Weighted Fair Division of Indivisible Items: A Review

Warut Suksompong

National University of Singapore

Abstract

Fair division is a longstanding problem in economics and has recently received substantial interest in computer science. Several applications of fair division involve agents with unequal entitlements represented by weights. We review work on weighted fair division of indivisible items, discuss the range of weighted fairness notions that have been proposed, and highlight a number of open questions.

1 Introduction

Given a set of valuable resources, how can we allocate it to interested agents with possibly differing preferences in a fair manner? This problem, also known as *fair division*, is fundamental in society and ubiquitous in everyday life. Applications of fair division are wide-ranging and include splitting inheritance among relatives, dividing usage rights of a facility between investors, as well as allocating supplies among organizations or communities. While fair division has long been studied by economists [Brams and Taylor, 1996, Robertson and Webb, 1998, Moulin, 2003], in the past decade it has also received substantial interest from computer scientists [Bouveret et al., 2016, Markakis, 2017, Aleksandrov and Walsh, 2020, Aziz, 2020, Suksompong, 2021, Amanatidis et al., 2023, Nguyen and Rothe, 2023].

Fair division may naturally evoke the idea of treating all involved agents equally, providing them with an equal opportunity to receive their desired resources. However, in several scenarios, fairness entails handling the agents unequally based on their *entitlements*. For instance, in inheritance division, closer relatives are typically more entitled to the inheritance than more distant ones. Another example is a situation where investors have made different investments into a facility—those who contributed more would understandably feel aggrieved if they did not receive a greater share of the usage rights. In a similar vein, larger organizations or communities normally deserve more of the allotted supplies and personnel. These entitlements can be represented by numerical *weights*, and the resulting problem is therefore referred to as *weighted fair division*. Clearly, approaches that do not take the weights into account cannot provide reasonable fairness guarantees in the aforementioned scenarios. Weighted fair division dates back to at least the seminal books by Brams and Taylor [1996] and Robertson and Webb [1998], and constitutes one of the most significant extensions of the basic fair division setting.[1] It encompasses the well-studied problem of *apportionment*, which corresponds to the special case where the items are identical, for example, seats in a parliament [Balinski and Young, 2001, Pukelsheim, 2014].

In this review, we provide a summary of recent developments in weighted fair division. Like the general fair division literature, the vast majority of work on weighted fair division in the past few years has concentrated on the allocation of *indivisible items*. Therefore, our focus will also be on this practically important case.[2] As we shall see, each fairness notion in the unweighted setting can usually be

[1] Other extensions include assuming that the items arrive in an online manner [Aleksandrov and Walsh, 2020], introducing constraints on the permissible allocations [Suksompong, 2021], and considering items of mixed nature [Liu et al., 2024].

[2] For *divisible items*, commonly modeled as a *cake*, recent work in the weighted setting has considered the query complexity of finding a fair allocation [Cseh and Fleiner, 2020] and the number of cuts necessary [Segal-Halevi, 2019a, Crew et al., 2020].

generalized to the weighted setting in multiple ways, and there are often trade-offs between different generalizations. The richness of the weighted setting is also reflected by the fact that while some unweighted approaches continue to ensure strong fairness guarantees in the presence of weights, others cannot even be extended to incorporate weights in a sensible manner. In addition to discussing and comparing the key notions and approaches in weighted fair division, we will also draw attention to several unresolved questions in this domain.

2 Preliminaries

In this section, we state definitions and notation for the fair allocation of indivisible items.

2.1 Unweighted Setting

Before we formally introduce weighted fair division, we first describe the simpler *unweighted setting*. There is a set of agents $N = \{1, 2, \ldots, n\}$ and a set of items $M = \{o_1, o_2, \ldots, o_m\}$. A *bundle* refers to a subset of items in M. Each agent $i \in N$ has a utility $u_i(M')$ for each bundle $M' \subseteq M$, where $u_i(\emptyset) = 0$. When M' consists of a single item o, we will write $u_i(o)$ instead of $u_i(\{o\})$ for convenience. Except in Section 5, we will assume that the items are *goods*. In particular, the agents' utility functions are *monotonic*, i.e., $u_i(M') \leq u_i(M'')$ for any $M' \subseteq M'' \subseteq M$; this implies that $u_i(o) \geq 0$ for every $o \in M$. Unless stated otherwise, we also assume that these functions are *additive*, i.e., $u_i(M') = \sum_{o \in M'} u_i(o)$ for every bundle M'. An *instance* of unweighted fair division is described by N, M, and $(u_i)_{i \in N}$. We will use g instead of o when discussing goods.

Given an instance, we are interested in allocating the items to the agents in a fair manner. An *allocation* $A = (A_1, \ldots, A_n)$ is a list of n bundles such that $A_i \cap A_j = \emptyset$ for all $i \neq j$; for each i, the bundle A_i is allocated to agent i. An allocation is said to be *complete* if $A_i \cup \cdots \cup A_n = M$. Unless stated otherwise, we assume that allocations must be complete. To measure the fairness of allocations, two of the oldest and most fundamental benchmarks are envy-freeness and proportionality.

Definition 2.1. In the unweighted setting, an allocation A is said to satisfy

- *envy-freeness (EF)* if for every pair of agents $i, j \in N$, it holds that $u_i(A_i) \geq u_i(A_j)$;
- *proportionality (PROP)* if for every agent $i \in N$, it holds that $u_i(A_i) \geq \frac{1}{n} \cdot u_i(M)$.

With indivisible goods, neither envy-freeness nor proportionality can always be satisfied—the simplest counterexample is when there are two agents and only one valuable good. It is therefore natural to relax these notions by allowing an "up to one good" approximation [Lipton et al., 2004, Budish, 2011, Aziz et al., 2022].

Definition 2.2. In the unweighted setting, an allocation A is said to satisfy

- *envy-freeness up to one good (EF1)* if for every pair of agents $i, j \in N$ for which $A_j \neq \emptyset$, there exists a good $g \in A_j$ such that $u_i(A_i) \geq u_i(A_j \setminus \{g\})$;
- *proportionality up to one good (PROP1)* if for every agent $i \in N$ for which $A_i \neq M$, there exists a good $g \notin A_i$ such that $u_i(A_i \cup \{g\}) \geq \frac{1}{n} \cdot u_i(M)$.

Every instance with arbitrary monotonic utilities admits an EF1 allocation [Lipton et al., 2004]. Moreover, if the utilities are additive, EF1 implies PROP1; this also means that a PROP1 allocation is guaranteed to exist under the additivity assumption.

Another route for circumventing the possible non-existence of a proportional allocation is to weaken the definition of the required share. By far the most extensively studied notion resulting from this approach is maximin share fairness [Budish, 2011]. Denote by $\Pi(M, n)$ the collection of all ordered partitions of M into n subsets.

Definition 2.3. In the unweighted setting, the *maximin share (MMS)* of an agent i is defined as

$$\mathrm{MMS}_i := \max_{(S_1,\dots,S_n)\in\Pi(M,n)} \min_{j\in\{1,\dots,n\}} u_i(S_j).$$

An allocation is said to be an *MMS allocation* if it gives every agent at least her MMS.

Even though an MMS allocation sometimes fails to exist, there is always an allocation that provides each agent with more than $3/4$ times her MMS [Kurokawa et al., 2018, Akrami and Garg, 2024].

Besides fairness, another important property in allocation problems is economic efficiency, which is often captured by the notion of Pareto optimality. An allocation A is said to be *Pareto optimal (PO)* if no other allocation makes at least one agent better off and no agent worse off. That is, that does not exist another allocation A' such that $u_i(A'_i) \geq u_i(A_i)$ for every $i \in N$ and the inequality is strict for at least one $i \in N$.

2.2 Weighted Setting

We are now ready to introduce the *weighted setting*, which is the focus of this review. In the weighted setting, each agent $i \in N$ is additionally endowed with a *weight* $w_i > 0$ representing her entitlement to the resource. Hence, besides N, M, and $(u_i)_{i\in N}$, an instance is also described by the weights $(w_i)_{i\in N}$. The unweighted setting corresponds to the special case where all weights are equal. For convenience, we denote $w_N := \sum_{i\in N} w_i$. Both envy-freeness and proportionality can be generalized to the weighted setting in an intuitive way.

Definition 2.4. In the weighted setting, an allocation A is said to satisfy

- *weighted envy-freeness (WEF)* if for every pair of agents $i, j \in N$, it holds that $\frac{u_i(A_i)}{w_i} \geq \frac{u_i(A_j)}{w_j}$;
- *weighted proportionality (WPROP)* if for every agent $i \in N$, it holds that $u_i(A_i) \geq \frac{w_i}{w_N} \cdot u_i(M)$.

Example 2.5. Consider three agents with weights $(w_1, w_2, w_3) = (1, 2, 4)$. In a WEF allocation, agent 1 finds her bundle to be worth no less than half of agent 2's bundle, while agent 2 believes that his bundle is at least twice as valuable as agent 1's. In a WPROP allocation, agent 3's utility for her own bundle is at least $4/7$ times her utility for the entire set of goods.

Since neither envy-freeness nor proportionality is always satisfiable, the same is true for their weighted generalizations. As we shall discuss, extending EF1, PROP1, and MMS to the weighted setting is a significantly more subtle task than for the former two notions. For brevity, unless stated otherwise, we consider the weighted setting in the rest of this review.

3 Fairness up to One Good

In this section, we review work on weighted fairness up to one good. By extending EF1 and PROP1 along with certain approaches for satisfying these notions, we will see that strong fairness guarantees can be attained in the weighted setting. Recall our assumptions from Section 2 that utilities are additive and allocations are complete unless stated otherwise.

3.1 WEF1 and WWEF1

Given the definitions of EF1 (Definition 2.2) and WEF (Definition 2.4), Chakraborty et al. [2021a] combined them into WEF1 in a natural manner.

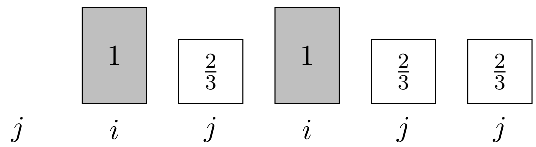

Figure 1: Illustration for the proof of Theorem 3.2.

Definition 3.1. An allocation A is said to satisfy *weighted envy-freeness up to one good (WEF1)* if for every pair of agents $i, j \in N$ for which $A_j \neq \emptyset$, there exists a good $g \in A_j$ such that

$$\frac{u_i(A_i)}{w_i} \geq \frac{u_i(A_j \setminus \{g\})}{w_j}.$$

In the unweighted setting, EF1 can be satisfied via the *round-robin algorithm*, which lets the agents take turns picking their favorite good from the remaining goods in cyclic order, breaking ties arbitrarily, until the goods run out. The proof is succinct: Consider any pair of agents i and j. If i picks before j, then i does not envy j in each 'round'; since utilities are additive, i also does not envy j overall. On the other hand, if i picks after j, then a similar argument shows that i does not envy j if we ignore j's first pick. Hence, in either case, EF1 from i towards j is fulfilled.

Can we generalize the round-robin algorithm to the weighted setting so as to satisfy WEF1? At first glance, it may seem unclear how to construct a suitable picking sequence based on the weights. However, as Chakraborty et al. [2021a] suggested, a general and intuitive way to devise such a sequence is to assign each pick to an agent who has picked least frequently so far, where the frequency is scaled by the agent's weight. For example, if there are three agents with weights $(w_1, w_2, w_3) = (1, 2, 4)$, a corresponding picking sequence is $(1, 2, 3, 3, 2, 3, 3, 1, 2, 3, \dots)$. Chakraborty et al. showed that the output of the constructed picking sequence is guaranteed to satisfy WEF1.

Theorem 3.2 ([Chakraborty et al., 2021a]). *Consider a picking sequence such that each pick is assigned to an agent i who minimizes t_i/w_i, where t_i denotes the number of times agent i has picked up to that point. An allocation chosen by this picking sequence always satisfies WEF1. Thus, there exists a polynomial-time algorithm for computing a WEF1 allocation.*

The proof of Theorem 3.2 is more involved than the aforementioned succinct proof of its unweighted counterpart. To explain it, we first describe a useful way to visualize the unweighted proof. Suppose we consider envy from agent i towards agent j. Every time i picks a good, we give her a bucket with 1 unit of water. Every time j picks a good from the second time onwards, we give him an empty bucket of capacity 1. Agent i is allowed to pour water from any bucket into any of j's buckets that comes later in the picking sequence. Since i values a good that she picks at least as much as any good that j picks in a later turn, in order to establish EF1, it suffices to show that i can fill up all of j's buckets using such operations. A similar idea can be used in the weighted setting, except that in order to account for the weights, every time j picks after the first time, we give her a bucket of capacity w_i/w_j instead of 1—see Figure 1 for an illustration when $(w_i, w_j) = (2, 3)$. However, unlike in the unweighted setting, where i can accomplish this task by simply pouring all the water from each of her buckets into j's subsequent bucket, in the weighted case, i may need to pour water from a bucket into several of j's buckets.

Besides the round-robin algorithm, two other prominent ways of satisfying EF1 in the unweighted setting are the *envy cycle elimination algorithm* [Lipton et al., 2004] and the *maximum Nash welfare (MNW)* solution [Caragiannis et al., 2019]. As Chakraborty et al. [2021a] observed, the envy cycle elimination algorithm does not admit a natural generalization to the weighted setting. Indeed, in the unweighted setting, if agent i envies agent j, then i would prefer to swap bundles with j. However, in the weighted case, even if i envies j with respect to the weights, it could be that i's weight is larger than j, and i in fact does not want to make the swap. Hence, the idea of eliminating envy cycles fundamentally fails in the presence of weights.

Does the MNW solution fare any better? In the unweighted setting, it chooses an allocation that maximizes the Nash welfare, which corresponds to the product $\prod_{i \in N} u_i(A_i)$ of the agents' utilities.[3] With weights, MNW can be generalized to *maximum weighted Nash welfare (MWNW)*, where the weighted Nash welfare is given by the product $\prod_{i \in N} u_i(A_i)^{w_i}$ with the weights in the exponents.[4] Given that MNW is known to satisfy EF1 [Caragiannis et al., 2019], it perhaps comes as a surprise that MWNW in fact fails to satisfy WEF1 [Chakraborty et al., 2021a].

Example 3.3. Consider two agents with weights $(w_1, w_2) = (1, 4)$ and seven identical goods, each yielding utility 1 to each agent. MWNW assigns one good to agent 1 and the remaining six goods to agent 2, which results in a violation of WEF1 from agent 1 towards agent 2.

From this counterexample, one can see that the WEF1 condition is rather stringent when $w_i < w_j$, as the utility of the removed good is scaled by $1/w_j$. This motivates the following relaxation of WEF1.

Definition 3.4. An allocation A is said to satisfy *weak weighted envy-freeness up to one good (WWEF1)* if for every pair of agents $i, j \in N$ for which $A_j \neq \emptyset$, there exists a good $g \in A_j$ such that

$$\frac{u_i(A_i)}{w_i} \geq \frac{u_i(A_j \setminus \{g\})}{w_j} \quad \text{or} \quad \frac{u_i(A_i \cup \{g\})}{w_i} \geq \frac{u_i(A_j)}{w_j}.$$

WWEF1 allows two options for eliminating envy from agent i towards agent j: either remove one good from j's bundle (as in WEF1), or copy a good from j's bundle into i's bundle. The former option is more potent when $w_i > w_j$, while the latter is more effective if $w_i < w_j$. It turns out that this weaker fairness benchmark is met by MWNW.

Theorem 3.5 ([Chakraborty et al., 2021a])**.** *Every MWNW allocation satisfies WWEF1 and PO.*

Theorem 3.5 generalizes an influential result of Caragiannis et al. [2019] from the unweighted setting and provides a concrete fairness guarantee for MWNW. In part due to this guarantee, the computation of an MWNW solution has recently been a subject of interest [Garg et al., 2022, Brown et al., 2024, Feng and Li, 2024].

Chakraborty et al. [2021a] further showed, by adapting a market-based algorithm of Barman et al. [2018], that an allocation fulfilling WEF1 and PO always exists and can be computed in pseudo-polynomial time. Whether such an allocation can be computed in polynomial time remains open even in the unweighted case.[5]

3.2 WPROP1

Next, we consider WPROP1, which is an intuitive combination of PROP1 (Definition 2.2) and WPROP (Definition 2.4) proposed by Aziz et al. [2020].

Definition 3.6. An allocation A is said to satisfy *weighted proportionality up to one good (WPROP1)* if for every agent $i \in N$ for which $A_i \neq M$, there exists a good $g \notin A_i$ such that

$$u_i(A_i \cup \{g\}) \geq \frac{w_i}{w_N} \cdot u_i(M).$$

[3] Ties can be broken arbitrarily except when the maximum possible product is 0, in which case more care is needed. Also, to be more precise, the Nash welfare is defined as the *geometric mean* of the utilities, but an allocation maximizing that objective is the same as one maximizing the product objective.

[4] To see why having weights in the exponents is sensible, observe that if the goods were divisible and homogeneous, this definition of MWNW would allocate the goods to the agents precisely in proportion to their respective weights.

[5] When there are two agents, Chakraborty et al. [2021a] showed that this can be done using a weighted variant of the classic *adjusted winner procedure* [Brams and Taylor, 1996].

In fact, Aziz et al. defined WPROP1 in a more general setting, allowing items to yield positive utilities to some agents and negative utilities to others. They demonstrated that when using WPROP1 as the benchmark, the computation of a fair and efficient allocation becomes tractable.

Theorem 3.7 ([Aziz et al., 2020])**.** *There exists a strongly polynomial-time algorithm that computes an allocation satisfying WPROP1 and PO.*

The proof of Theorem 3.7 relies on first treating the items as divisible, starting from a WPROP division of the items, and conducting sequential cyclic trades to arrive at a Pareto-dominating allocation. The resulting allocation satisfies a stronger property than PO called "fractional PO", and it can be rounded into an integral allocation that maintains fractional PO and at the same time satisfies WPROP1.

While EF1 implies PROP1 and WEF implies WPROP, Chakraborty et al. [2021a] showed that, surprisingly, WEF1 does not imply WPROP1. Worse still, there exist instances in which the two properties are conflicting.

Example 3.8. Consider an instance with $n = 10$ agents with weights $(w_1, \ldots, w_{10}) = (1, 1, \ldots, 1, 100)$ and $m = 10$ identical goods, each yielding utility 1 to each agent. In every WEF1 allocation, each agent must receive exactly one good; otherwise the WEF1 condition from an agent with no good towards an agent with more than one good would be violated. On the other hand, a WPROP1 allocation necessarily assigns at least nine goods to the last agent. Hence, no allocation satisfies both WEF1 and WPROP1.

The inherent incompatibility between WEF1 and WPROP1 is highly counterintuitive and prompts a need to revisit the definitions of these two properties, as we shall discuss next.

3.3 WEF(x, y) and WPROP(x, y)

An alternative formulation of WEF1 is that the *weighted envy* from agent i towards agent j, defined as the quantity $\max\left\{0, \frac{u_i(A_j)}{w_j} - \frac{u_i(A_i)}{w_i}\right\}$, is allowed to be up to $\frac{u_i(g)}{w_j}$ for some $g \in A_j$. From this perspective, there is no clear reason why $u_i(g)$ should be multiplied by $1/w_j$ instead of $1/w_i$.[6] Chakraborty et al. [2022] proposed a general definition where the multiplier is an arbitrary linear combination of $1/w_j$ and $1/w_i$.

Definition 3.9. Let $x, y \geq 0$ be real numbers. An allocation A is said to satisfy *WEF*(x, y) if for every pair of agents $i, j \in N$ for which $A_j \neq \emptyset$, there exists a good $g \in A_j$ such that

$$\frac{u_i(A_i) + y \cdot u_i(g)}{w_i} \geq \frac{u_i(A_j) - x \cdot u_i(g)}{w_j},$$

or equivalently,

$$\frac{u_i(A_j)}{w_j} - \frac{u_i(A_i)}{w_i} \leq \left(\frac{y}{w_i} + \frac{x}{w_j}\right) \cdot u_i(g).$$

Note that WEF and WEF1 correspond to WEF$(0, 0)$ and WEF$(1, 0)$, respectively. WEF$(1, 1)$ allows one good to be transferred from j's bundle to i's bundle in order to eliminate the weighted envy, and has therefore been called "transfer weighted envy-freeness up to one good" [Chakraborty et al., 2021a] or "weighted envy-freeness up to one transfer" [Aziz et al., 2023b]. If all weights are equal, the WEF(x, y) notion depends only on the sum $x + y$; in particular, if $x + y = 1$, the notion reduces to EF1. However, when weights can be different, each choice of (x, y), even with the same sum, leads to a different condition: a larger x induces a stronger guarantee for agents with lower weights, while a larger y ensures less envy for those with higher weights.

[6] For WWEF1, the multiplier is $1/\min\{w_i, w_j\}$.

If $x+y<1$, an instance with two equal-weight agents and one valuable good shows that WEF(x, y) cannot always be satisfied. In light of this, the strongest existence guarantee one could hope for is when $x+y=1$. It turns out that this guarantee indeed holds, and moreover there is an inherent trade-off between different values of (x, y).

Theorem 3.10 ([Chakraborty et al., 2022]). *For each $x \in [0,1]$, there exists an allocation satisfying WEF$(x, 1-x)$, which can be computed in polynomial time.*

On the other hand, for each pair of distinct $x, x' \in [0,1]$, there exists an instance with identical goods, each yielding utility 1 to each agent, such that no allocation satisfies both WEF$(x, 1-x)$ and WEF$(x', 1-x')$.

It is worth noting that WEF$(x, 1-x)$ is strictly stronger than WWEF1 for every x. The algorithm for computing a WEF$(x, 1-x)$ allocation is a generalization of the picking sequence in Theorem 3.2—instead of the ratio t_i/w_i, one can use $(t_i+(1-x))/w_i$. Interestingly, when all goods are identical, these picking sequences belong to a well-studied class of apportionment methods called *divisor methods*. As for the incompatibility, the intuition is similar to that of Example 3.8: there is an inevitable compromise between ensuring that low-weight agents are not empty-handed and providing high-weight agents with their due share.

Chakraborty et al. [2022] defined WPROP(x, y) in an analogous manner as WEF(x, y).

Definition 3.11. Let $x, y \geq 0$ be real numbers. An allocation A is said to satisfy *WPROP(x, y)* if for every agent $i \in N$ for which $A_i \neq M$, there exists a good $g \notin A_i$ such that

$$u_i(A_i) + y \cdot u_i(g) \geq \frac{w_i}{w_N} \cdot (u_i(M) - n \cdot x \cdot u_i(g)).$$

WPROP is the same as WPROP$(0,0)$, but WPROP1 corresponds to WPROP$(0,1)$ rather than WPROP$(1,0)$. This explains the counterintuitive phenomenon described at the end of Section 3.2: there is a 'mismatch' between WEF1 and WPROP1, which leads to the incompatibility between the two notions. Indeed, Chakraborty et al. [2022] showed that WEF(x, y) implies WPROP(x, y) for all x, y; this generalizes the facts that EF1 implies PROP1 and WEF implies WPROP. They also established analogous results for WPROP(x, y) as for WEF(x, y) (Theorem 3.10), and showed that MWNW does not imply WEF$(x, 1-x)$ or WPROP$(x, 1-x)$ for any x. This raises our first open problem.

Open problem 1. *For each $x \in [0,1)$, does there always exist an allocation satisfying WEF$(x, 1-x)$ and PO? For each $x \in (0,1)$, does there always exist an allocation satisfying WPROP$(x, 1-x)$ and PO?*

While the case $x+y=1$ is the most interesting in terms of existence by itself, when other requirements are added, it may become necessary to consider larger values of $x+y$. Indeed, this is the case in the "best of both worlds" framework, which aims to offer fairness guarantees for random allocations both before and after the randomization (so-called *ex-ante* and *ex-post*, respectively) [Aziz et al., 2023a]. Aziz et al. [2023b] and Hoefer et al. [2023] showed that while ex-ante WEF and ex-post WEF(x, y) can be satisfied simultaneously when $x=y=1$, the same is not true for any other pair $x, y \in [0,1]$.

3.4 WEFX

Besides EF1, another well-studied relaxation of envy-freeness in the unweighted setting is *envy-freeness up to any good (EFX)*, which requires the envy from one agent towards another agent to disappear as soon as *any* good in the latter agent's bundle is removed. EFX can be extended to accommodate weights in a similar way as EF1.

Definition 3.12. An allocation A is said to satisfy *weighted envy-freeness up to any good (WEFX)* if for every pair of agents $i, j \in N$ and every good $g \in A_j$, it holds that

$$\frac{u_i(A_i)}{w_i} \geq \frac{u_i(A_j \setminus \{g\})}{w_j}.$$

With equal weights, an EFX allocation is guaranteed to exist for up to three agents, but the existence question is open beyond three agents [Akrami et al., 2023]. By contrast, Hajiaghayi et al. [2024] showed that with unequal weights, a WEFX allocation may not exist even for two agents.

Example 3.13. Consider an instance with $n = 2$ agents with weights $(w_1, w_2) = (3, 4)$, $m = 4$ goods, and the following utilities:

	g_1	g_2	g_3	g_4
Agent 1	0	2	3	3
Agent 2	0	0	1	1

Since agent 2 has a higher weight, for WEFX to be satisfied, she must receive either g_3 and g_4 together, or one of g_3, g_4 along with both of g_1, g_2. However, one can check that in either case, WEFX is violated for agent 1.

In view of this example, one could consider relaxing WEFX to α-*WEFX*, where $\alpha < 1$ is a multiplicative factor on the right-hand side of the inequality in Definition 3.12.[7] By using similar utilities as in Example 3.13, Hajiaghayi et al. [2024] showed that one cannot always attain α-WEFX for any $\alpha > 0.786$, and proved guaranteed existence for a factor α that degrades with m, the number of goods.

Open problem 2. *For $n = 2$ agents, is there a positive constant α such that an α-WEFX allocation always exists? What about for larger n?*

3.5 Non-Additive Utilities

Although utilities are frequently assumed to be additive in the fair division literature, certain goods in practical applications may exhibit complementarity or substitution effects, resulting in non-additive utilities. In the unweighted setting, an EF1 allocation is guaranteed to exist even for arbitrary monotonic utilities [Lipton et al., 2004]. However, as Chakraborty et al. [2021a] showed, this is not true for WEF1.

Example 3.14. Consider an instance with $n = 2$ agents with weights $(w_1, w_2) = (1, 2)$, and $m = 6$ goods. Agent 1 has an additive utility function with utility 1 for every good, while agent 2 has utility 0 for the empty bundle and utility 1 for any other bundle. If agent 1 receives more than one good, agent 2 has weighted envy towards agent 1 even after removing any good from agent 1's bundle. On the other hand, if agent 1 receives at most one good, WEF1 from agent 1 towards agent 2 is violated. Thus, there is no WEF1 allocation in this instance.

By increasing the number of goods in Example 3.14, one can show that the impossibility persists even if WEF1 is relaxed to WWEF1 and the 'up to one good' condition is weakened to 'up to c goods' for any constant c. Note that the utility functions in this example are particularly simple: both of them are *binary submodular*, also known as *matroid-rank*.[8] Such utility functions are relevant when allocating

[7]Multiplicative approximations have also been studied for EFX in the unweighted setting [Amanatidis et al., 2020, Plaut and Roughgarden, 2020].

[8]A monotonic utility function u is called *binary* if the utility of any set of goods increases by 0 or 1 upon the addition of a new good. It is called *submodular* if the gain in utility from adding a good to a set is at most the gain from adding the same good to its subset.

course slots to students or dividing public housing estates among ethnic groups [Babaioff et al., 2021b, Benabbou et al., 2021]

Montanari et al. [2024] proposed two families of notions tailored for submodular utilities. The first family, TWEF(x, y), is based on the concept of "transferability" studied in the unweighted setting by Benabbou et al. [2021]. TWEF(x, y) stipulates that an agent i should not be considered to envy another agent j if i's utility would not increase upon transferring all of j's goods to i's bundle, and therefore rules out such scenarios as Example 3.14, where agent 2 remains envious even when already at maximum utility. The second family, WMEF(x, y), is an extension of the notion *marginal EF1 (MEF1)* of Caragiannis et al. [2019] from the unweighted setting. The idea is that, instead of comparing i's utility for her own bundle with her utility for j's bundle as in WEF1, we compare it with her marginal utility of j's bundle given i's bundle (i.e., $u_i(A_i \cup A_j) - u_i(A_i)$); for submodular utilities, this results in a weaker benchmark. Importantly, when utilities are additive, both TWEF(x, y) and WMEF(x, y) reduce to WEF(x, y), which in turn reduces to EF1 if all weights are equal and $x + y = 1$.

Montanari et al. [2024] showed that TWEF$(x, 1 - x)$ can be satisfied for every x when utilities are matroid-rank,[9] and the same holds for WMEF$(x, 1 - x)$ when utilities are general submodular. As a result, meaningful fairness guarantees can be made in weighted fair division even for non-additive utilities. Interestingly, these authors also demonstrated that for matroid-rank utilities, the *maximum weighted harmonic welfare (MWHW)* rule, defined based on harmonic numbers, offers better guarantees than MWNW. Viswanathan and Zick [2023] proved that, under the matroid-rank assumption, there exists a polynomial-time algorithm that can optimize a range of objectives, including MWNW and MWHW.

Despite these positive results for submodular utilities, it remains intriguingly open whether similar fairness guarantees can be extended to other significant classes of utilities.

Open problem 3. *For supermodular utilities, does there exist a meaningful envy-freeness notion that can always be satisfied and moreover reduces to EF1 in the unweighted additive setting? What about for subadditive utilities?*

3.6 Monotonicity and Strategyproofness

In addition to fairness and efficiency, two other sets of desirable properties for allocation rules are monotonicity and strategyproofness. In particular, *resource-monotonicity* stipulates that when an extra good is added, no agent's utility should decrease as a result. Likewise, *population-monotonicity* means that introducing an additional agent should not benefit any existing agent, while *weight-monotonicity* states that if an agent's weight increases, her utility should not go down. On the other hand, *strategyproofness* requires that agents do not have an incentive to misreport their utility functions. A stronger version, *group-strategyproofness*, asserts that no group of agents can misreport in such a way that all agents in the group strictly benefit.

While the monotonicity properties are intuitive and one may be tempted to think that any reasonable allocation rule should satisfy them, this is in fact not the case. As Chakraborty et al. [2021b] showed, even in the unweighted setting, prominent rules such as MNW and envy cycle elimination fail resource-monotonicity. Nevertheless, these authors proved that several picking sequences—including those in Theorems 3.2 and 3.10—satisfy resource- and population-monotonicity for any number of agents n, as well as weight-monotonicity for $n = 2$. Combined with their fairness guarantees and their simplicity, this establishes picking sequences as strong candidate rules for weighted fair division.

Open problem 4. *For each $n \geq 3$, does there exist a rule that satisfies resource-, population-, and weight-monotonicity along with WWEF1?*

As for strategyproofness, with equal weights and additive utilities, it is known that no rule satisfies both EF1 and strategyproofness [Amanatidis et al., 2017], and the only rule that is both Pareto

[9] However, instead of requiring allocations to be complete, they ensured other efficiency properties.

optimal and strategyproof is serial dictatorship [Klaus and Miyagawa, 2002]. Nevertheless, Suksompong and Teh [2023] showed that with matroid-rank utilities and possibly different weights, MWNW and MWHW (with specific tie-breaking) satisfy all three monotonicity properties as well as group-strategyproofness.[10] Together with the fairness and efficiency guarantees of these rules, this implies that virtually all desirable properties can be fulfilled simultaneously under matroid-rank utilities.

4 Share-Based Notions

In this section, we turn our attention to share-based notions. As mentioned in Section 2, the most widely studied share-based notion in the unweighted setting is maximin share fairness (MMS). Like EF1 and PROP1, we will see that there are several valid ways to extend MMS to the weighted setting.

4.1 WMMS and NMMS

The first weighted extension of MMS was proposed by Farhadi et al. [2019].

Definition 4.1. The *weighted maximin share (WMMS)* of an agent i is defined as

$$\text{WMMS}_i := \max_{(S_1,\dots,S_n)\in\Pi(M,n)} \min_{j\in\{1,\dots,n\}} \frac{w_i}{w_j} \cdot u_i(S_j).$$

For $c \leq 1$, we say that an allocation is *c-WMMS* if it yields utility at least $c \cdot \text{WMMS}_i$ for every agent i. We will use analogous terminology for the remaining share-based notions.

Intuitively, the definition of WMMS involves finding the 'most proportional' allocation with respect to all agents' weights and agent i's utility function. In contrast to MMS, no constant approximation of WMMS can be guaranteed.

Theorem 4.2 ([Farhadi et al., 2019])**.** *There always exists a $1/n$-WMMS allocation, and the factor $1/n$ cannot be improved in the worst case. Moreover, such an allocation can be found in polynomial time.*

To find a $1/n$-WMMS allocation, one can use the simple (unweighted) round-robin algorithm, with an additional specification that the agents pick in non-increasing order of weight. Nevertheless, a drawback of WMMS is that the WMMS of an agent depends not only on the agent's weight (relative to the total weight), but also on the weight of every other agent. Moreover, the definition of WMMS is rather difficult to understand, especially when compared to MMS. With these observations in mind, Chakraborty et al. [2022] proposed an alternative extension of MMS.

Definition 4.3. The *normalized maximin share (NMMS)* of an agent i is defined as

$$\text{NMMS}_i := \frac{w_i}{w_N} \cdot n \cdot \text{MMS}_i.$$

The NMMS of an agent depends only on the agent's relative weight, and since NMMS is simply an appropriately scaled version of MMS, it is arguably easier to understand. Chakraborty et al. [2022] showed that $1/n$-NMMS can be guaranteed (by the same algorithm as $1/n$-WMMS), and this factor is again tight. Perhaps more notably, they showed that WEF1 implies $1/n$-NMMS, which means that any algorithm that yields WEF1 (cf. Section 3.1) ensures $1/n$-NMMS as well.[11] This generalizes a result of Amanatidis et al. [2018] that EF1 implies $1/n$-MMS in the unweighted setting.

[10]They showed this for a broader class of rules called *weighted additive welfarist rules*. Their result generalizes earlier results by Halpern et al. [2020] and Suksompong and Teh [2022] in more restricted settings.

[11]On the other hand, WEF1 does not imply any positive approximation of WMMS.

4.2 OMMS and APS

The fact that neither WMMS nor NMMS admits a constant approximation raises the question of whether there exist any natural shares that allow such an approximation. As it turns out, the answer to this question is positive, as exhibited by the following two shares.

Definition 4.4. For any positive integers $\ell \leq d$, let

$$\text{MMS}_i^{\ell\text{-out-of-}d} := \max_{S \in \Pi(M,d)} \min_{T \in \text{Unions}(S,\ell)} u_i(T),$$

where $\text{Unions}(S, \ell)$ denotes the collection of all unions of ℓ bundles from the d-partition S.[12] Then, the *ordinal maximin share (OMMS)* of an agent i is defined as

$$\text{OMMS}_i := \max_{\ell,d:\, \frac{\ell}{d} \leq \frac{w_i}{w_N}} \text{MMS}_i^{\ell\text{-out-of-}d}.$$

Definition 4.5. The *AnyPrice share (APS)* of an agent i is defined as

$$\text{APS}_i := \max_{P \in \text{AllowedCollections}(M,w_i)} \min_{T \in P} u_i(T).$$

Here, $\text{AllowedCollections}(M, w_i)$ includes all collections P of bundles such that some assignment of weights to the bundles in P satisfies the following two properties: (i) the total weight of all bundles in P is w_N; (ii) for each good, the total weight of the bundles to which the good belongs is at most w_i.

OMMS was implicitly considered by Babaioff et al. [2021c] and explicitly studied by Segal-Halevi [2019b].[13] APS was introduced by Babaioff et al. [2021a], who established several important properties of it. While Babaioff et al. [2021a] defined APS for non-additive utilities, their main results hold only for additive utilities. Note that of the four weighted share-based notions, APS is the only one that does not reduce to MMS in the unweighted setting.

Theorem 4.6 ([Babaioff et al., 2021a]). *There always exists a $3/5$-APS allocation, and such an allocation can be computed in polynomial time.*

Babaioff et al. showed Theorem 4.6 by analyzing a bidding game in which the weights serve as budgets and at each round the highest bidder wins, taking any goods she wants and paying her bid for each good taken. The same authors also proved that an agent's APS is always at least her OMMS, which means that a $3/5$-OMMS allocation always exists as well. Hence, a significant fraction of both OMMS and APS can be guaranteed for arbitrary instances.

Open problem 5. *What is the largest constant c such that a c-APS allocation always exists? What about c-OMMS?*

5 Chores

Up until now, our focus has been on the case of goods, which is also the primary focus of the fair division literature. Nevertheless, the opposite case of *chores*—items that yield disutility to agents—is also relevant in several applications and has received due interest. We assume throughout this section that utilities are additive and $u_i(o) \leq 0$ for all $i \in N$ and $o \in M$.

With chores, the weights are interpreted as obligations rather than entitlements: the higher the weight, the more obligation the agent has. The definition of WEF (Definition 2.4) can be used as is. However, WEF1 needs to be modified so that instead of an agent removing a good from another agent's bundle, she removes a chore from her own bundle.

[12] Thus, the canonical MMS corresponds to 1-out-of-n MMS.

[13] It was also called *pessimistic share* by Babaioff et al. [2021a].

Definition 5.1. In chore division, an allocation A is said to satisfy *weighted envy-freeness up to one chore (WEF1)* if for every pair of agents $i, j \in N$ for which $A_i \neq \emptyset$, there exists a chore $o \in A_i$ such that $\frac{u_i(A_i \setminus \{o\})}{w_i} \geq \frac{u_i(A_j)}{w_j}$.

By modifying the picking sequence from the case of goods (Theorem 3.2), both Wu et al. [2023] and Hajiaghayi et al. [2024] established the existence of a WEF1 allocation.

Theorem 5.2 ([Wu et al., 2023, Hajiaghayi et al., 2024])**.** *In chore division, a WEF1 allocation always exists and can be computed in polynomial time.*

As in the case of goods, Hajiaghayi et al. [2024] showed that a WEFX allocation does not necessarily exist for chores. Aziz et al. [2020] proved that WPROP1 and PO can be satisfied together even for combinations of goods and chores. Li et al. [2022] demonstrated that WPROPX, a stronger notion than WPROP1, can always be fulfilled and moreover provides a constant approximation to APS for chores. Feige and Huang [2023] established the existence of an allocation that gives every agent no worse than 1.733 times her APS; they achieved this by reducing any given instance to one in which all agents agree on the ordering of the chores, considering a fractional allocation of the chores and rounding it into an integral allocation, and running a picking sequence defined based on this integral allocation. Aziz et al. [2019] investigated WMMS in chore division, although their positive results are restricted to the case of two agents or binary valuations.

Open problem 6. *In chore division, what is the best approximation of WMMS that can be guaranteed?*

6 Conclusion

Weighted fair division constitutes a significant extension of the basic fair division framework and has garnered considerable attention in recent times. As we have seen throughout this review, the weighted setting is often richer and more challenging than its unweighted counterpart. Indeed, each notion from the unweighted case can typically be extended to accommodate weights in multiple ways. Moreover, several methods that work well in the absence of weights cease to do so when weights are involved.

Among a wide range of potential future directions, we highlight two that we find particularly interesting.

Open problem 7. *Although MMS has been explored in the context of non-additive utilities [Barman and Krishnamurthy, 2020, Ghodsi et al., 2022], the study of weighted share-based notions has been mostly restricted to additive utilities. Are there appropriate share-based notions for non-additive utilities in the weighted setting?*

Open problem 8. *The notion WEF1 can be extended to handle combinations of goods and chores in a similar way as in the unweighted setting [Aziz et al., 2022]. Can this extension, or some relaxation of it, always be satisfied?*

Acknowledgments

The author thanks Ioannis Caragiannis for encouraging him to write this review and the anonymous referees for helpful comments, and acknowledges support from the Singapore Ministry of Education under grant number MOE-T2EP20221-0001 and an NUS Start-up Grant.

References

Hannaneh Akrami and Jugal Garg. Breaking the $3/4$ barrier for approximate maximin share. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 74–91, 2024.

Hannaneh Akrami, Noga Alon, Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, and Ruta Mehta. EFX: a simpler approach and an (almost) optimal guarantee via rainbow cycle number. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, page 61, 2023.

Martin Aleksandrov and Toby Walsh. Online fair division: a survey. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 13557–13562, 2020.

Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 545–562, 2017.

Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. Comparing approximate relaxations of envy-freeness. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 42–48, 2018.

Georgios Amanatidis, Apostolos Ntokos, and Evangelos Markakis. Multiple birds with one stone: Beating $1/2$ for EFX and GMMS via envy cycle elimination. *Theoretical Computer Science*, 841:94–109, 2020.

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

Haris Aziz. Developments in multi-agent fair allocation. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 13563–13568, 2020.

Haris Aziz, Hau Chan, and Bo Li. Weighted maxmin fair share allocation of indivisible chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 46–52, 2019.

Haris Aziz, Hervé Moulin, and Fedor Sandomirskiy. A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation. *Operations Research Letters*, 48(5):573–578, 2020.

Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi-Agent Systems*, 36(1):3:1–3:21, 2022.

Haris Aziz, Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of both worlds: Ex ante and ex post fairness in resource allocation. *Operations Research*, 2023a.

Haris Aziz, Aditya Ganguly, and Evi Micha. Best of both worlds fairness under entitlements. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 941–948, 2023b.

Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair-share allocations for agents with arbitrary entitlements. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, page 127, 2021a.

Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5119–5126, 2021b.

Moshe Babaioff, Noam Nisan, and Inbal Talgam-Cohen. Competitive equilibrium with indivisible goods and generic budgets. *Mathematics of Operations Research*, 46(1):382–403, 2021c.

Michel L. Balinski and H. Peyton Young. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Brookings Institution Press, 2001.

Siddharth Barman and Sanath Kumar Krishnamurthy. Approximation algorithms for maximin fair division. *ACM Transactions on Economics and Computation*, 8(1):5:1–5:28, 2020.

Siddharth Barman, Sanath Kumar Krisnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 557–574, 2018.

Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, and Yair Zick. Finding fair and efficient allocations for matroid rank valuations. *ACM Transactions on Economics and Computation*, 9(4): 21:1–21:41, 2021.

Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12, pages 284–310. Cambridge University Press, 2016.

Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

Adam Brown, Aditi Laddha, Madhusudhan Reddy Pittu, and Mohit Singh. Approximation algorithms for the weighted Nash social welfare via convex and non-convex programs. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1307–1327, 2024.

Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3):12:1–12:32, 2019.

Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. Weighted envy-freeness in indivisible item allocation. *ACM Transactions on Economics and Computation*, 9(3):18:1–18:39, 2021a.

Mithun Chakraborty, Ulrike Schmidt-Kraepelin, and Warut Suksompong. Picking sequences and monotonicity in weighted fair division. *Artificial Intelligence*, 301:103578, 2021b.

Mithun Chakraborty, Erel Segal-Halevi, and Warut Suksompong. Weighted fairness notions for indivisible items revisited. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pages 4949–4956, 2022.

Logan Crew, Bhargav Narayanan, and Sophie Spirkl. Disproportionate division. *Bulletin of the London Mathematical Society*, 52(5):885–890, 2020.

Ágnes Cseh and Tamás Fleiner. The complexity of cake cutting with unequal shares. *ACM Transactions on Algorithms*, 16(3):29:1–29:21, 2020.

Alireza Farhadi, Mohammad Ghodsi, MohammadTaghi Hajiaghayi, Sebastien Lahaie, David Pennock, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research*, 64:1–20, 2019.

Uriel Feige and Xin Huang. On picking sequences for chores. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, pages 626–655, 2023.

Yuda Feng and Shi Li. A note on approximating weighted Nash social welfare with additive valuations. In *Proceedings of the 51st International Colloquium on Automata, Languages, and Programming (ICALP)*, 2024. Forthcoming.

Jugal Garg, Edin Husić, Aniket Murhekar, and László Végh. Tractable fragments of the maximum Nash welfare problem. In *Proceedings of the 18th Conference on Web and Internet Economics (WINE)*, pages 362–363, 2022.

Mohammad Ghodsi, MohammadTaghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods: Beyond additive valuations. *Artificial Intelligence*, 303: 103633, 2022.

MohammadTaghi Hajiaghayi, Max Springer, and Hadi Yami. Almost envy-free allocations of indivisible goods or chores with entitlements. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 9901–9908, 2024.

Daniel Halpern, Ariel D. Procaccia, Alexandros Psomas, and Nisarg Shah. Fair division with binary valuations: One rule to rule them all. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*, pages 370–383, 2020.

Martin Hoefer, Marco Schmalhofer, and Giovanna Varricchio. Best of both worlds: Agents with entitlements. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 564–572, 2023.

Bettina Klaus and Eiichi Miyagawa. Strategy-proofness, solidarity, and consistency for multiple assignment problems. *International Journal of Game Theory*, 30(3):421–435, 2002.

David Kurokawa, Ariel D. Procaccia, and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM*, 64(2):8:1–8:27, 2018.

Bo Li, Yingkai Li, and Xiaowei Wu. Almost (weighted) proportional allocations for indivisible chores. In *Proceedings of the 31st ACM Web Conference (WWW)*, pages 122–131, 2022.

Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 125–131, 2004.

Shengxin Liu, Xinhang Lu, Mashbat Suzuki, and Toby Walsh. Mixed fair division: a survey. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 22641–22649, 2024.

Evangelos Markakis. Approximation algorithms and hardness results for fair division. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 12, pages 231–247. AI Access, 2017.

Luisa Montanari, Ulrike Schmidt-Kraepelin, Warut Suksompong, and Nicholas Teh. Weighted envy-freeness for submodular valuations. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 9865–9873, 2024.

Hervé Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.

Trung Thanh Nguyen and Jörg Rothe. Complexity results and exact algorithms for fair division of indivisible items: a survey. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 6732–6740, 2023.

Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2):1039–1068, 2020.

Friedrich Pukelsheim. *Proportional Representation: Apportionment Methods and Their Applications*. Springer, 2014.

Jack Robertson and William Webb. *Cake-Cutting Algorithms: Be Fair if You Can*. Peters/CRC Press, 1998.

Erel Segal-Halevi. Cake-cutting with different entitlements: How many cuts are needed? *Journal of Mathematical Analysis and Applications*, 480(1):123382, 2019a.

Erel Segal-Halevi. The maximin share dominance relation. *CoRR*, abs/1912.08763v1, 2019b.

Warut Suksompong. Constraints in fair division. *ACM SIGecom Exchanges*, 19(2):46–61, 2021.

Warut Suksompong and Nicholas Teh. On maximum weighted Nash welfare for binary valuations. *Mathematical Social Sciences*, 117:101–108, 2022.

Warut Suksompong and Nicholas Teh. Weighted fair division with matroid-rank valuations: Monotonicity and strategyproofness. *Mathematical Social Sciences*, 126:48–59, 2023.

Vignesh Viswanathan and Yair Zick. A general framework for fair allocation with matroid rank valuations. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, pages 1129–1152, 2023.

Xiaowei Wu, Cong Zhang, and Shengwei Zhou. Weighted EF1 allocations for indivisible chores. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, page 1155, 2023.