

Weighted Envy-Freeness for Submodular Valuations

Luisa Montanari¹, Ulrike Schmidt-Kraepelin², Warut Suksompong³, Nicholas Teh⁴

¹Technische Universität Berlin, Germany

²TU Eindhoven, The Netherlands

³National University of Singapore, Singapore

⁴University of Oxford, UK

Abstract

We investigate the fair allocation of indivisible goods to agents with possibly different entitlements represented by weights. Previous work has shown that guarantees for additive valuations with existing envy-based notions cannot be extended to the case where agents have matroid-rank (i.e., binary submodular) valuations. We propose two families of envy-based notions for matroid-rank and general submodular valuations, one based on the idea of transferability and the other on marginal values. We show that our notions can be satisfied via generalizations of rules such as picking sequences and maximum weighted Nash welfare. In addition, we introduce welfare measures based on harmonic numbers, and show that variants of maximum weighted harmonic welfare offer stronger fairness guarantees than maximum weighted Nash welfare under matroid-rank valuations.

1 Introduction

Fair division refers to the study of how to fairly allocate resources among agents with possibly differing preferences. Over the 75 years since Steinhaus (1948) initiated a mathematical framework of fair division, the field has given rise to numerous fairness notions and procedures for computing fair outcomes in a variety of scenarios (Brams and Taylor 1996; Robertson and Webb 1998). For instance, in the common scenario of allocating indivisible goods, the notion *envy-freeness up to one good (EF1)* has emerged as a standard benchmark. An allocation of the goods satisfies EF1 if any envy that an agent has toward another agent can be eliminated by removing some good in the latter agent’s bundle. Even when agents have arbitrary monotonic valuations over the goods, an EF1 allocation always exists and can be found in polynomial time (Lipton et al. 2004).

The definitions of many fairness notions in the literature, including EF1, inherently assume that all agents have the same entitlement to the resource. Recently, several researchers have examined a more general model in which different agents may have different *weights* reflecting their entitlements to the goods (Farhadi et al. 2019; Aziz, Moulin, and Sandomirskiy 2020; Babaioff, Ezra, and Feige 2021b; Babaioff, Nisan, and Talgam-Cohen 2021; Chakraborty et al. 2021; Chakraborty, Schmidt-Kraepelin,

and Suksompong 2021; Suksompong and Teh 2022, 2023; Hoefer, Schmalhofer, and Varricchio 2023; Scarlett, Teh, and Zick 2023; Viswanathan and Zick 2023a). This model allows us to capture settings such as inheritance division, in which relatives are typically entitled to unequal shares of the legacy, as well as resource allocation among groups or organizations of different sizes. Chakraborty et al. (2021) generalized EF1 to *weighted EF1 (WEF1)*: for instance, if Alice’s weight is three times as high as Bob’s, then WEF1 stipulates that, after removing some good from Bob’s bundle, Alice should have at least three times as much value for her own bundle as for Bob’s. The same authors demonstrated that if agents have *additive* valuations over the goods, a complete WEF1 allocation always exists and can be computed efficiently.¹ However, they provided the following example showing that existence is no longer guaranteed once we move beyond additivity.

Example 1 (Chakraborty et al. (2021)). Consider an instance with $n = 2$ agents whose weights are $w_1 = 1$ and $w_2 = 2$, and $m \geq 6$ goods. Agent 1 has an additive valuation with value 1 for every good, whereas agent 2 has value 0 for the empty bundle and 1 for any nonempty bundle.

If agent 1 is allocated more than one good, then agent 2 has weighted envy toward agent 1 even after removing any good from agent 1’s bundle. Thus, assuming that all goods need to be allocated, agent 2 must obtain at least $m - 1$ goods in a WEF1 allocation. Again, this causes weighted envy according to WEF1, this time from agent 1 toward agent 2. Hence, no complete WEF1 allocation exists in this instance.

The impossibility result illustrated in this example still holds even if WEF1 is relaxed to *weak WEF1 (WWEF1)*, whereby an agent is allowed to either remove a good from the other agent’s bundle or copy one such good into her own bundle, and stands in contrast to the aforementioned EF1 guarantee in the unweighted setting (Lipton et al. 2004). In fact, the impossibility persists even with $WWEF_c$ for any constant c (Chakraborty et al. 2021, Sec. 8). In light of these observations, Chakraborty et al. left open the direction of identifying appropriate envy-based notions for non-additive valuations. We also remark that relaxing WEF1 using a multiplicative approximation studied in the unweighted setting by, e.g., Amanatidis, Birmpas, and Markakis (2018) and

¹An allocation is called *complete* if it allocates all of the goods.

Plaut and Roughgarden (2020) does not help circumvent this counterexample either.²

Note that the valuations in Example 1 are particularly simple: both agents have *binary submodular* valuations, that is, submodular valuations³ in which the marginal gain from receiving any single good is either 0 or 1. Binary submodular valuations are also known as *matroid-rank valuations*, and have been studied in a number of recent fair division papers, mostly in the unweighted setting (Babaioff, Ezra, and Feige 2021a; Barman and Verma 2021, 2022; Benabbou et al. 2021; Goko et al. 2022; Suksompong and Teh 2023; Viswanathan and Zick 2023a,b).⁴ Such valuations arise in settings such as the allocation of course slots to students, or apartments in public housing estates to ethnic groups (Benabbou et al. 2021). General submodular valuations have likewise received interest among fair division researchers, for example in the context of a (non-envy-based) notion called maximin share fairness (Barman and Krishnamurthy 2020; Ghodsi et al. 2022; Ben Uziah and Feige 2023).

In this paper, we explore weighted envy-freeness for both matroid-rank and general submodular valuations. We propose new envy-based notions and show that they can be satisfied in these settings, not only via extensions of existing algorithms, but also via new rules. For the sake of generality, we define our notions based on the notion $\text{WEF}(x, 1 - x)$ of Chakraborty, Segal-Halevi, and Suksompong (2022). With additive valuations, given a parameter $x \in [0, 1]$, $\text{WEF}(x, 1 - x)$ allows each agent i to subtract x times the value of some good in another agent j 's bundle from i 's value for this bundle, and add $(1 - x)$ times the value of this good to the value of i 's own bundle. WEF1 corresponds to $\text{WEF}(1, 0)$, and higher values of x yield notions that favor lower-weight agents. To obtain more intuition on $\text{WEF}(x, 1 - x)$, consider the following example.

Example 2. Let $n \geq 2$, and suppose that there are n identical goods with value 1 each, and n agents with additive valuations such that $w_1 = \dots = w_{n-1} = 1$ and $w_n = n + 1$. If one wants to ensure that each agent receives nonzero value, the only way is to allocate one good to every agent—this is the only $\text{WEF}(1, 0)$ allocation. However, agent n may reasonably object to this allocation, given that her weight is larger than the weight of all other agents combined. In particular, she may demand that all goods be allocated to her—this allocation is the unique one fulfilling $\text{WEF}(0, 1)$. Hence, $\text{WEF}(x, 1 - x)$ captures the (inevitable) trade-off between satisfying lower-weight agents and higher-weight ones.

Chakraborty, Segal-Halevi, and Suksompong (2022) showed that for any instance with additive valuations and

²Specifically, for any $r \in (0, 1]$, if $w_1 = r/2$, $w_2 = 1$, $m > 2 + 2/r^2$, and the valuations of the two agents are as in Example 1, one can check that there is no r - WEF1 allocation.

³See the definition of submodularity in Section 2.

⁴Exceptions are the recent works of Suksompong and Teh (2023) and Viswanathan and Zick (2023a), which deal with the weighted setting. In fact, Viswanathan and Zick (2023a) pointed out that the “main limitation” of their approach is that it cannot be used to achieve envy-based fairness properties; this is precisely the issue that we address in our paper.

any $x \in [0, 1]$, a complete $\text{WEF}(x, 1 - x)$ allocation always exists; on the other hand, they proved that for any distinct x and x' , there is an instance with binary additive valuations and identical goods in which no complete allocation satisfies both $\text{WEF}(x, 1 - x)$ and $\text{WEF}(x', 1 - x')$.

1.1 Our Contributions

In Section 2, we introduce two new families of weighted envy-freeness notions. The first family, $\text{TWEF}(x, 1 - x)$, is based on the concept of *transferability*:⁵ we consider the condition $\text{TWEF}(x, 1 - x)$ from agent i to agent j to be violated only if the $\text{WEF}(x, 1 - x)$ condition between i and j fails *and* i 's value for her own bundle increases if all goods from j 's bundle are transferred to i 's bundle. $\text{TWEF}(x, 1 - x)$ handles instances such as the one in Example 1, where an agent could be unsatisfied with respect to $\text{WEF}(x, 1 - x)$ even if she already receives her maximum possible utility. Our second family, $\text{WMEF}(x, 1 - x)$, is an extension of the notion *marginal EF1* (MEF1) of Caragiannis et al. (2019) from the unweighted setting. The idea is that, instead of agent i considering her value for agent j 's bundle as in $\text{WEF}(x, 1 - x)$, agent i considers her *marginal* value of j 's bundle when added to i 's own bundle. While $\text{TWEF}(x, 1 - x)$ is stronger than $\text{WMEF}(x, 1 - x)$, we show that the former notion is suitable primarily for matroid-rank valuations, whereas the latter can be guaranteed even for general submodular valuations. Note that when valuations are additive, both $\text{TWEF}(x, 1 - x)$ and $\text{WMEF}(x, 1 - x)$ reduce to $\text{WEF}(x, 1 - x)$, which in turn reduces to EF1 if all agents have equal weights.

In Sections 3 and 4, we allow agents to have arbitrary submodular valuations. In Section 3, we investigate *picking sequences*, which let agents take turns picking a good according to a specified agent ordering until the goods run out. While previous work on picking sequences typically assumed that agents have additive valuations, this assumption may be violated in real-world applications of picking sequences, such as the allocation of ministries to political parties. We adjust picking sequences to submodular valuations by letting agents pick a good with the highest *marginal* gain in each of their turns. We show that for every x , the output of the adjusted version of the picking sequence proposed by Chakraborty, Segal-Halevi, and Suksompong (2022) with parameter x satisfies $\text{WMEF}(x, 1 - x)$; this generalizes their result from the weighted additive setting. As a corollary, in the unweighted submodular setting, the adjusted version of the commonly studied round-robin algorithm produces an MEF1 allocation. In Section 4, we consider the *maximum weighted Nash welfare* (MWNW) rule, which chooses an allocation that maximizes the weighted product of the agents' utilities. Although prior results rule out the possibility for each x that MWNW implies $\text{WMEF}(x, 1 - x)$, we show that an MWNW allocation always satisfies a relaxation of $\text{WMEF}(x, 1 - x)$ called *weak weighted MEF1* (WWMEF1). This extends a corresponding result of Chakraborty et al. (2021) from the weighted additive setting, which in turn gen-

⁵This concept has been discussed by Benabbou et al. (2021) and Chakraborty et al. (2021).

eralizes the prominent result by Caragiannis et al. (2019) in the unweighted additive setting.

Next, in Sections 5 and 6, we focus on agents with matroid-rank valuations—as we discussed earlier, this class of valuations has been studied in several recent papers. In Section 5, we extend the “transfer algorithm” of Benabbou et al. (2021) from the unweighted setting, and prove that our algorithm returns a clean⁶ TWEF($x, 1 - x$) allocation that maximizes the unweighted utilitarian welfare. While Benabbou et al.’s potential function argument can be generalized to show that our algorithm terminates, it is insufficient for establishing polynomial-time termination in our setting with different weights; hence, we devise a more elaborate argument for this purpose. Finally, in Section 6, we introduce new welfare measures based on harmonic numbers and their variants.⁷ Perhaps surprisingly, we demonstrate that under matroid-rank valuations, the maximum-welfare rules based on our measures offer stronger fairness guarantees than MWNW. In particular, while MWNW does not imply WEF($x, 1 - x$) for any x even with binary additive valuations and identical goods (Chakraborty, Segal-Halevi, and Suksompong 2022), we prove that a clean *maximum weighted harmonic welfare* allocation parameterized by x satisfies TWEF($x, 1 - x$) for matroid-rank valuations (and therefore WEF($x, 1 - x$) for binary additive valuations).⁸

2 Preliminaries

Let $N = [n]$ be the set of agents and $G = \{g_1, \dots, g_m\}$ be the set of indivisible goods, where $[k] := \{1, \dots, k\}$ for any positive integer k . A *bundle* refers to a subset of G . Each agent $i \in N$ has a *weight* $w_i > 0$ representing her entitlement, and a *valuation function* (or *utility function*) $v_i : 2^G \rightarrow \mathbb{R}_{\geq 0}$. The setting where all of the weights are equal is sometimes referred to as the *unweighted setting*. For convenience, we write $v_i(g)$ instead of $v_i(\{g\})$ for a single good g . We assume throughout the paper that v_i is

- *monotone*: $v_i(G') \leq v_i(G'')$ for all $G' \subseteq G'' \subseteq G$;
- *submodular*: $v_i(G' \cup \{g\}) - v_i(G') \geq v_i(G'' \cup \{g\}) - v_i(G'')$ for all $G' \subseteq G'' \subseteq G$ and $g \in G \setminus G''$;
- *normalized*: $v_i(\emptyset) = 0$.

The function v_i is said to be *matroid-rank* (or *binary submodular*) if it is submodular and $v_i(G' \cup \{g\}) - v_i(G') \in \{0, 1\}$ for all $G' \subseteq G$ and $g \in G \setminus G'$. Moreover, v_i is *additive* if $v_i(G') = \sum_{g \in G'} v_i(g)$ for all $G' \subseteq G$, and

⁶An allocation is *clean* if no good can be discarded from an agent’s bundle without decreasing the agent’s utility (Benabbou et al. 2021). The term *non-redundant* has also been used with the same meaning (Babaioff, Ezra, and Feige 2021a).

⁷The harmonic welfare measure is the basis of the *proportional approval voting* (PAV) rule in the context of *approval-based committee voting* (see, e.g., the book by Lackner and Skowron (2023)). To the best of our knowledge, we are the first to consider this measure in the context of fair division.

⁸To further exhibit the potential of harmonic welfare, we show in the full version of our paper (Montanari et al. 2022) that, in the unweighted additive setting, if each agent’s value for each good is an integer, then a maximum harmonic welfare allocation always satisfies EF1.

binary additive if it is additive and $v_i(g) \in \{0, 1\}$ for all $g \in G$. An *instance* consists of the set of agents N , the set of goods G , and the agents’ weights $(w_i)_{i \in N}$ and valuation functions $(v_i)_{i \in N}$.

An *allocation* $\mathcal{A}$ is a list of bundles $(A_1, \dots, A_n)$ such that no two bundles overlap, where bundle A_i is assigned to agent i . The allocation is *complete* if $\bigcup_{i \in N} A_i = G$. It is *Pareto-optimal* (PO) if there does not exist another allocation $\mathcal{A}'$ such that $v_i(A'_i) \geq v_i(A_i)$ for all $i \in N$ and the inequality is strict for at least one $i \in N$; such an allocation $\mathcal{A}'$ is said to *Pareto-dominate* $\mathcal{A}$. We denote by $N_A^+ \subseteq N$ the subset of agents receiving positive utility from $\mathcal{A}$. The *unweighted utilitarian welfare* of $\mathcal{A}$ is defined as $\sum_{i \in N} v_i(A_i)$.

For a bundle $G' \subseteq G$, we define the *marginal gain* of a good $g \notin G'$ for agent i as $\Delta_i^+(G', g) := v_i(G' \cup \{g\}) - v_i(G')$. Similarly, the *marginal loss* of a good $g \in G'$ for agent i is defined as $\Delta_i^-(G', g) := v_i(G') - v_i(G' \setminus \{g\})$. An allocation $\mathcal{A}$ is called *clean* (or *non-redundant*) if for any $i \in N$ and any $g \in A_i$, it holds that $\Delta_i^-(A_i, g) > 0$. For matroid-rank valuations, $\mathcal{A}$ is clean if and only if $v_i(A_i) = |A_i|$ for all $i \in N$ (Benabbou et al. 2021, Prop. 3.3). Clean allocations are common in the study of matroid-rank valuations (Babaioff, Ezra, and Feige 2021a; Benabbou et al. 2021; Barman and Verma 2022; Goko et al. 2022; Suksompong and Teh 2023; Viswanathan and Zick 2023a,b). While clean allocations may be incomplete, achieving completeness along with certain properties under matroid-rank valuations can be surprisingly challenging—we refer to the discussion by Benabbou et al. (2021, p. 21).

We now introduce our first family of fairness notions, TWEF(x, y).

Definition 3 (TWEF(x, y)). For $x, y \in [0, 1]$, an allocation $\mathcal{A}$ is said to satisfy *transferable WEF*(x, y) (TWEF(x, y)) if, for each pair of agents $i, j \in N$, either $v_i(A_i) = v_i(A_i \cup A_j)$ or there exists $g \in A_j$ such that

$$\frac{v_i(A_i) + y \cdot \Delta_i^+(A_i, g)}{w_i} \geq \frac{v_i(A_j) - x \cdot \Delta_i^-(A_j, g)}{w_j}.$$

By submodularity and monotonicity, the condition $v_i(A_i) = v_i(A_i \cup A_j)$ is equivalent to the requirement that $v_i(A_i) = v_i(A_i \cup \{g\})$ for every $g \in A_j$.

For any x and y , if valuations are additive, then TWEF(x, y) reduces to the notion WEF(x, y) of Chakraborty, Segal-Halevi, and Suksompong (2022). Like Chakraborty et al., we will mostly be concerned with the case where $y = 1 - x$. As we will see, TWEF($x, 1 - x$) is a useful notion for matroid-rank valuations. However, like WEF($x, 1 - x$), it can be too demanding for general submodular valuations. For instance, in Example 1, if agent 2 has value $1 + (|G'| - 1) \cdot \varepsilon$ for any nonempty bundle G' , where $\varepsilon > 0$ is a small constant, then the condition $v_i(A_i) = v_i(A_i \cup A_j)$ becomes impotent and a complete TWEF($x, 1 - x$) allocation does not exist for any x . The second family of notions that we propose, which is based on the *marginal EF1* (MEF1) notion of Caragiannis et al. (2019),⁹ does not suffer from this shortcoming.

⁹In the unweighted setting, an allocation satisfies MEF1 if for

Definition 4 ($WMEF(x, y)$). For $x, y \in [0, 1]$, an allocation $\mathcal{A}$ is said to satisfy $WMEF(x, y)$ if, for each pair of agents $i, j \in N$, either $A_j = \emptyset$ or there exists $g \in A_j$ such that

$$\frac{v_i(A_i) + y \cdot \Delta_i^+(A_i, g)}{w_i} \geq \frac{v_i(A_i \cup A_j) - v_i(A_i) - x \cdot \Delta_i^-(A_i \cup A_j, g)}{w_j}.$$

If valuations are additive, $WMEF(x, y)$ reduces to $WEF(x, y)$ for any x and y . On the other hand, if all agents have the same weight, $WMEF(x, 1 - x)$ reduces to MEF1 only if $x = 1$. The following proposition, whose proof can be found in the full version of our paper (Montanari et al. 2022), establishes an implication relationship between our two families of notions.

Proposition 5. For all $x, y \in [0, 1]$, every $TWEF(x, y)$ allocation is also $WMEF(x, y)$.

Since the valuations that we consider in this paper are not necessarily additive, in order to reason about the running time of algorithms, we make the standard assumption that an algorithm can query the value of any agent i for any bundle $G' \subseteq G$ in constant time.

3 Picking Sequences

In this section, we investigate *picking sequences*, which are procedures wherein agents take turns picking a good according to a specified agent ordering until there are no more goods left. For brevity, we will say that a picking sequence satisfies a fairness notion if the allocation that it returns always satisfies that notion.

With additive valuations, Chakraborty, Segal-Halevi, and Suksompong (2022) showed that for each $x \in [0, 1]$, a picking sequence that assigns each subsequent pick to an agent $i \in N$ with the smallest ratio $\frac{t_i + (1-x)}{w_i}$, where t_i denotes the number of times agent i has picked so far, satisfies $WEF(x, 1 - x)$. Our main result of this section extends their result to submodular valuations. We make the specification that, in each turn, the agent picks a good that yields the highest *marginal gain* with respect to the agent's current bundle, breaking ties arbitrarily. More formally, if it is agent i 's turn, then i chooses a good g that maximizes $\Delta_i^+(A_i, g)$, where A_i is the set of goods that i picked in previous turns.

Theorem 6. Let $x \in [0, 1]$. Consider a picking sequence π_x such that, in each turn, the pick is assigned to an agent $i \in N$ with the smallest ratio $\frac{t_i + (1-x)}{w_i}$, where t_i denotes the number of times agent i has picked so far, and the agent picks a good that yields the highest marginal gain. Then, under submodular valuations, π_x satisfies $WMEF(x, 1 - x)$.

For any x and agents with equal weights, π_x encompasses the popular *round-robin algorithm* where the agents take turns in the order $1, 2, \dots, n, 1, 2, \dots, n, 1, 2, \dots$, and $WMEF(1, 0)$ reduces to MEF1 of Caragiannis et al. (2019). We therefore have the following corollary in the unweighted setting, which is also new to the best of our knowledge.

all $i, j \in N$, either $A_j = \emptyset$ or there exists $g \in A_j$ such that $v_i(A_i) \geq v_i(A_i \cup A_j \setminus \{g\}) - v_i(A_i)$.

Corollary 7. Assume that all agents have equal weights and submodular valuations. Suppose that in each turn of the round-robin algorithm, the picking agent picks a good with the highest marginal gain. Then, the algorithm returns a complete MEF1 allocation.

As Corollary 7 admits a more direct proof, which also illustrates the ideas that we use to show Theorem 6, we present the proof of Corollary 7 here and leave the proof of Theorem 6 to the full version of our paper (Montanari et al. 2022).

Proof of Corollary 7. Let $\mathcal{A}$ be the allocation produced by the round-robin algorithm, and consider any $i, j \in N$. Assume without loss of generality that $i < j$.

We first establish the MEF1 condition from i toward j . Let $k := |A_j| \leq |A_i|$, and suppose that agent j picks the goods in the order $c_1, c_2, \dots, c_k$. Let $b_1, b_2, \dots, b_k$ be the first k goods picked by agent i in this order. For $0 \leq \ell \leq k$, let $B_\ell = \{b_1, \dots, b_\ell\}$ and $C_\ell = \{c_1, \dots, c_\ell\}$ (so $B_0 = C_0 = \emptyset$). For $1 \leq \ell \leq k$, since agent i picks b_ℓ when c_ℓ is also available, it must be that

$$v_i(B_\ell) - v_i(B_{\ell-1}) \geq v_i(B_{\ell-1} \cup \{c_\ell\}) - v_i(B_{\ell-1}).$$

Moreover, since $B_{\ell-1} \subseteq A_i \subseteq A_i \cup C_{\ell-1}$, submodularity implies that

$$\begin{aligned} v_i(B_{\ell-1} \cup \{c_\ell\}) - v_i(B_{\ell-1}) \\ \geq v_i(A_i \cup C_{\ell-1} \cup \{c_\ell\}) - v_i(A_i \cup C_{\ell-1}). \end{aligned}$$

Combining the previous two inequalities yields

$$\begin{aligned} v_i(B_\ell) - v_i(B_{\ell-1}) \\ \geq v_i(A_i \cup C_{\ell-1} \cup \{c_\ell\}) - v_i(A_i \cup C_{\ell-1}). \end{aligned}$$

Summing this over all $\ell \in [k]$, we get $v_i(B_k) \geq v_i(A_i \cup C_k) - v_i(A_i)$. Since $B_k \subseteq A_i$ and $C_k = A_j$, it follows that $v_i(A_i) \geq v_i(A_i \cup A_j) - v_i(A_i)$, and the MEF1 condition from i to j is fulfilled.

The proof for the MEF1 condition from j toward i is almost identical: by ignoring the first good g picked by agent i and applying the same argument as before, we arrive at $v_j(A_j) \geq v_j(A_j \cup (A_i \setminus \{g\})) - v_j(A_j)$. Thus, the MEF1 condition is again satisfied. $\square$

In the full version, we provide an example showing that the condition MEF1 in Corollary 7 cannot be strengthened to EF1, even when agents have matroid-rank valuations.

4 Nash Welfare

In this section, we turn our attention to maximum weighted Nash welfare (MWNW), a weighted extension of the well-studied maximum Nash welfare (MNW). MWNW has been examined in several recent papers (Chakraborty et al. 2021; Garg, Husić, and Végh 2021; Garg et al. 2022; Suksompong and Teh 2022; Viswanathan and Zick 2023a).

Definition 8 (MWNW). Given an instance, an allocation $\mathcal{A}$ is a *maximum weighted Nash welfare (MWNW)* allocation if it maximizes the *weighted Nash welfare* $WNW(\mathcal{A}) := \prod_{i \in N} v_i(A_i)^{w_i}$. If the highest possible weighted Nash welfare is 0, an MWNW allocation should maximize the number of agents receiving positive utility and, subject to that, maximize the weighted Nash welfare of these agents.

Chakraborty, Segal-Halevi, and Suksompong (2022) showed that, for each $x \in [0, 1]$, there exists an instance with binary additive valuations and identical goods such that every MWNW allocation is not $\text{WEF}(x, 1 - x)$. As a consequence, MWNW allocations cannot always satisfy $\text{WMEF}(x, 1 - x)$ for submodular valuations. On the other hand, Chakraborty et al. (2021) proved that, under additive valuations, MWNW allocations satisfy *weak WEF1* (*WWEF1*), which is weaker than $\text{WEF}(x, 1 - x)$ for every x but still reduces to EF1 in the unweighted additive setting. We extend their result to the weighted submodular setting via a natural generalization of *WWEF1*.

Definition 9 (*WWEF1*). An allocation $\mathcal{A}$ is said to satisfy *weak weighted marginal envy-freeness up to one good* (*WWEF1*) if for each pair of agents i, j with $A_j \neq \emptyset$, there exists a good $g \in A_j$ such that

$$\begin{aligned} \text{either } \frac{v_i(A_i)}{w_i} &\geq \frac{v_i(A_i \cup A_j \setminus \{g\}) - v_i(A_i)}{w_j} \\ \text{or } \frac{v_i(A_i \cup \{g\})}{w_i} &\geq \frac{v_i(A_i \cup A_j) - v_i(A_i)}{w_j}. \end{aligned}$$

Theorem 10. *Under submodular valuations, every MWNW allocation satisfies WWEF1 and PO.*

The proof of Theorem 10 can be found in the full version of our paper (Montanari et al. 2022).

Viswanathan and Zick (2023a) showed that if agents have matroid-rank valuations, an MWNW allocation can be found in polynomial time. On the other hand, with equal-weight agents and additive valuations, even approximating the maximum Nash welfare is computationally difficult (Lee 2017).

5 Transfer Algorithm

For agents with equal weights and matroid-rank valuations, Benabbou et al. (2021, Algorithm 1) proposed a “transfer algorithm” that computes a clean, utilitarian welfare-maximizing EF1 allocation in polynomial time. In this section, we extend their algorithm to the weighted setting. Our algorithm is presented as Algorithm 1 below; we argue in the proof of Theorem 11 that the algorithm is well-defined.

Algorithm 1: For finding a clean $\text{TWEF}(x, 1 - x)$ allocation maximizing $\sum_{i \in N} v_i(A_i)$

Compute a clean allocation $\mathcal{A}$ that maximizes the unweighted utilitarian welfare.

while there exist $i, j \in N$ such that $\text{TWEF}(x, 1 - x)$ from i to j fails with respect to $\mathcal{A}$ **do**

Find a good $g \in A_j$ with $\Delta_i^+(A_i, g) = 1$.

$A_i \leftarrow A_i \cup \{g\}; A_j \leftarrow A_j \setminus \{g\}$.

end while

Theorem 11. *Suppose that all agents have matroid-rank valuations, and let $x \in [0, 1]$. Algorithm 1 with parameter x returns a clean $\text{TWEF}(x, 1 - x)$ (and therefore $\text{WMEF}(x, 1 - x)$) allocation that maximizes the unweighted utilitarian welfare among all allocations in polynomial time.*

Since any allocation maximizing the unweighted utilitarian welfare is PO, the allocation output by Algorithm 1 is also PO. In the unweighted setting, Benabbou et al. (2021) exhibited polynomial-time termination of their algorithm using the potential function $\Phi(\mathcal{A}) := \sum_{i \in N} v_i(A_i)^2$. As $\Phi(\mathcal{A})$ is always an integer between 0 and m^2 and decreases with every transfer, the number of transfers made by their algorithm is at most m^2 . While we can also establish termination of our weighted algorithm by modifying the potential function to $\Phi(\mathcal{A}) = \sum_{i \in N} \frac{v_i(A_i)^2 + (1-2x) \cdot v_i(A_i)}{w_i}$, this argument does not yield a polynomial upper bound on the number of transfers, because the potential function may decrease by a very small amount depending on the weights. Therefore, we will instead employ a different, more refined, argument to show that our algorithm terminates in polynomial time as well.

Proof of Theorem 11. By Proposition 5, it suffices to prove the statement for $\text{TWEF}(x, 1 - x)$.

First, we claim that each transfer maintains the welfare optimality and cleanness of the allocation. Indeed, $v_j(A_j)$ decreases by 1 because the previous allocation is clean, while $v_i(A_i)$ increases by 1 due to the algorithm’s choice of the good $g \in A_j$. Hence, $\sum_{k \in N} v_k(A_k)$ remains the same. Moreover, since $v_k(A_k) = |A_k|$ for all $k \in N$, the allocation remains clean.

If the $\text{TWEF}(x, 1 - x)$ condition from agent i to agent j fails at some point during the execution of the algorithm, it must be that $v_i(A_i) < v_i(A_i \cup A_j)$ and for every $g \in A_j$ we have

$$\begin{aligned} \frac{v_i(A_i) + (1-x) \cdot \Delta_i^+(A_i, g)}{w_i} &< \frac{v_i(A_j) - x \cdot \Delta_i^-(A_j, g)}{w_j} \\ &= \frac{v_i(A_j \setminus \{g\}) + (1-x) \cdot \Delta_i^-(A_j, g)}{w_j} \\ &\leq \frac{v_j(A_j \setminus \{g\}) + (1-x) \cdot \Delta_j^-(A_j, g)}{w_j}, \end{aligned} \quad (1)$$

where the latter inequality follows from cleanness. Since $v_i(A_i) < v_i(A_i \cup A_j)$, by submodularity, there exists $g^* \in A_j$ such that $\Delta_i^+(A_i, g^*) = 1$; in particular, the algorithm is well-defined. Plugging this good g^* into (1) and using the cleanness of $\mathcal{A}$, we get

$$\frac{|A_i| + (1-x)}{w_i} < \frac{|A_j| - x}{w_j}. \quad (2)$$

If the algorithm terminates, then $\text{TWEF}(x, 1 - x)$ is satisfied for all pairs of agents i, j . We will show that the algorithm always terminates, and moreover does so in polynomial time. The initial clean allocation $\mathcal{A}$ can be found in polynomial time (Benabbou et al. 2021). Checking whether $\text{TWEF}(x, 1 - x)$ fails for some pair i, j (and, if so, finding a valid transfer) can be done in polynomial time. It therefore remains to argue that the number of transfers is polynomial. For ease of understanding, we will formulate this argument in terms of cupboards and balls.

Associate each $k \in N$ with a cupboard consisting of m shelves at height $\frac{1-x}{w_k}, \frac{2-x}{w_k}, \dots, \frac{m-x}{w_k}$, respectively. For the

clean allocation $\mathcal{A}$ at the beginning of the algorithm and each $k \in N$, place one ball on each of the $|A_k|$ lowest shelves¹⁰ of cupboard k . Whenever a good is transferred from A_j to A_i , move the highest ball in cupboard j to the lowest shelf without a ball in cupboard i . This means that the ball is moved from height $\frac{|A_j|-x}{w_j}$ to height $\frac{|A_i|+1-x}{w_i}$; by (2), the height of the ball decreases. Since there are m balls and at most mn heights of the shelves, the number of transfers is at most m^2n , which is indeed polynomial.¹¹ This concludes the proof. $\square$

6 Harmonic Welfare

Recall from Section 4 that an MWNW allocation maximizes the product $\prod_{i \in N} v_i(A_i)^{w_i}$, or equivalently, the sum $\sum_{i \in N} w_i \cdot \ln v_i(A_i)$. Since $\ln k$ is approximately the k -th harmonic number $H_k := 1 + \frac{1}{2} + \dots + \frac{1}{k}$ for each positive integer k , one could also consider a *maximum weighted harmonic welfare (MWHW)* allocation, defined as an allocation that maximizes the sum $\sum_{i \in N} w_i \cdot H_{v_i(A_i)}$, where $H_0 = 0$. Interestingly, we show in this section that for agents with matroid-rank valuations, MWHW outperforms MWNW in terms of fairness. Specifically, even though for each $x \in [0, 1]$ there exists an instance with binary additive valuations and identical goods in which every MWNW allocation fails $\text{WEF}(x, 1-x)$ (Chakraborty, Segal-Halevi, and Suksompong 2022), we show that a clean MWHW allocation satisfies $\text{TWEF}(0, 1)$ for matroid-rank valuations (and therefore $\text{WEF}(0, 1)$ for binary additive valuations). More generally, we define a class of modified harmonic numbers parameterized by x such that a clean maximum-welfare allocation based on each x satisfies $\text{TWEF}(x, 1-x)$.

Definition 12 (Modified harmonic numbers). Let $k \in \mathbb{Z}_{\geq 0}$. For $x \in [0, 1]$, the number $H_{k,x}$ is defined by

$$H_{k,x} = \begin{cases} \frac{1}{1-x} + \frac{1}{2-x} + \dots + \frac{1}{k-x} & \text{if } k \geq 1; \\ 0 & \text{if } k = 0, \end{cases}$$

whereas for $x = 1$, $H_{k,x}$ is defined by

$$H_{k,1} = \begin{cases} 1 + \frac{1}{2} + \dots + \frac{1}{k-1} & \text{if } k \geq 2; \\ 0 & \text{if } k = 1; \\ -\infty & \text{if } k = 0. \end{cases}$$

Note that the numbers $H_{k,0}$ correspond to the canonical harmonic numbers H_k , and for each x the sequence $H_{0,x}, H_{1,x}, \dots$ is increasing. We define a maximum weighted harmonic welfare allocation parameterized by x . Recall that $N_{\mathcal{A}}^+$ denotes the set of agents who receive positive utility from the allocation $\mathcal{A}$.

¹⁰The sum of the heights of all balls is $\sum_{i \in N} \frac{|A_i|^2 + (1-2x) \cdot |A_i|}{2w_i}$, which is exactly half of the potential function mentioned before the proof.

¹¹Note that if all agents have equal weights, the number of different shelf heights is only m . The number of transfers is then bounded by m^2 , which matches the bound provided by Benabbou et al. (2021).

Definition 13 (MWHW_x). For $x \in [0, 1]$, given an instance with matroid-rank valuations, an allocation $\mathcal{A}$ is an MWHW_x allocation if it maximizes the sum $\text{WHW}_x(\mathcal{A}) := \sum_{i \in N} w_i \cdot H_{v_i(A_i),x}$.

For $x = 1$, $\mathcal{A}$ is an MWHW_1 allocation if it maximizes the number of agents receiving positive utility and, subject to that, maximizes the sum $\sum_{i \in N_{\mathcal{A}}^+} w_i \cdot H_{v_i(A_i),1}$.

The quantity $H_{v_i(A_i),x}$ is well-defined because, for matroid-rank valuations, $v_i(A_i)$ is always a non-negative integer. We now prove the efficiency and fairness guarantees of MWHW_x allocations, starting with efficiency.

Theorem 14. Let $x \in [0, 1]$. Under matroid-rank valuations, every MWHW_x allocation is PO.

Proof. Let $\mathcal{A}$ be an MWHW_x allocation. For $x < 1$, if $\mathcal{A}$ is Pareto-dominated by another allocation $\mathcal{A}'$, then $\sum_{i \in N} w_i \cdot H_{v_i(A'_i),x} > \sum_{i \in N} w_i \cdot H_{v_i(A_i),x}$, a contradiction.

Consider now the case $x = 1$. If $\mathcal{A}$ were not PO, there would exist an allocation $\hat{\mathcal{A}}$ such that $v_j(\hat{A}_j) > v_j(A_j)$ for some $j \in N$ and $v_i(\hat{A}_i) \geq v_i(A_i)$ for every $i \in N \setminus \{j\}$. If $j \in N \setminus N_{\mathcal{A}}^+$, we would have $v_i(\hat{A}_i) > 0$ for every $i \in N_{\mathcal{A}}^+ \cup \{j\}$, contradicting the assumption that $N_{\mathcal{A}}^+$ is the largest subset of agents to whom it is possible to give positive utility simultaneously. On the other hand, if $j \in N_{\mathcal{A}}^+$, we would have $\sum_{i \in N_{\mathcal{A}}^+} w_i \cdot H_{v_i(\hat{A}_i),1} > \sum_{i \in N_{\mathcal{A}}^+} w_i \cdot H_{v_i(A_i),1}$, again a contradiction. Therefore, $\mathcal{A}$ is PO. $\square$

For the fairness guarantee, we will make an assumption that the allocation is clean; we shall demonstrate later that this assumption is necessary. We also remark that given any MWHW_x allocation, one can easily obtain a clean MWHW_x allocation in which every agent receives the same utility as before by iteratively removing any good that does not contribute to its owner's utility until no such good exists.

Theorem 15. Let $x \in [0, 1]$. Under matroid-rank valuations, every clean MWHW_x allocation satisfies $\text{TWEF}(x, 1-x)$ (and therefore $\text{WMEF}(x, 1-x)$).

Proof. By Proposition 5, it suffices to prove the statement for $\text{TWEF}(x, 1-x)$.

Let $\mathcal{A}$ be a clean MWHW_x allocation. Assume for contradiction that for some $i, j \in N$, the $\text{TWEF}(x, 1-x)$ condition from i to j is violated. This means that $v_i(A_i) < v_i(A_i \cup A_j)$ and for every $g \in A_j$ it holds that

$$\frac{v_i(A_i) + (1-x) \cdot \Delta_i^+(A_i, g)}{w_i} < \frac{v_i(A_j) - x \cdot \Delta_i^-(A_j, g)}{w_j}.$$

By the same argument as in the proof of Theorem 11, this implies that

$$\frac{v_i(A_i) + (1-x)}{w_i} < \frac{v_j(A_j) - x}{w_j}. \quad (3)$$

Also, since $v_i(A_i) < v_i(A_i \cup A_j)$, submodularity implies that there exists a good $g^* \in A_j$ such that $\Delta_i^+(A_i, g^*) = 1$.

We now consider two cases depending on whether $x = 1$.

Case 1: $0 \leq x < 1$. If we transfer g^* from A_j to A_i , we obtain an allocation $\mathcal{A}'$ in which $v_i(A'_i) = v_i(A_i) + 1$, $v_j(A'_j) = v_j(A_j) - 1$, and $v_k(A'_k) = v_k(A_k)$ for all $k \in N \setminus \{i, j\}$. Since $\mathcal{A}$ is an MWHW_x allocation, it must be that

$$\begin{aligned} w_i \cdot H_{v_i(A_i), x} + w_j \cdot H_{v_j(A_j), x} \\ \geq w_i \cdot H_{v_i(A_i)+1, x} + w_j \cdot H_{v_j(A_j)-1, x}. \end{aligned}$$

This is equivalent to

$$w_j \cdot \frac{1}{v_j(A_j) - x} - w_i \cdot \frac{1}{v_i(A_i) + 1 - x} \geq 0.$$

Algebraic manipulation gives us

$$\frac{v_i(A_i) + 1 - x}{w_i} \geq \frac{v_j(A_j) - x}{w_j},$$

which contradicts (3).

Case 2: $x = 1$. From (3), we have that

$$\frac{v_i(A_i)}{w_i} < \frac{v_j(A_j) - 1}{w_j}. \quad (4)$$

Since $v_i(A_i) \geq 0$ and $v_j(A_j)$ is an integer, it must be that $v_j(A_j) \geq 2$. If $v_i(A_i) = 0$, we can transfer g^* from A_j to A_i and increase the number of agents with positive utility, contradicting the assumption that $\mathcal{A}$ is an MWHW_1 allocation. Hence, $v_i(A_i) \geq 1$.

The rest of the argument proceeds in a similar way as in Case 1. If we transfer g^* from A_j to A_i , we obtain an allocation $\mathcal{A}'$ in which $v_i(A'_i) = v_i(A_i) + 1$, $v_j(A'_j) = v_j(A_j) - 1$, and $v_k(A'_k) = v_k(A_k)$ for all $k \in N \setminus \{i, j\}$. Note that the number of agents with positive utility is the same in $\mathcal{A}$ and $\mathcal{A}'$. Since $\mathcal{A}$ is an MWHW_1 allocation, it must be that

$$\begin{aligned} w_i \cdot H_{v_i(A_i), 1} + w_j \cdot H_{v_j(A_j), 1} \\ \geq w_i \cdot H_{v_i(A_i)+1, 1} + w_j \cdot H_{v_j(A_j)-1, 1}. \end{aligned}$$

This is equivalent to

$$w_j \cdot \frac{1}{v_j(A_j) - 1} - w_i \cdot \frac{1}{v_i(A_i)} \geq 0.$$

Algebraic manipulation gives us

$$\frac{v_i(A_i)}{w_i} \geq \frac{v_j(A_j) - 1}{w_j},$$

which contradicts (4). $\square$

We now exhibit the necessity of the cleanness condition in Theorem 15.

Proposition 16. *There exists an instance and an allocation such that, for every $x \in [0, 1]$, the allocation is MWHW_x but does not satisfy $\text{TWEF}(x, 1 - x)$.*

Proof. Consider an instance with $n = 2$ agents whose weights are $w_1 = 1$ and $w_2 = 2$, and $m = 6$ goods. Agent 1 has an additive valuation with value 1 for g_1 and 0 for the remaining goods. Agent 2's valuation v_2 is given by

$$v_2(S) = \begin{cases} \min\{3, |S|\} & \text{if } g_1 \notin S; \\ \min\{4, |S|\} & \text{if } g_1 \in S, \end{cases}$$

for each bundle $S \subseteq G$. One can check that v_2 is matroid-rank; we leave the details to the full version of our paper (Montanari et al. 2022).

Fix $x \in [0, 1]$. If $x = 1$, every MWHW_x allocation must give g_1 to agent 1, which leaves agent 2 with a utility of at most 3. Else, for $x < 1$, the maximum weighted harmonic welfare achievable by giving g_1 to agent 1 is

$$\frac{1}{1-x} + 2 \cdot \left(\frac{1}{1-x} + \frac{1}{2-x} + \frac{1}{3-x} \right),$$

whereas the maximum achievable by giving g_1 to agent 2 is

$$2 \cdot \left(\frac{1}{1-x} + \frac{1}{2-x} + \frac{1}{3-x} + \frac{1}{4-x} \right).$$

Since $\frac{1}{1-x} = \frac{2}{2-2x} > \frac{2}{4-x}$, every MWHW_x allocation must again give g_1 to agent 1. In particular, for every $x \in [0, 1]$, the allocation $\mathcal{A} = (\{g_1, g_2, g_3\}, \{g_4, g_5, g_6\})$ is an MWHW_x allocation.

To finish the proof, we show that $\mathcal{A}$ violates the $\text{TWEF}(x, 1-x)$ condition from agent 2 toward agent 1. Note that $v_2(A_2) = 3 < 4 = v_2(A_2 \cup A_1)$. Moreover, for each $g \in A_1$, it holds that

$$\begin{aligned} \frac{v_2(A_2) + (1-x) \cdot \Delta_2^+(A_2, g)}{w_2} \\ \leq \frac{3 + (1-x)}{2} < 3 - x = \frac{v_2(A_1) - x \cdot \Delta_2^-(A_1, g)}{w_1}. \end{aligned}$$

Hence, the $\text{TWEF}(x, 1-x)$ condition from agent 2 to agent 1 is not satisfied. $\square$

By applying results from the recent work of Viswanathan and Zick (2023a), we show in the full version of our paper (Montanari et al. 2022) that, for each $x \in [0, 1]$, an MWHW_x allocation (which additionally maximizes the unweighted utilitarian welfare across all allocations) can be found in polynomial time.

Finally, we remark that it may be interesting to consider harmonic welfare beyond binary valuations. In the full version of our paper, we prove that for agents with equal weights and additive valuations, if the value of every agent for every good is an integer (in which case the harmonic welfare is well-defined), then an allocation maximizing the harmonic welfare is always EF1.

7 Conclusion

In this paper, we have embarked on a study of weighted envy-freeness for non-additive valuations by focusing on the important class of submodular valuations. We proposed two families of envy-based notions: $\text{TWEF}(x, 1-x)$, which is suitable for matroid-rank (i.e., binary submodular) valuations, and $\text{WMEF}(x, 1-x)$, which is useful even for arbitrary submodular valuations. We demonstrated that our notions can be satisfied via procedures ranging from picking sequences to welfare maximization. To the best of our knowledge, these are the first notions that can always be satisfied in the weighted setting under submodular (or matroid-rank) valuations and moreover reduce to EF1 in the unweighted additive setting. An interesting direction in light of our work is to consider other valuation classes, e.g., supermodular or subadditive valuations.

Acknowledgments

This work was partially supported by the Singapore Ministry of Education under grant number MOE-T2EP20221-0001 and by an NUS Start-up Grant. Ulrike Schmidt-Kraepelin was supported by the *Deutsche Forschungsgemeinschaft* (under grant BR 4744/2-1), the *Centro de Modelamiento Matemático (CMM)* (under grant FB210005, BASAL funds for center of excellence from ANID-Chile) and *ANID-Chile* (grant ACT210005). Moreover, she was supported by the National Science Foundation under Grant No. DMS-1928930 and by the Alfred P. Sloan Foundation under grant G-2021-16778 while being in residence at the Simons Laufer Mathematical Sciences Institute (formerly MSRI) in Berkeley, California, during the Fall 2023 semester. We thank the anonymous reviewers for their valuable feedback.

References

- Amanatidis, G.; Birmpas, G.; and Markakis, E. 2018. Comparing approximate relaxations of envy-freeness. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 42–48.
- Aziz, H.; Moulin, H.; and Sandomirskiy, F. 2020. A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation. *Operations Research Letters*, 48(5): 573–578.
- Babaioff, M.; Ezra, T.; and Feige, U. 2021a. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 5119–5126.
- Babaioff, M.; Ezra, T.; and Feige, U. 2021b. Fair-share allocations for agents with arbitrary entitlements. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 127.
- Babaioff, M.; Nisan, N.; and Talgam-Cohen, I. 2021. Competitive equilibrium with indivisible goods and generic budgets. *Mathematics of Operations Research*, 46(1): 382–403.
- Barman, S.; and Krishnamurthy, S. K. 2020. Approximation algorithms for maximin fair division. *ACM Transactions on Economics and Computation*, 8(1): 5:1–5:28.
- Barman, S.; and Verma, P. 2021. Existence and computation of maximin fair allocations under matroid-rank valuations. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 169–177.
- Barman, S.; and Verma, P. 2022. Truthful and fair mechanisms for matroid-rank valuations. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 4801–4808.
- Ben Uziah, G.; and Feige, U. 2023. On fair allocation of indivisible goods to submodular agents. *CoRR*, abs/2303.12444.
- Benabbou, N.; Chakraborty, M.; Igarashi, A.; and Zick, Y. 2021. Finding fair and efficient allocations for matroid rank valuations. *ACM Transactions on Economics and Computation*, 9(4): 21:1–21:41.
- Brams, S. J.; and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3): 12:1–12:32.
- Chakraborty, M.; Igarashi, A.; Suksompong, W.; and Zick, Y. 2021. Weighted envy-freeness in indivisible item allocation. *ACM Transactions on Economics and Computation*, 9(3): 18:1–18:39.
- Chakraborty, M.; Schmidt-Kraepelin, U.; and Suksompong, W. 2021. Picking sequences and monotonicity in weighted fair division. *Artificial Intelligence*, 301: 103578.
- Chakraborty, M.; Segal-Halevi, E.; and Suksompong, W. 2022. Weighted fairness notions for indivisible items revisited. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 4949–4956.
- Farhadi, A.; Ghodsi, M.; Hajiaghayi, M.; Lahaie, S.; Pennock, D.; Seddighin, M.; Seddighin, S.; and Yami, H. 2019. Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research*, 64: 1–20.
- Garg, J.; Husić, E.; Murhekar, A.; and Végh, L. 2022. Tractable fragments of the maximum Nash welfare problem. In *Proceedings of the 18th Conference on Web and Internet Economics (WINE)*, 362–363.
- Garg, J.; Husić, E.; and Végh, L. A. 2021. Approximating Nash social welfare under Rado valuations. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, 1412–1425.
- Ghodsi, M.; HajiAghayi, M.; Seddighin, M.; Seddighin, S.; and Yami, H. 2022. Fair allocation of indivisible goods: Beyond additive valuations. *Artificial Intelligence*, 303: 103633.
- Goko, H.; Igarashi, A.; Kawase, Y.; Makino, K.; Sumita, H.; Tamura, A.; Yokoi, Y.; and Yokoo, M. 2022. Fair and truthful mechanism with limited subsidy. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 534–542.
- Hoefer, M.; Schmalhofer, M.; and Varricchio, G. 2023. Best of both worlds: Agents with entitlements. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 564–572.
- Lackner, M.; and Skowron, P. 2023. *Multi-Winner Voting with Approval Preferences*. Springer.
- Lee, E. 2017. APX-hardness of maximizing Nash social welfare with indivisible items. *Information Processing Letters*, 122: 17–20.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, 125–131.
- Montanari, L.; Schmidt-Kraepelin, U.; Suksompong, W.; and Teh, N. 2022. Weighted envy-freeness for submodular valuations. *CoRR*, abs/2209.06437.

- Plaut, B.; and Roughgarden, T. 2020. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2): 1039–1068.
- Robertson, J.; and Webb, W. 1998. *Cake-Cutting Algorithms: Be Fair if You Can*. Peters/CRC Press.
- Scarlett, J.; Teh, N.; and Zick, Y. 2023. For one and all: Individual and group fairness in the allocation of indivisible goods. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2466–2468.
- Steinhaus, H. 1948. The problem of fair division. *Econometrica*, 16(1): 101–104.
- Suksompong, W.; and Teh, N. 2022. On maximum weighted Nash welfare for binary valuations. *Mathematical Social Sciences*, 117: 101–108.
- Suksompong, W.; and Teh, N. 2023. Weighted fair division with matroid-rank valuations: Monotonicity and strategyproofness. *Mathematical Social Sciences*, 126: 48–59.
- Viswanathan, V.; and Zick, Y. 2023a. A general framework for fair allocation with matroid rank valuations. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, 1129–1152.
- Viswanathan, V.; and Zick, Y. 2023b. Yankee Swap: a fast and simple fair allocation mechanism for matroid rank valuations. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 179–187.