hervé Moulin, A new solution to the random assignment problem, *J. Econ. Theory* 100 (2) (2001) 295–328.
- [56] Anna Bogomolnaia, Hervé Moulin, Random matching under dichotomous preferences, *Econometrica* 72 (1) (2004) 257–279.
- [57] Anna Bogomolnaia, Hervé Moulin, Fedor Sandomirskiy, Elena Yanovskaya, Competitive division of a mixed Manna, *Econometrica* 85 (6) (2017) 1847–1871.
- [58] Anna Bogomolnaia, Hervé Moulin, Fedor Sandomirskiy, Elena Yanovskaya, Dividing goods and bads under additive utilities, *Soc. Choice Welf.* 52 (2019) 395–417.
- [59] Shant Boodaghians, Bhaskar Ray Chaudhury, Ruta Mehta, Polynomial time algorithms to find an approximate competitive equilibrium for chores, in: *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2022, pp. 2285–2302.
- [60] Sylvain Bouveret, Jérôme Lang, Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity, *J. Artif. Intell. Res.* 32 (2008) 525–564.
- [61] Sylvain Bouveret, Jérôme Lang, Manipulating picking sequences, in: *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*, 2014, pp. 141–146.
- [62] Sylvain Bouveret, Michel Lemaître, Characterizing conflicts in fair division of indivisible goods using a scale of criteria, *Auton. Agents Multi-Agent Syst.* 30 (2) (2016) 259–290.
- [63] Sylvain Bouveret, Yann Chevaleyre, Nicolas Maudet, Fair allocation of indivisible goods, in: Ariel D. Procaccia, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang (Eds.), *Handbook of Computational Social Choice*, Cambridge University Press, 2016, pp. 284–310, chapter 12.
- [64] Steven J. Brams, Alan D. Taylor, *Fair Division: From Cake-Cutting to Dispute Resolution*, Cambridge University Press, 1996.
- [65] Johannes Brustle, Jack Dippel, Vishnu V. Narayan, Mashbat Suzuki, Adrian Vetta, One dollar each eliminates envy, in: *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 2020, pp. 23–39.
- [66] Erik Budish, The combinatorial assignment problem: approximate competitive equilibrium from equal incomes, *J. Polit. Econ.* 119 (6) (2011) 1061–1103.
- [67] Ioannis Caragiannis, Stavros Ioannidis, Computing envy-freeable allocations with limited subsidies, in: *Proceedings of the 17th International Conference on Web and Internet Economics (WINE)*, 2021, pp. 522–539.
- [68] Ioannis Caragiannis, Christoc Kaklamani, Panagiotis Kanellopoulos, Maria Kyropoulou, On low-envy truthful allocations, in: *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT)*, 2009, pp. 111–119.
- [69] Ioannis Caragiannis, Christos Kaklamani, Panagiotis Kanellopoulos, Maria Kyropoulou, The efficiency of fair division, *Theory Comput. Syst.* 50 (4) (2012) 589–610.
- [70] Ioannis Caragiannis, Nick Gravin, Xin Huang, 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, pp. 527–545.
- [71] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, Junxing Wang, The unreasonable fairness of maximum Nash welfare, *ACM Trans. Econ. Comput.* 7 (3) (2019) 12:1–12:32.
- [72] Ioannis Caragiannis, Panagiotis Kanellopoulos, Maria Kyropoulou, On interim envy-free allocation lotteries, in: *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 2021, pp. 264–284.
- [73] Ioannis Caragiannis, Jugal Garg, Nidhi Rathi, Eklaty Sharma, Giovanna Varricchio, Existence and computation of epistemic efx, *arXiv preprint, arXiv: 2206.01710*, 2022.
- [74] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, Yair Zick, Weighted envy-freeness in indivisible item allocation, *ACM Trans. Econ. Comput.* 9 (3) (2021) 18:1–18:39.
- [75] Hau Chan, Jing Chen, Bo Li, Xiaowei Wu, Maximin-aware allocations of indivisible goods, in: *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019, pp. 137–143.
- [76] Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, EFX exists for three agents, in: *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 2020, pp. 1–19.
- [77] Bhaskar Ray Chaudhury, Jugal Garg, Peter McGlaughlin, Ruta Mehta, Competitive allocation of a mixed Manna, in: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2021, pp. 1405–1424.
- [78] Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta, Pranabendu Misra, Improving EFX guarantees through rainbow cycle number, in: *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 2021, pp. 310–311.
- [79] Bhaskar Ray Chaudhury, Jugal Garg, Ruta Mehta, Fair and efficient allocations under subadditive valuations, in: *AAAI, AAAI Press*, 2021, pp. 5269–5276.
- [80] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, Alkmini Sgouritsa, A little charity guarantees almost envy-freeness, *SIAM J. Comput.* 50 (4) (2021) 1336–1358.
- [81] Bhaskar Ray Chaudhury, Jugal Garg, Peter McGlaughlin, Ruta Mehta, Competitive equilibrium with chores: combinatorial algorithm and hardness, *CoRR, arXiv:2205.11363 [abs]*, 2022.
- [82] Yiling Chen, Nisarg Shah, Ignorance is often bliss: Envy with incomplete information, Technical report, Working paper, Harvard University, 2017.
- [83] Richard Cole, Vasilis Gkatzelis, Approximating the Nash social welfare with indivisible items, *SIAM J. Comput.* 47 (3) (2018) 1211–1236.
- [84] Richard Cole, Nikhil R. Devanur, Vasilis Gkatzelis, Kamal Jain, Tung Mai, Vijay V. Vazirani, Sadra Yazdanbod, Convex program duality, Fisher markets, and Nash social welfare, in: *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, 2017, pp. 459–460.
- [85] Vincent Conitzer, Rupert Freeman, Nisarg Shah, Fair public decision making, in: *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, 2017, pp. 629–646.
- [86] Vincent Conitzer, Rupert Freeman, Nisarg Shah, Jennifer Wortman Vaughan, Group fairness for the allocation of indivisible goods, in: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019, pp. 1853–1860.
- [87] Sina Dehghani, Alireza Farhadi, Mohammad Taghi Hajiaghayi, Hadi Yami, Envy-free chore division for an arbitrary number of agents, in: *Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2018.
- [88] John P. Dickerson, Jonathan R. Goldman, Jeremy Karp, Ariel D. Procaccia, Tuomas Sandholm, The computational rise and fall of fairness, in: *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 2014, pp. 1405–1411.
- [89] L.E. Dubins, E.H. Spanier, How to cut a cake fairly, *Am. Math. Mon.* 68 (1P1) (1961) 1–17.

- [90] Sorush Ebadian, Dominik Peters, Nisarg Shah, How to fairly allocate easy and difficult chores, in: Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2022, pp. 372–380.
- [91] Brandon Fain, Kamesh Munagala, Nisarg Shah, Fair allocation of indivisible public goods, in: Proceedings of the 2018 ACM Conference on Economics and Computation (EC), 2018, pp. 575–592.
- [92] Alireza Farhadi, Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Sébastien Lahaie, David M. Pennock, Masoud Seddighin, Saeed Seddighin, Hadi Yami, Fair allocation of indivisible goods to asymmetric agents, *J. Artif. Intell. Res.* 64 (2019) 1–20.
- [93] Alireza Farhadi, Mohammad Taghi Hajiaghayi, Mohamad Latifian, Masoud Seddighin, Hadi Yami, Almost envy-freeness, envy-rank, and Nash social welfare matchings, in: Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI), 2021, pp. 5355–5362.
- [94] Uriel Feige, On allocations that maximize fairness, in: Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms, 2008, pp. 287–293.
- [95] Uriel Feige, Maximin fair allocations with two item values. Unpublished manuscript, <https://www.wisdom.weizmann.ac.il/~feige/mypapers/MMSab.pdf>, 2022.
- [96] Uriel Feige, Alexey Norikin, Improved maximin fair allocation of indivisible items to three agents, *CoRR*, arXiv:2205.05363 [abs], 2022.
- [97] Uriel Feige, Ariel Sapir, Lali Tauber, A tight negative example for MMS fair allocations, in: Proceedings of the 17th International Conference on Web and Internet Economics (WINE), 2021, pp. 355–372.
- [98] Rupert Freeman, Sujoy Sikdar, Rohit Vaish, Lirong Xia, Equitable allocations of indivisible goods, in: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), 2019, pp. 280–286.
- [99] Rupert Freeman, Nisarg Shah, Rohit Vaish, Best of both worlds: ex-ante and ex-post fairness in resource allocation, in: Proceedings of the 21st ACM Conference on Economics and Computation (EC), 2020, pp. 21–22.
- [100] George Gamow, Marvin Stern, *Puzzle-Math*, Viking Press, 1958.
- [101] Jugal Garg, Aniket Murhekar, Computing fair and efficient allocations with few utility values, in: Proceedings of the 14th International Symposium on Algorithmic Game Theory (SAGT), 2021, pp. 345–359.
- [102] Jugal Garg, Aniket Murhekar, Computing Pareto-optimal and almost envy-free allocations of indivisible goods, arXiv preprint, arXiv:2204.14229, 2022.
- [103] Jugal Garg, Setareh Taki, An improved approximation algorithm for maximin shares, *Artif. Intell.* 300 (2021) 103547.
- [104] Jugal Garg, Martin Hoefer, Kurt Mehlhorn, Satiation in Fisher markets and approximation of Nash social welfare, arXiv preprint, arXiv:1707.04428, 2017.
- [105] Jugal Garg, Pooja Kulkarni, Rucha Kulkarni, Approximating Nash social welfare under submodular valuations through (un)matchings, in: *SODA, SIAM*, 2020, pp. 2673–2687.
- [106] Jugal Garg, Pooja Kulkarni, Aniket Murhekar, On fair and efficient allocations of indivisible public goods, in: Proceedings of the 41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS), 2021, pp. 22:1–22:19.
- [107] Jugal Garg, Edin Husic, Wenzheng Li, László A. Végh, Jan Vondrák, Approximating Nash social welfare by matching and local search, *CoRR* (in press in *STOC 2023*), arXiv:2211.03883 [abs], 2022, <https://doi.org/10.1145/3564246.3585255>.
- [108] Jugal Garg, Aniket Murhekar, John Qin, Fair and efficient allocations of chores under bivalued preferences, in: Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI), 2022, pp. 5043–5050.
- [109] Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, Hadi Yami, Fair allocation of indivisible goods: improvement, *Math. Oper. Res.* 46 (3) (2021) 1038–1053.
- [110] Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, Hadi Yami, Fair allocation of indivisible goods: beyond additive valuations, *Artif. Intell.* 303 (2022) 103633.
- [111] Paul W. Goldberg, Kasper Høgh, Alexandros Hollender, The frontier of intractability for EFX with two agents, *CoRR*, arXiv:2301.10354 [abs], 2023.
- [112] Laurent Gourvès, Jérôme Monnot, On maximin share allocations in matroids, *Theor. Comput. Sci.* 754 (2019) 50–64.
- [113] Laurent Gourvès, Jérôme Monnot, Lydia Tiliane, Near fairness in matroids, in: Proceedings of the 21st European Conference on Artificial Intelligence (ECAI), 2014, pp. 393–398.
- [114] Daniel Halpern, Nisarg Shah, Fair division with subsidy, in: Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT), 2019, pp. 374–389.
- [115] Daniel Halpern, Nisarg Shah, Fair and efficient resource allocation with partial information, in: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), 2021, pp. 224–230.
- [116] Daniel Halpern, Ariel D. Procaccia, Alexandros Psomas, Nisarg Shah, Fair division with binary valuations: one rule to rule them all, in: Proceedings of the 16th International Conference on Web and Internet Economics (WINE), 2020, pp. 370–383.
- [117] Jiafan He, Ariel D. Procaccia, Alexandros Psomas, David Zeng, Achieving a fairer future by changing the past, in: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), 2019, pp. 343–349.
- [118] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, Hejun Wang, Lirong Xia, Fair division through information withholding, in: Proceedings of the AAAI Conference on Artificial Intelligence, 2020, pp. 2014–2021.
- [119] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, Lirong Xia, Fair and efficient allocations under lexicographic preferences, in: Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI), 2021, pp. 5472–5480.
- [120] Xin Huang, Pinyan Lu, An algorithmic framework for approximating maximin share allocation of chores, in: Proceedings of the 22nd ACM Conference on Economics and Computation (EC), 2021, pp. 630–631.
- [121] Aanund Hylland, Richard Zeckhauser, The efficient allocation of individuals to positions, *J. Polit. Econ.* 87 (2) (1979) 293–314.
- [122] Ian A. Kash, Ariel D. Procaccia, Nisarg Shah, No agent left behind: dynamic fair division of multiple resources, *J. Artif. Intell. Res.* 51 (2014) 579–603.
- [123] Rucha Kulkarni, Ruta Mehta, Setareh Taki, Indivisible mixed Manna: on the computability of MMS+PO allocations, in: Proceedings of the 22nd ACM Conference on Economics and Computation (EC), 2021, pp. 683–684.
- [124] David Kurokawa, Fair Division in Game Theoretic Settings, PhD thesis, Carnegie Mellon University, 2017.
- [125] David Kurokawa, Ariel D. Procaccia, Junxing Wang, When can the maximin share guarantee be guaranteed?, in: Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI), 2016, pp. 523–529.
- [126] David Kurokawa, Ariel D. Procaccia, Junxing Wang, Fair enough: guaranteeing approximate maximin shares, *J. ACM* 65 (2) (2018) 8:1–8:27.
- [127] Maria Kyropoulou, Warut Suksompong, Alexandros A. Voudouris, Almost envy-freeness in group resource allocation, *Theor. Comput. Sci.* 841 (2020) 110–123.
- [128] Jérôme Lang, Jörg Rothe, Fair division of indivisible goods, in: Jörg Rothe (Ed.), *Economics and Computation*, Springer, 2016, pp. 493–550, chapter 8.
- [129] Jan Karel Lenstra, David B. Shmoys, Éva Tardos, Approximation algorithms for scheduling unrelated parallel machines, *Math. Program.* 46 (1990) 259–271.
- [130] Bo Li, Yingkai Li, Xiaowei Wu, Almost (weighted) proportional allocations for indivisible chores, in: Proceedings of the ACM Web Conference 2022 (WWW), 2022, pp. 122–131.
- [131] Bo Li, Fangxiao Wang, Yu Zhou, Fair allocation of indivisible chores: beyond additive valuations, *CoRR*, arXiv:2205.10520 [abs], 2022.
- [132] Claudia Lindner, Jörg Rothe, Fair division of divisible goods, in: Jörg Rothe (Ed.), *Economics and Computation*, Springer, 2016, pp. 395–491, chapter 7.

- [133] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, Amin Saberi, On approximately fair allocations of indivisible goods, in: Proceedings of the 5th ACM Conference on Electronic Commerce (EC), 2004, pp. 125–131.
- [134] Pasin Manurangsi, Warut Suksompong, When do envy-free allocations exist?, *SIAM J. Discrete Math.* 34 (3) (2020) 1505–1521.
- [135] Pasin Manurangsi, Warut Suksompong, Closing gaps in asymptotic fair division, *SIAM J. Discrete Math.* 35 (2) (2021) 668–706.
- [136] Pasin Manurangsi, Warut Suksompong, Almost envy-freeness for groups: improved bounds via discrepancy theory, in: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), 2021, pp. 335–341.
- [137] Evangelos Markakis, Approximation algorithms and hardness results for fair division with indivisible goods, in: Endriss Ulle (Ed.), Trends in Computational Social Choice, AI Access, 2017, pp. 231–247, chapter 12.
- [138] Hervé Moulin, Fair Division and Collective Welfare, MIT Press, 2003.
- [139] Hervé Moulin, Fair division in the age of Internet, *Annu. Rev. Econ.* (2018).
- [140] Noam Nisan, Bidding and allocation in combinatorial auctions, in: Proceedings of the 2nd ACM Conference on Electronic Commerce (EC), 2000, pp. 1–12.
- [141] Hoon Oh, Ariel D. Procaccia, Warut Suksompong, Fairly allocating many goods with few queries, *SIAM J. Discrete Math.* 35 (2) (2021) 788–813.
- [142] Benjamin Plaut, Tim Roughgarden, Almost envy-freeness with general valuations, *SIAM J. Discrete Math.* 34 (2) (2020) 1039–1068.
- [143] Ariel D. Procaccia, Cake cutting algorithms, in: Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, Ariel D. Procaccia (Eds.), Handbook of Computational Social Choice, Cambridge University Press, 2016, pp. 311–330, chapter 13.
- [144] Ariel D. Procaccia, Technical perspective: an answer to fair division’s most enigmatic question, *Commun. ACM* 63 (4) (2020) 118.
- [145] Alexandros Psomas, Paritosh Verma, Fair and efficient allocations without obvious manipulations, *CoRR*, arXiv:2206.11143 [abs], 2022.
- [146] Seddighin Masoud, Saeed Seddighin, Improved maximin guarantees for subadditive and fractionally subadditive fair allocation problem, in: Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI), 2022, pp. 5183–5190.
- [147] Masoud Seddighin, Hamed Saleh, Mohammad Ghodsi, Maximin share guarantee for goods with positive externalities, *Soc. Choice Welf.* 56 (2) (2021) 291–324.
- [148] Erel Segal-Halevi, Warut Suksompong, Democratic fair allocation of indivisible goods, *Artif. Intell.* 277 (2019).
- [149] Hugo Steinhaus, Sur la division pragmatique, *Econometrica* 17 (Supplement) (1949) 315–319.
- [150] Warut Suksompong, Approximate maximin shares for groups of agents, *Math. Soc. Sci.* 92 (2018) 40–47.
- [151] Warut Suksompong, Constraints in fair division, *SIGames Exch.* 19 (2) (2021) 46–61.
- [152] Ankang Sun, Bo Chen, Xuan Vinh Doan, Connections between fairness criteria and efficiency for allocating indivisible chores, in: Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2021, pp. 1281–1289.
- [153] Zhiyi Tan, Yong Wu, Optimal semi-online algorithms for machine covering, *Theor. Comput. Sci.* 372 (1) (2007) 69–80.
- [154] Hal R. Varian, Equity, envy and efficiency, *J. Econ. Theory* 9 (1974) 63–91.
- [155] Toby Walsh, Fair division: the computer scientist’s perspective, in: Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI), 2020, pp. 4966–4972.
- [156] Gerhard J. Woeginger, A polynomial-time approximation scheme for maximizing the minimum machine completion time, *Oper. Res. Lett.* 20 (4) (1997) 149–154.
- [157] Sheung Man Yuen, Warut Suksompong, Extending the characterization of maximum Nash welfare, *Econ. Lett.* 224 (2023) 111030.
- [158] David Zeng, Alexandros Psomas, Fairness-efficiency tradeoffs in dynamic fair division, in: Proceedings of the 21st ACM Conference on Economics and Computation (EC), 2020, pp. 911–912.
- [159] Shengwei Zhou, Xiaowei Wu, Approximately EFX allocations for indivisible chores, in: Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI), 2022, pp. 783–789.