Weighted Envy-Freeness in Indivisible Item Allocation

Mithun Chakraborty*
University of Michigan
Ann Arbor, Michigan, USA
dcsmc@umich.edu

Ayumi Igarashi
University of Tokyo
Tokyo, Japan
igarashi@mist.i.u-tokyo.ac.jp

Warut Suksompong
University of Oxford
Oxford, United Kingdom
warut.suksompong@cs.ox.ac.uk

Yair Zick
National University of Singapore
Singapore
dcsyaz@nus.edu.sg

ABSTRACT

In this paper, we introduce and analyze new envy-based fairness concepts for agents with *weights* that quantify their entitlements in the allocation of indivisible items. We propose two variants of weighted envy-freeness up to one item (WEF1) – *strong* (where the envy can be eliminated by removing an item from the envied agent's bundle) and *weak* (where the envy can be eliminated either by removing an item as in the strong version or by replicating an item from the envied agent's bundle in the envious agent's bundle). We prove that for additive valuations, an allocation that is both Pareto optimal and strongly WEF1 always exists; however, an allocation that maximizes the weighted Nash social welfare may not be strongly WEF1 but always satisfies the weak version of the property. Moreover, we establish that a generalization of the round-robin picking sequence produces in polynomial time a strongly WEF1 allocation for an arbitrary number of agents; for two agents, we can efficiently achieve both strong WEF1 and Pareto optimality by adapting the classic adjusted winner algorithm. We also explore the connections of WEF1 with approximations to the weighted versions of two other fairness concepts: proportionality and the maximin share guarantee.

ACM Reference Format:
Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. 2020. Weighted Envy-Freeness in Indivisible Item Allocation. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 9 pages.

1 INTRODUCTION

The fair allocation of resources to interested parties is a central issue in economics, with several fairness criteria having been studied in the vast literature on fair division. One of the most fundamental criteria is *envy-freeness*, which requires that all agents find their assigned bundle to be the best among all assigned bundles [18, 27].

Envy-freeness is a compelling notion when all agents have the same entitlement to the resource; indeed, in an envy-free allocation, no agent would rather take the place of another agent with respect to the assigned bundles. However, in many division problems, different agents have varying claims on the resource. For instance, consider a facility that has been jointly funded by Alice and Bob, where Alice's contribution to the facility is twice of Bob's. One could then expect Alice to be envious if she does not value her share at least twice as much as Bob's share when they divide the usage of the facility. Besides this interpretation as the *cost* of participating in the resource allocation "game", the weights may also represent other publicly known and accepted measures of entitlement such as *eligibility* or *merit*. Envy-freeness can be naturally extended to this general setting with weights. When the resource to be allocated is infinitely divisible (e.g., time to use a facility), it is known that a weighted envy-free allocation exists for any set of weights and valuations [24, 28].

In this paper, we study weighted envy-freeness for the ubiquitous situation where the resource consists of *indivisible* items. Envy-free allocations cannot always be attained in such situations even in the absence of weights, thereby prompting the need for approximate notions that can be satisfied. The goal of our work is to extend these notions to the weighted case, and explore their relationship to other important welfare notions such as Pareto optimality.

1.1 Our Contributions

In Section 2, we propose two extensions of the popular fairness concept *envy-freeness up to one item (EF1)* [10, 19] to the weighted setting: (strong) weighted envy-freeness up to one item (WEF1) and weak weighted envy-freeness up to one item (WWEF1). We assume for most of the paper that the agents are endowed with additive valuation functions. In addition to some negative results, we establish in Section 3 that a WEF1 allocation always exists and can be computed efficiently using a weight-based picking sequence — this generalizes a well-known result from the unweighted setting. In Section 4, we show by adapting an algorithm of Barman et al. [5] that the existence of a Pareto optimal and WEF1 allocation is always guaranteed for any number of agents; in particular, for two agents, we exhibit that a weighted variant of the classic adjusted winner procedure allows us to efficiently compute an allocation that is both WEF1 and Pareto optimal. In Section 5, we prove that while an allocation with maximum weighted Nash welfare may fail to satisfy WEF1, such an allocation is both Pareto optimal and WWEF1, thereby generalizing an important result of Caragiannis et al. [11]. We explore the relationship between weighted envy-freeness and the weighted versions of other fairness concepts in Section 6, and show through experiments in Section 7 that envy-freeness is often harder to satisfy in the weighted setting than in the unweighted setting.

*This work was done while the author was at the National University of Singapore.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

Finally, we conclude in Section 8 with some thoughts on extending our ideas and results beyond additive valuation functions.

1.2 Related Work

There is a long line of work on fair division of indivisible items; see, e.g., the survey by Bouveret et al. [8] for an overview. Prior work on the fair allocation of indivisible items to asymmetric agents has tackled fairness concepts that are not based on envy. Farhadi et al. [17] introduce and study *weighted maxmin share (WMMS)* fairness, a generalization of an earlier fairness notion of Budish [10]. Aziz et al. [2] explore WMMS fairness in the allocation of indivisible *chores* — items that, in contrast to goods, are valued negatively by the agents — where agents' weights can be interpreted as their shares of the workload. Babaioff et al. [4] study competitive equilibrium for agents with different budgets. Recently, Aziz et al. [3] propose a polynomial-time algorithm for computing an allocation of a pool of goods and chores that satisfies both Pareto optimality and weighted *proportionality up to one item (PROP1)* for agents with asymmetric weights. Unequal entitlements have also been considered in the context of divisible items with respect to *proportionality* [15, 24, 25].

In addition to expressing the entitlement of individual agents, weights can also be applied to settings where each agent represents a group of individuals [6, 7] — here, the *size* of a group can be used as its weight.[1] Specifically, in the model of Benabbou et al. [7], groups correspond to ethnic groups (namely, the major ethnic groups in Singapore, e.g., Chinese, Malay, and Indian). Maintaining provable fairness guarantees amongst the ethnic groups is highly desirable; in fact, it is one of the major tenets of Singaporean society.

2 PRELIMINARIES

Throughout the paper, we denote by $[r]$ the set $\{1, 2, \ldots, r\}$ for any positive integer r. We are given a set $N = [n]$ of *agents*, and a set $O = \{o_1, \ldots, o_m\}$ of *items* or *goods*. Subsets of O are referred to as *bundles*, and each agent $i \in N$ has a *valuation function* $v_i : 2^O \to \mathbb{R}_{\geq 0}$ over bundles; the valuation function for every $i \in N$ is *normalized* (i.e., $v_i(\emptyset) = 0$) and *monotone* (i.e., $v_i(S) \leq v_i(T)$ whenever $S \subseteq T$). We denote $v_i(\{o\})$ simply by $v_i(o)$ for any $i \in N$ and $o \in O$.

An *allocation* A of the items to the agents is a collection of n disjoint bundles $A_1, \ldots, A_n$ such that $\bigcup_{i \in N} A_i \subseteq O$; the bundle A_i is allocated to agent i and $v_i(A_i)$ is agent i's *realized valuation* under A. Given an allocation A, we denote by A_0 the set $O \backslash (\bigcup_{i \in N} A_i)$, and its elements are referred to as *withheld items*. An allocation A is said to be *complete* if $A_0 = \emptyset$, and *incomplete* otherwise.

In our asymmetric setting, each agent $i \in N$ has a fixed weight $w_i > 0$: these weights regulate how agents value their own allocated bundles relative to those of other agents, and hence bear on the overall (subjective) fairness of an allocation. More precisely, we define the *weighted envy* of agent i towards agent j under an allocation A as $\max\left\{0, \frac{v_i(A_j)}{w_j} - \frac{v_i(A_i)}{w_i}\right\}$. An allocation is *weighted envy-free* (WEF) if no agent has positive weighted envy towards another agent. Weighted envy-freeness reduces to traditional envy-freeness when $w_i = w, \forall i \in N$ for some positive real constant w. Since a complete envy-free allocation may not always exist (see, e.g., [8]), it follows trivially that a complete WEF allocation may not exist in general.

We now state the main definitions of our paper, which naturally extend *envy-freeness up to one item (EF1)* [10, 19] to the weighted setting.

Definition 2.1. An allocation A is said to be *(strongly) weighted envy-free up to one item* (WEF1) if for any pair of agents i, j with $A_j \neq \emptyset$, there exists an item $o \in A_j$ such that

$$\frac{v_i(A_i)}{w_i} \geq \frac{v_i(A_j \backslash \{o\})}{w_j}.$$

More generally, A is said to be *weighted envy-free up to c items* (WEFc) for an integer $c \geq 1$ if for any pair of agents i, j, there exists a subset $S_c \subseteq A_j$ of size at most c such that

$$\frac{v_i(A_i)}{w_i} \geq \frac{v_i(A_j \backslash S_c)}{w_j}.$$

Definition 2.2. An allocation A is said to be *weakly weighted envy-free up to one item* (WWEF1) if for any pair of agents i, j with $A_j \neq \emptyset$, there exists an item $o \in A_j$ such that

$$\text{either } \frac{v_i(A_i)}{w_i} \geq \frac{v_i(A_j \backslash \{o\})}{w_j} \text{ or } \frac{v_i(A_i \cup \{o\})}{w_i} \geq \frac{v_i(A_j)}{w_j}.$$

A valuation function $v : 2^O \to \mathbb{R}_{\geq 0}$ is said to be *additive* if $v(S) = \sum_{o \in S} v(o)$ for every $S \subseteq O$. We will assume additive valuations for most of the paper; this is a very common assumption in the fair division literature. Under this assumption, both WEF1 and WWEF1 reduce to EF1 in the unweighted setting. Moreover, one can check that under additive valuations, an allocation satisfies WWEF1 if and only if for any pair of agents i, j with $A_j \neq \emptyset$, there exists an item $o \in A_j$ such that

$$\frac{v_i(A_i)}{w_i} \geq \frac{v_i(A_j)}{w_j} - \frac{v_i(o)}{\min\{w_i, w_j\}}.$$

The criterion WEF1 can be criticized as being too demanding in certain circumstances, when the weight of the envied agent is much larger than that of the envying agent. To illustrate this, consider a problem instance where agent 1 has an additive valuation function and is indifferent among all items taken individually, e.g., $v_1(o) = 1$ for every $o \in O$. Now, if $w_1 = 1$ and $w_2 = 100$, then eliminating one item from agent 2's bundle reduces agent 1's weighted valuation of this bundle by merely 0.01. As such, we might trigger a substantial adverse effect on the welfare/efficiency of the allocation by aiming to (approximately) eliminate 1's weighted envy towards 2. This line of thinking was our motivation for introducing the weak weighted envy-freeness concept. We also note that WWEF1 can be viewed as a stronger version of what we refer to as *transfer weighted envy-freeness up to one item*: agent i is transfer weighted envy-free up to one item towards agent j under the allocation A if there is an item $o \in A_j$ that would eliminate the weighted envy of i towards j upon being transferred from A_j to A_i, i.e., $v_i(A_i \cup \{o\}) \geq \frac{w_i}{w_j} \cdot v_i(A_j \backslash \{o\})$.

In addition to fairness, we often want our allocation to satisfy an efficiency criterion. One important such criterion is Pareto optimality. An allocation A' is said to *Pareto dominate* an allocation

[1] Note that in this model, each group has a valuation function that represents the overall preference of its members. Other group fairness notions do not assume the existence of such aggregate functions and instead take directly into account the preferences of the individual agents in each group [14, 20, 26].

A if $v_i(A'_i) \geq v_i(A_i)$ for all agents $i \in N$ and $v_j(A'_j) > v_j(A_j)$ for some agent $j \in N$. An allocation is *Pareto optimal* (or PO for short) if it is not Pareto dominated by any other allocation.

Allocations maximizing the *Nash welfare* — defined as $\mathsf{NW}(A) := \prod_{i \in N} v_i(A_i)$ — are known to offer strong guarantees with respect to both fairness and efficiency in the unweighted setting [11]. For our weighted setting, we define a natural extension called *weighted Nash welfare*: $\mathsf{WNW}(A) := \prod_{i \in N} v_i(A_i)^{w_i}$. Since it is possible that the maximum attainable $\mathsf{WNW}(A)$ is 0, we define a *maximum weighted Nash welfare* or MWNW allocation along the lines of [11] as follows: given a problem instance, we find a maximum subset of agents, say $N_{\max} \subseteq N$, to which we can allocate bundles of positive value, and compute an allocation to the agents in $N_{\max}$ that maximizes $\prod_{i \in N_{\max}} v_i(A_i)^{w_i}$. To see why the notion of MWNW makes intuitive sense, consider a setting where agents have a value of 1 for each item they receive; furthermore, assume that the number of items is exactly $\sum_{i=1}^{n} w_i$. In this case, one can verify (using standard calculus) that an allocation maximizing MWNW assigns to agent i exactly w_i items. Indeed, following the interpretation of w_i as the number of members of group i, the expression $v_i(A_i)^{w_i}$ can be thought of as each member of group i deriving the same value from the set A_i; the group's overall Nash welfare is thus $v_i(A_i)^{w_i}$.

We also examine the extent to which (approximate) weighted envy-freeness relates to the weighted versions of two other key fairness notions: *proportionality* and the *maximin share guarantee*.

An allocation A is said to be *weighted proportional* or WPROP if, for every agent $i \in N$, $v_i(A_i) \geq \frac{w_i}{\sum_{j \in N} w_j} v_i(O)$. For additive valuations, an allocation A is *weighted proportional up to c items* or WPROPc if, for every $i \in N$, there exists a subset of items not allocated to i, i.e., $S_c \subseteq O \setminus A_i$, of size at most c such that $v_i(A_i) \geq \frac{w_i}{\sum_{j \in N} w_j} \cdot v_i(O) - v_i(S_c)$; this is a natural extension of the (weighted) PROP1 concept in [3, 13].

Let $\Pi(O)$ denote the collection of all (ordered) n-partitions of the set of items O, or, in other words, the collection of all complete allocations of O to n agents. Then, the *weighted maximin share* [17] of agent i is defined as:

$$\mathsf{WMMS}_i := \max_{(A_1, A_2, \ldots, A_n) \in \Pi(O)} \min_{j \in N} \frac{w_i}{w_j} v_i(A_j).$$

An allocation A is called WMMS if $v_i(A_i) \geq \mathsf{WMMS}_i$ for every $i \in N$; for any approximation ratio $\alpha \in (0, 1]$, A is called α-WMMS if $v_i(A_i) \geq \alpha \cdot \mathsf{WMMS}_i$ for every $i \in N$.

Due to space constraints, some proofs are omitted and deferred to the full version of our paper [12].

3 WEF1 ALLOCATIONS

Although it is known that for arbitrary monotone valuation functions, a complete (unweighted) EF1 allocation always exists and can be efficiently computed assuming polynomial-time oracle access to the valuation functions [19], this fact does not imply the existence of the more general complete WEF1 allocations for arbitrary weights and valuation functions.

As far as envy in the traditional sense is concerned, what an agent actually "envies" is an allocated bundle regardless of who owns that bundle. However, both the subjective valuations of allocated bundles and the relative weights interact in non-trivial ways to determine weighted envy. It is easy to see that weighted envy of i towards j does not imply (traditional) envy of i towards j, and vice versa. A crucial implication is that even if agent i's bundle is replaced with the bundle of an agent j towards whom i has weighted envy, i's realized valuation (and hence the ratio of her realized valuation to her weight) may decrease as a result. Indeed, consider a problem instance with $N = [2]$, $O = \{o_1, o_2, o_3\}$; weights $w_1 = 3$ and $w_2 = 1$; and identical, additive valuation functions such that $v_i(o) = 1$, $\forall i \in N$, $\forall o \in O$. Under the complete allocation with $A_1 = \{o_1, o_2\}$, agent 1 has weighted envy towards agent 2 since $v_1(A_2)/w_2 = 1/1 = 1 > 2/3 = v_1(A_1)/w_1$, although agent 1 would not prefer to replace A_1 with A_2 since that reduces her realized valuation from 2 to 1. On the other hand, agent 2 could benefit from replacing A_2 with A_1 even though she does not have weighted envy towards agent 1. As such, the natural extension of Lipton et al. [19]'s seminal envy cycle elimination algorithm does not guarantee a complete WEF1 allocation except in special cases. One such special case is when the agents all have identical valuations.

Proposition 3.1. *The weighted version of Lipton et al.'s envy cycle elimination algorithm (where an edge exists from agent i to agent j if and only if i has weighted envy towards j) produces a complete* WEF1 *allocation whenever agents have identical (not necessarily additive) valuations, i.e.,* $v_i(S) = v(S)$ *for some* $v : 2^O \to \mathbb{R}_{\geq 0}$, $\forall i \in N$, $\forall S \subseteq O$.

Proof. By construction of the algorithm [19], the (incomplete) allocation at the end of each iteration is guaranteed to be WEF1 as long as we can find an agent, say i, towards whom no other agent has weighted envy at the beginning of the iteration: we give the item under consideration to agent i and thus any resulting weighted envy towards i can be eliminated by removing this item. If there is no unenvied agent, then the weighted envy graph consists of at least one cycle; however, under identical valuations, the envy graph cannot have cycles. Indeed, suppose that agents $1, 2, \ldots, \ell$ form a cycle (in that order) for some $\ell \in [n]$. Since agents have identical valuations, it must be that $v(A_1)/w_1 < v(A_2)/w_2 < \cdots < v(A_\ell)/w_\ell < v(A_1)/w_1$, a contradiction. □

Unfortunately, the positive results for utilizing the envy cycle elimination algorithm end with Proposition 3.1.

Proposition 3.2. *If agents do not have identical valuation functions, then the weighted version of Lipton et al. [19]'s envy cycle elimination algorithm may not produce a complete* WEF1 *allocation, even in a problem instance with two agents and additive valuations.*

3.1 Picking Sequence Protocols

When all agents have equal weight and additive valuations, it is well-known that a round-robin algorithm, wherein the agents take turns picking an item, produces an EF1 allocation. This is in fact easy to see: If agent i is ahead of agent j in the ordering, then in every "round", i picks an item that she likes at least as much as j's pick; by additivity, i does not envy j. On the other hand, if agent i picks after agent j, then by considering the first round to begin at i's first pick, it follows that i does not envy j up to the first item that j picks.

We show next that in the general setting with weights, we can construct a weight-dependent picking sequence which guarantees

WEF1 for any number of agents and arbitrary weights. However, unlike in the unweighted case, the proof is much less straightforward and requires making several intricate arguments.

Theorem 3.3. *For any number of agents with additive valuations and arbitrary positive real weights, a complete* WEF1 *allocation always exists and can be computed in polynomial time.*

Algorithm 1 Pick the Least Weight-Adjusted Frequent Picker

Remaining items $\widehat{O} \leftarrow O$.
Bundles $A_i \leftarrow \emptyset, \forall i \in N$.
$t_i \leftarrow 0, \forall i \in N$. /***number of times each agent has picked so far***/
while $\widehat{O} \neq \emptyset$ **do**
  $i^* \leftarrow \arg\min_{i \in N} \frac{t_i}{w_i}$, breaking ties lexicographically.
  $o^* \leftarrow \arg\max_{o \in \widehat{O}} v_{i^*}(o)$, breaking ties arbitrarily.
  $A_{i^*} \leftarrow A_{i^*} \cup \{o^*\}$.
  $\widehat{O} \leftarrow \widehat{O} \backslash \{o^*\}$.
  $t_{i^*} \leftarrow t_{i^*} + 1$.
end while

Proof. Our proof is constructive: we construct a picking sequence such that at each turn, an agent with the lowest weight-adjusted picking frequency picks the next item (Algorithm 1). We claim that after the allocation of each item, for any agent i, every other agent is weighted envy-free towards i up to the item that i picked first.

To this end, first note that choosing an agent who has had the minimum (weight-adjusted) number of picks thus far ensures that the first n picks are a round-robin over all of the agents; in this phase, the allocation is obviously WEF1 since each agent has at most one item at any point. We will show that, after this phase, the algorithm generates a picking sequence over the agents with the following property:

Lemma 3.4. *Consider an agent i chosen by Algorithm 1 to pick an item at some iteration t, and suppose that this is not her first pick. Let t_i and t_j be the numbers of times agent i and some other agent j appear in the prefix of iteration t in the sequence respectively, not including iteration t itself. Then $\frac{t_j}{t_i} \geq \frac{w_j}{w_i}$.*

This property is sufficient to ensure that the latest picker does not attract weighted envy up to more than one item towards herself after her latest pick:

Lemma 3.5. *Suppose that, for every iteration t in which agent i picks an item after her first pick, the numbers of times that agent i and some other agent j appear in the prefix of the iteration in the sequence, not including iteration t itself — t_i and t_j respectively — satisfy the relation $\frac{t_j}{t_i} \geq \frac{w_j}{w_i}$. Then, in the partial allocation up to and including i's latest pick, agent j is weighted envy-free towards i up to the first item i picked.*

Obviously, directly after an agent picks an item, her envy towards other agents cannot get any worse than before. Since the partial allocation after the initial round-robin phase is WEF1 and every agent is weighted envy-free up to one item towards every subsequent picker due to Lemmas 3.4 and 3.5, the allocation is WEF1 at every iteration, and in particular at the end of the algorithm. Hence, for the proof of correctness, it suffices to prove the two lemmas.

Proof of Lemma 3.4. Since agent i is picked at iteration t, it must be the case that $i \in \arg\min_{k \in N} \frac{t_k}{w_k}$. This means that $\frac{t_i}{w_i} \leq \frac{t_j}{w_j}$, i.e., $\frac{t_j}{t_i} \geq \frac{w_j}{w_i}$ since $t_i > 0$. □

Proof of Lemma 3.5. Let $\gamma := \frac{w_j}{w_i}$. Consider any iteration t in which agent i is chosen after her first pick. Let agent j's values for the items allocated to agent i in the latter's second, third, ..., $(t_i + 1)^{\text{st}}$ picks (the last one occurring at the iteration t under consideration) be $\beta_1, \beta_2, \ldots, \beta_{t_i}$ respectively. If o^* is the first item picked by agent i and A^t the partial allocation up to and including iteration t, then clearly $v_j(A_i^t \backslash \{o^*\}) = \sum_{x=1}^{t_i} \beta_x$. Let the number of times agent j appears in the prefix of agent i's second pick be τ_1; that between agent i's second and third picks be τ_2; ...; that between agent i's t_i^{th} and $(t_i + 1)^{\text{st}}$ picks be τ_{t_i}. Let agent j's values for the items she herself picked during phase $x \in [t_i]$ be $\alpha_1^x, \alpha_2^x, \ldots, \alpha_{\tau_x}^x$ respectively, where the phases are defined as in the previous sentence. Then, $v_j(A_j^t) = \sum_{x=1}^{t_i} \sum_{y=1}^{\tau_x} \alpha_y^x$. Now, since $\sum_{x=1}^{r} \tau_x$ and r are the numbers of times agents j and i appear in the prefix of the latter's $(r+1)^{\text{st}}$ pick respectively, the condition of the lemma dictates that

$$\sum_{x=1}^{r} \tau_x \geq r\gamma \quad \forall r \in [t_i]. \tag{1}$$

Note that $\tau_1 \geq \gamma > 0$; however, τ_x can be zero for $x \in \{2, 3, \ldots, t_i\}$ — this corresponds to the scenario where agent i picked more than once without agent j picking in between. Moreover, every time agent j was chosen, she picked one of the items she values the most among those available, including the items picked by agent i later. Hence, if $\tau_x > 0$ for some $x \in [t_i]$, then

$$\alpha_y^x \geq \max\{\beta_x, \beta_{x+1}, \ldots, \beta_{t_i}\} \quad \forall y \in [\tau_x]$$
$$\Rightarrow \quad \sum_{y=1}^{\tau_x} \alpha_y^x \geq \tau_x \max\{\beta_x, \beta_{x+1}, \ldots, \beta_{t_i}\}. \tag{2}$$

Note that Inequality (2) holds trivially if $\tau_x = 0$ since both sides are zero; hence it holds for every $x \in [t_i]$.

We claim that for each $r \in [t_i]$,

$$\sum_{x=1}^{r} \sum_{y=1}^{\tau_x} \alpha_y^x \geq \gamma \sum_{x=1}^{r} \beta_x + \left(\sum_{x=1}^{r} \tau_x - r\gamma \right) \max\{\beta_r, \beta_{r+1}, \ldots, \beta_{t_i}\}.$$

To prove the claim, we proceed by induction on r. For the base case $r = 1$, we have from Inequality (2) that

$$\sum_{y=1}^{\tau_1} \alpha_y^1 \geq \tau_1 \max\{\beta_1, \beta_2, \ldots, \beta_{t_i}\}$$
$$\geq \gamma\beta_1 + (\tau_1 - \gamma) \max\{\beta_1, \beta_2, \ldots, \beta_{t_i}\}.$$

For the inductive step, assume that the claim holds for $r - 1$; we will prove it for r. We have

$$\sum_{x=1}^{r} \sum_{y=1}^{\tau_x} \alpha_y^x = \sum_{x=1}^{r-1} \sum_{y=1}^{\tau_x} \alpha_y^x + \sum_{y=1}^{\tau_r} \alpha_y^r$$

$$
\begin{aligned}
&\geq \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^{r-1} \tau_x - (r-1)\gamma \right) \max\{\beta_{r-1}, \beta_r, \ldots, \beta_{t_i}\} \\
&\quad + \sum_{y=1}^{\tau_r} \alpha_y^r \\
&\geq \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^{r-1} \tau_x - (r-1)\gamma \right) \max\{\beta_{r-1}, \beta_r, \ldots, \beta_{t_i}\} \\
&\quad + \tau_r \max\{\beta_r, \beta_{r+1}, \ldots, \beta_{t_i}\} \\
&\geq \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^{r-1} \tau_x - (r-1)\gamma \right) \max\{\beta_r, \beta_{r+1}, \ldots, \beta_{t_i}\} \\
&\quad + \tau_r \max\{\beta_r, \beta_{r+1}, \ldots, \beta_{t_i}\} \\
&= \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^{r} \tau_x - (r-1)\gamma \right) \max\{\beta_r, \beta_{r+1}, \ldots, \beta_{t_i}\} \\
&= \gamma \sum_{x=1}^{r-1} \beta_x + \gamma \max\{\beta_r, \beta_{r+1}, \ldots, \beta_{t_i}\} \\
&\quad + \left(\sum_{x=1}^{r} \tau_x - r\gamma \right) \max\{\beta_r, \beta_{r+1}, \ldots, \beta_{t_i}\} \\
&\geq \gamma \sum_{x=1}^{r-1} \beta_x + \gamma \beta_r + \left(\sum_{x=1}^{r} \tau_x - r\gamma \right) \max\{\beta_r, \beta_{r+1}, \ldots, \beta_{t_i}\} \\
&= \gamma \sum_{x=1}^{r} \beta_x + \left(\sum_{x=1}^{r} \tau_x - r\gamma \right) \max\{\beta_r, \beta_{r+1}, \ldots, \beta_{t_i}\},
\end{aligned}
$$

where the first inequality follows directly from the inductive hypothesis, the second from Inequality (2), and the third from Inequality (1) as well as the fact that

$$\max\{\beta_{r-1}, \beta_r, \ldots, \beta_{t_i}\} \geq \max\{\beta_r, \beta_{r+1}, \ldots, \beta_{t_i}\}.$$

This completes the induction and establishes the claim.

Now, taking $r = t_i$ in the claim, we get

$$
\begin{aligned}
\sum_{x=1}^{t_i} \sum_{y=1}^{\tau_x} \alpha_y^x &\geq \gamma \sum_{x=1}^{t_i} \beta_x + \left(\sum_{x=1}^{t_i} \tau_x - t_i \gamma \right) \beta_{t_i} \\
&\geq \gamma \sum_{x=1}^{t_i} \beta_x,
\end{aligned}
$$

where we use Inequality (1) again for the second inequality. This implies that $v_j(A_j^t) \geq \frac{w_j}{w_i} \cdot v_j(A_i^t \backslash \{o^*\})$, i.e., agent j is weighted envy-free towards agent i up to one item, concluding the proof of the lemma and therefore the proof of correctness. ◻

For the time-complexity, note that there are $O(m)$ iterations of the **while** loop. In each iteration, determining the next picker takes $O(n)$ time, while letting the picker pick her favorite item takes $O(m)$ time. Since we may assume that $m > n$ (otherwise it suffices to allocate at most one item to every agent), the algorithm runs in time $O(m^2)$. ◻

If w_i equals a positive constant w for every $i \in N$, then Algorithm 1 degenerates into the traditional round-robin procedure which is guaranteed to return an EF1 allocation for additive valuations, but may not be PO; as such, Algorithm 1 may not produce a PO allocation either. This is easily seen in the following example: $N = [2]$, $O = \{o_1, o_2\}$; $w_1 = w_2 = 1$; $v_1(o_1) = 0.5$, $v_1(o_2) = 0.5$, $v_2(o_1) = 0.8$, $v_2(o_2) = 0.2$. With lexicographic tie-breaking for both agents and items, our algorithm will give us $A_1 = \{o_1\}$ and $A_2 = \{o_2\}$, which is Pareto dominated by $A_1' = \{o_2\}$ and $A_2' = \{o_1\}$.

4 WEF1 AND PO ALLOCATIONS

Our next question is whether WEF1 notion can be achieved in conjunction with economic efficiency. When agents have equal weight, it is known that fairness and efficiency are compatible: Caragiannis et al. [11] show that an allocation maximizing the Nash social welfare satisfies PO and EF1. Unfortunately, this approach is not applicable to our setting: we show that the MWNW allocation may fail to be WEF1. In fact, we prove an even stronger negative result: for any fixed c, the allocation may fail to be WEFc.

Proposition 4.1. *Let c be an arbitrary positive integer. There exists a problem instance with two agents having identical additive valuations for which any* MWNW *allocation is not* WEFc.

Given that an MWNW allocation may not be WEF1 in our setting, a natural question is whether there is an alternative approach to guarantee the existence of a PO and WEF1 allocation. We first show that this is indeed the case for two agents: we establish that an allocation satisfying PO and WEF1 exists and can be computed in polynomial time for two agents with additive valuations, by adapting the classic adjusted winner procedure [9] to the weighted setting.

Theorem 4.2. *For two agents with additive valuations and arbitrary positive real weights, a complete* WEF1 *and* PO *allocation always exists and can be computed in polynomial time.*

Algorithm 2 Weighted Adjusted Winner

Require: $\frac{v_1(o_1)}{v_2(o_1)} \geq \frac{v_1(o_2)}{v_2(o_2)} \geq \cdots \geq \frac{v_1(o_m)}{v_2(o_m)}$ w.l.o.g.

$d \leftarrow 1$.
while $\frac{1}{w_1} \sum_{r=1}^{d} v_1(o_r) < \frac{1}{w_2} \sum_{r=d+2}^{m} v_1(o_r)$ **do**
    $d \leftarrow d + 1$.
end while
$A_1 \leftarrow \{o_1, \ldots, o_d\}$.
$A_2 \leftarrow \{o_{d+1}, \ldots, o_m\}$.

Proof sketch. Our proof is constructive: the Weighted Adjusted Winner algorithm as delineated in Algorithm 2 produces an allocation satisfying the theorem statement whenever $v_1(o) > 0$ and $v_2(o) > 0$ for every $o \in O$.[2] Note that we can safely disregard items valued at zero by both agents; in case there is an item valued positively by only one agent, we initialize the bundle A_i with items valued positively by agent $i \in \{1, 2\}$ only, i.e., $A_1^0 = \{o \in O : v_1(o) > 0, v_2(o) = 0\}$ and $A_2^0 = \{o \in O : v_2(o) > 0, v_1(o) = 0\}$, then run Algorithm 2 on the remaining items and use its output (A_1, A_2) to augment the respective bundles. The proof of correctness is a natural extension of that for the original adjusted winner procedure [9]: if the **while** loop terminates with $d = d^*$, then agent 1 (resp., agent 2) becomes weighted envy-free towards agent 2 (resp.,

[2]For $d \geq m - 1$, we set the right-hand side of the **while** loop condition to zero.

agent 1) upon dropping item o_{d^*+1} (resp., o_{d^*}) from the latter's bundle. Since each new **while** loop condition can be checked in $O(1)$ time, the time complexity is dominated by the sorting of the m ratios and hence is $O(m \log m)$. □

4.1 WEF1 and PO Allocations for $n \geq 2$ Agents

Having resolved the existence question of PO and WEF1 for two agents, we now investigate whether such an allocation always exists for any number of agents, answering the question in the affirmative. To this end, we employ a weighted modification of the algorithm by Barman et al. [5], who design a pseudo-polynomial time algorithm to find a PO and EF1 allocation when agents have additive valuations in the unweighted setting. As in Barman et al. [5], we consider the artificial market where each item has a price and agents purchase a bundle of items with the highest ratio of value to price, called *bang per buck ratio*. This allows us to measure the degree of fairness of a given allocation in terms of the prices.

Formally, a *price vector* is an m-dimensional non-negative real vector $\boldsymbol{p} = (p_1, p_2, \ldots, p_m) \in \mathbb{R}^O_{\geq 0}$; we call p_o the *price* of item $o \in O$, and write $p(X) = \sum_{o \in X} p_o$. Let A be an allocation and $\boldsymbol{p}$ be a price vector. We call each $p(A_i)$ the *spending* and $\frac{1}{w_i} p(A_i)$ the *weighted spending* of agent i. We now define a weighted version of the price envy-freeness up to one item (pEF1) notion introduced by Barman et al. [5].

Definition 4.3. Given an allocation A and a price vector $\boldsymbol{p}$, we say that A is *weighted price envy-free up to one item* (WpEF1) with respect to $\boldsymbol{p}$ if for any pair of agents i, j, either $A_j = \emptyset$ or $\frac{1}{w_i} p(A_i) \geq \frac{1}{w_j} \min_{o \in A_j} p(A_j \setminus \{o\})$.

The *bang per buck ratio* of item o for agent i is $\frac{v_i(o)}{p_o}$; we write the maximum bang per buck ratio for agent i as $\alpha_i(\boldsymbol{p})$. We refer to the items with maximum bang per buck ratio for i as i's *MBB items* and denote the set of such items by $\mathrm{MBB}_i(\boldsymbol{p})$ for each $i \in N$. The following lemma is a straightforward adaptation of Lemma 4.1 in [5] to our setting; it ensures that one can obtain the property of WEF1 by balancing among the spending of agents under the MBB condition.

Lemma 4.4. *Given a complete allocation A and a price vector $\boldsymbol{p}$, suppose that allocation A satisfies* WpEF1 *with respect to $\boldsymbol{p}$ and agents are assigned to MBB items only, i.e., $A_i \subseteq \mathrm{MBB}_i(\boldsymbol{p})$ for each $i \in N$. Then A is* WEF1.

It is also known that if each agent i only purchases MBB items, i.e., i maximizes his utility under the budget $p(A_i)$, then the corresponding allocation is Pareto optimal.

Lemma 4.5 (First Welfare Theorem; Mas-Colell et al. [22], Chapter 16). *Given a complete allocation A and a price vector $\boldsymbol{p}$, suppose that agents are assigned to MBB items only, i.e., $A_i \subseteq \mathrm{MBB}_i(\boldsymbol{p})$ for each $i \in N$. Then A is* PO.

Now, the problem of finding a PO and WEF1 allocation reduces to that of finding an allocation and price vector pair satisfying the MBB condition and WpEF1. Similarly to Barman et al. [5], we develop an algorithm that alternates between two phases: the first phase consists of reallocating items from large to small spenders, and the second phase consists of increasing the prices of the items owned by small spenders. We show that by increasing prices gradually, the algorithm converges to an allocation which is guaranteed to be both PO and WEF1 by Lemmas 4.4 and 4.5.

Theorem 4.6. *For any number of agents with additive valuations and arbitrary positive real weights, there exists a* WEF1 *and* PO *allocation.*

5 WWEF1 AND PO ALLOCATIONS

In the previous section, we saw that MWNW allocations may fail to satisfy WEF1, showing that Caragiannis et al. [11]'s result from the unweighted setting does not extend to the weighted setting via WEF1. Do MWNW allocations provide any fairness guarantee? The answer is positive: we show that a MWNW allocation indeed satisfies WWEF1, a weaker fairness notion that also generalizes EF1.

Theorem 5.1. *For any number of agents with additive valuations and arbitrary positive real weights, a* MWNW *allocation is always* WWEF1 *and* PO.

Proof sketch. Let A be a MWNW allocation, with $N_{\max}$ being the subset of agents having strictly positive realized valuations under A. If it were not PO, there would exist an allocation $\hat{A}$ such that $v_i(\hat{A}_i) > v_i(A_i)$ for some $i \in N$ and $v_j(\hat{A}_j) \geq v_j(A_j)$ for every $j \in N \backslash \{i\}$. If $i \in N \backslash N_{\max}$, we would have $v_j(\hat{A}_j) > 0$ for every $j \in N_{\max} \cup \{i\}$, which contradicts the assumption that $N_{\max}$ is a largest subset of agents to which it is possible to give positive valuations simultaneously. If $i \in N_{\max}$, then $\prod_{j \in N_{\max}} v_j(\hat{A}_j)^{w_j} > \prod_{j \in N_{\max}} v_j(A_j)^{w_j}$, which violates the optimality of the right-hand side. This proves that A is PO.

As in Caragiannis et al. [11], we will start by proving that A is WWEF1 for the scenario $N_{\max} = N$ and then address the case $N_{\max} \neq N$. Assume that $N_{\max} = N$. If A is not WWEF1, then there exists a pair of agents $i, j \in N$ such that i has weak weighted envy towards j up to more than one item. Clearly, there must be at least two items in j's bundle that i values positively. Moreover, j must value these items positively as well — otherwise we can transfer them to i and obtain a Pareto improvement.

Let $A_j^i := \{o \in A_j : v_i(o) > 0\}$. We construct another allocation A' by transferring an item o^* (to be chosen later) from j to i so that $A'_i = A_i \cup \{o^*\}$, $A'_j = A_j \backslash \{o^*\}$, and $A'_r = A_r, \forall r \in N \backslash \{i, j\}$. We have

$$\begin{aligned}
\frac{\mathrm{WNW}(A')}{\mathrm{WNW}(A)} &= \left(\frac{v_i(A_i \cup \{o^*\})}{v_i(A_i)}\right)^{w_i} \left(\frac{v_j(A_j \backslash \{o^*\})}{v_j(A_j)}\right)^{w_j} \\
&= \left(\frac{v_i(A_i) + v_i(o^*)}{v_i(A_i)}\right)^{w_i} \left(\frac{v_j(A_j) - v_j(o^*)}{v_j(A_j)}\right)^{w_j} \\
&= \left(1 + \frac{v_i(o^*)}{v_i(A_i)}\right)^{w_i} \left(1 - \frac{v_j(o^*)}{v_j(A_j)}\right)^{w_j}.
\end{aligned}$$

First, note that $v_j(o) > 0, \forall o \in A_j^i$; otherwise the above ratio for o^* with $v_j(o^*) = 0$ equals $\left(1 + \frac{v_i(o^*)}{v_i(A_i)}\right)^{w_i} > 1$, contradicting the assumption that A is a MWNW allocation. However, even under this condition, we will show that if agents i, j violated the WWEF1 property, the above ratio would still exceed 1 for some item o^*.

Case I. $w_i \geq w_j$. Let us pick an item $o^* \in \arg\min_{o \in A_j^i} \frac{v_j(o)}{v_i(o)}$ specifically to transfer from j to i for changing the allocation from A to A'. This is well-defined by the definition of A_j^i. Consider

$$\left[\frac{\mathsf{WNW}(A')}{\mathsf{WNW}(A)}\right]^{\frac{1}{w_j}} = \left(1 + \frac{v_i(o^*)}{v_i(A_i)}\right)^{\frac{w_i}{w_j}} \left(1 - \frac{v_j(o^*)}{v_j(A_j)}\right)$$

where $1 - \frac{v_j(o^*)}{v_j(A_j)} > 0$ since $v_j(A_j) > v_j(o^*) > 0$, and

$$\left(1 + \frac{v_i(o^*)}{v_i(A_i)}\right)^{\frac{w_i}{w_j}} \geq \left(1 + \frac{w_i}{w_j} \cdot \frac{v_i(o^*)}{v_i(A_i)}\right)$$

from Bernoulli's inequality, since $\frac{v_i(o^*)}{v_i(A_i)} > 0$ and $\frac{w_i}{w_j} \geq 1$. As in Caragiannis et al. [11], simple algebra shows that

$$\begin{aligned} &\left(1 + \frac{w_i}{w_j} \cdot \frac{v_i(o^*)}{v_i(A_i)}\right)\left(1 - \frac{v_j(o^*)}{v_j(A_j)}\right) > 1 \\ \Leftrightarrow \quad & \frac{v_i(A_i)}{w_i} < \frac{v_i(o^*)}{v_j(o^*)}\left(\frac{v_j(A_j) - v_j(o^*)}{w_j}\right). \end{aligned} \tag{3}$$

The latter inequality is true under our assumptions for the following reasons: Since the "bigger" agent i has weak weighted envy towards the "smaller" agent j up to more than one item, $\frac{v_i(A_i)}{w_i} < \frac{v_i(A_j) - v_i(o^*)}{w_j}$; due to our choice of o^*,

$$\frac{v_j(o^*)}{v_i(o^*)} \leq \frac{\sum_{o \in A_j^i} v_j(o)}{\sum_{o \in A_j^i} v_i(o)} \leq \frac{\sum_{o \in A_j} v_j(o)}{\sum_{o \in A_j} v_i(o)} = \frac{v_j(A_j)}{v_i(A_j)},$$

since $\sum_{o \in A_j \setminus A_j^i} v_j(o) \geq 0$ and $\sum_{o \in A_j \setminus A_j^i} v_i(o) = 0$. Plugging $v_i(A_j) \leq \frac{v_i(o^*)}{v_j(o^*)} v_j(A_j)$ into the above strict inequality and simplifying, we obtain (3). But chaining all these inequalities together, we get

$$\left[\frac{\mathsf{WNW}(A')}{\mathsf{WNW}(A)}\right]^{\frac{1}{w_j}} > 1 \quad \Rightarrow \quad \mathsf{WNW}(A') > \mathsf{WNW}(A).$$

This is a contradiction, which shows that A is WWEF1 in this case.

Case II. $w_i < w_j$. We can pick an item $o^* \in \arg\max_{o \in A_j^i} \frac{v_i(o)}{v_j(o)}$ and proceed similarly to Case I, noting that the smaller agent i having weak weighted envy towards the bigger agent j up to more than one item implies having $\frac{v_i(A_i) + v_i(o^*)}{w_i} < \frac{v_i(A_j)}{w_j}$; this leads to the same contradiction $\mathsf{WNW}(A') > \mathsf{WNW}(A)$.

If $N_{\max} \subsetneq N$, our argument mirrors the corresponding part of the proof of Caragiannis et al. [11]'s Theorem 3.2. The key ideas are: there can be no (weighted) envy towards any $i \notin N_{\max}$ or weak weighted envy up to more than one item between any two agents $i, j \in N_{\max}$ (from the proof for $N_{\max} = N$); if some $i \in N \setminus N_{\max}$ were not weakly weighted envy-free up to one item towards $j \in N_{\max}$, we could transfer one item from j to i and keep both i and j's valuations positive, contradicting the maximality of $N_{\max}$. ◻

6 WEF1 AND OTHER FAIRNESS NOTIONS

An allocation that satisfies multiple fairness guarantees is naturally desirable but often elusive in the setting with indivisible items. Hence, we will now explore the implications of the WEF1 property for the other fairness criteria defined in Section 2.

For additive valuations, Aziz et al. [3] provide a polynomial-time algorithm for computing a PO and WPROP1 allocation, whereas we prove the existence of PO and WEF1 allocations in Section 4.1. It is straightforward to show that, in the unweighted scenario, any complete envy-free allocation is also proportional for subadditive valuations and any complete EF1 allocation is PROP1 for additive valuations (see, e.g., [3]). This begs the question: does the WEF1 property along with completeness also imply the WPROP1 condition? Unfortunately, the answer is no in general — in fact, we establish a stronger result in the following proposition.

Proposition 6.1. *For any number $n \geq 2$ of agents with additive valuations and arbitrary positive real weights, any complete* WEF1 *allocation is* WPROP$(n-1)$*. However, for any $n \geq 3$, there exists an instance in which no complete* WEF1 *allocation is* WPROP$(n-2)$.

For n symmetric (unweighted) agents with additive valuations, Amanatidis et al. [1, Prop. 3.6] show that any complete EF1 allocation is $\frac{1}{n}$-MMS and this approximation guarantee is tight. Moreover, as Caragiannis et al. [11, Thm. 4.1] prove, every maximum Nash welfare allocation, which is EF1 and PO, is also $\Theta(1/\sqrt{n})$-MMS. This means that, for a small number of agents, the EF1 property provides a reasonable approximation to MMS fairness. However, for any number of agents with asymmetric weights, the WEF1 condition does not imply any positive approximation of the WMMS guarantee, even in conjunction with Pareto optimality.

Proposition 6.2. *For any constant $\epsilon > 0$ and any number $n \geq 2$ of agents, there exists an instance with additive valuations in which some* PO *and* WEF1 *allocation is not ϵ-*WMMS.

For the special case of $n = 2$ agents, the first part of Proposition 6.1 implies that the output of the Weighted Adjusted Winner Algorithm in Section 4, which is a complete WEF1 allocation, is WPROP1; however, we show that it comes with no guarantee on the WMMS approximation.

Proposition 6.3. *The output of the Weighted Adjusted Winner procedure for two agents (Algorithm 2) is always* WPROP1*. However, for any constant $\epsilon > 0$, it is not necessarily ϵ-*WMMS.

7 EXPERIMENTS

Thus far, we have thoroughly investigated the existence and computational properties of approximations to *weighted envy-freeness* (WEF). While the WEF notion itself obviously cannot always be satisfied with indivisible items, it is still interesting to ask how "likely" it is for a problem instance with weighted agents to admit a WEF allocation, and compare that to the unweighted setting.

In this section, we approach this question experimentally by generating sets of 1000 instances with $n \in \{2, 3, 4, 5\}$ agents and $m \in \{2, 3, \ldots, 9\}$ items wherein each agent's value for each item is drawn independently from a distribution. We determine by exhaustive search over all allocations whether a WEF allocation exists for each instance for two different weight vectors: $w_i = 1$ (unweighted) and $w_i = i$ (weighted) for every $i \in N$. Figure 1 shows the results for two common distributions. The main takeaway is that weighted envy-free allocations are almost always harder to find than their unweighted counterparts. This illustrates the difficulty of achieving weighted envy-freeness and further justifies our quest for the (strong and weak) relaxations of the WEF property.

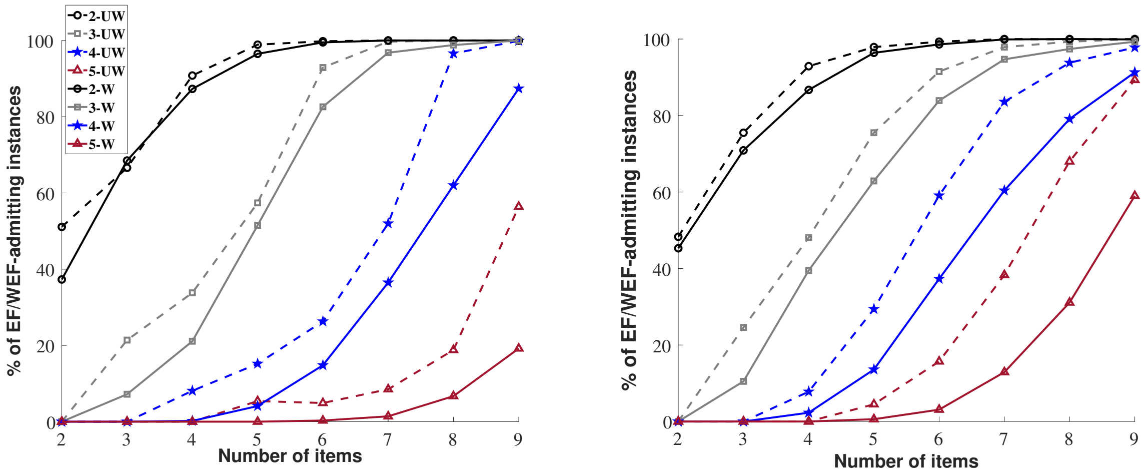

(a) Uniform ditribution on $[0, 1]$

(b) Exponential distribution with mean 1

Figure 1: Percentages of instances that admit EF and WEF allocations for two different valuation distributions in our experiments; n-UW, depicted by dashed curves (resp., n-W, depicted by solid curves) refers to a scenario with n unweighted (resp., weighted) agents with equal weights (resp., weights proportional to agent indices) in both graphs.

8 DISCUSSION AND FUTURE WORK

In this paper, we have defined and characterized natural extensions of envy-freeness up to one good (EF1) for agents with asymmetric weights. We conclude with some hurdles we faced while trying to extend weighted envy concepts beyond additive valuations. First, we show that even for simple non-additive valuations, the existence of a WEF1 or WWEF1 allocation can no longer be guaranteed.

Proposition 8.1. *There exists an instance with $n = 2$ such that one of the agents has a (normalized and monotone) submodular valuation, the other agent has an additive valuation, and a complete* WWEF1 *allocation (or a complete* WEF1 *allocation) does not exist.*

Proof. Suppose $N = [2]$; $w_1 = 1$ and $w_2 = 2$; $m > 5$; the valuation functions are: $v_1(S) = |S|$ and $v_2(S) = 1$ for every $S \in 2^O \setminus \emptyset$, $v_1(\emptyset) = v_2(\emptyset) = 0$. The functions are obviously normalized and monotone; $v_1(\cdot)$ is additive and $v_2(\cdot)$ is submodular (agent 2 is indifferent among all non-empty bundles). Note that for any allocated bundles A_1 and A_2 such that $|A_1| \geq 2$ and A_2 is non-empty, $v_2(A_2)/w_2 = 1/2$ but $v_2(A_1 \setminus \{o\})/w_1 = 1 > v_2(A_2)/w_2$ for every $o \in A_1$, and $v_2(A_2 \cup \{o\})/w_2 = 1/2 < v_2(A_1)/w_1$ for every $o \in A_1$. Thus, the only way to make agent 2 (weakly) weighted envy-free up to one item towards agent 1 is to ensure that $|A_1| \leq 1$. Assume without loss of generality that $A_1 = \{o_1\}$ (if $A_1 = \emptyset$, agent 1 will be even worse off in the argument that follows). To make the allocation complete, we must have $A_2 = O \setminus \{o_1\}$, so that $v_1(A_1) = 1$ and $v_1(A_2) = |A_2| = m - 1$. Since agent 1 has an additive valuation and a smaller weight than agent 2, she would be weakly weighted envy-free up to one item towards agent 2 if and only if there is an item $o \in A_2$ such that $v_1(A_1 \cup \{o\})/w_1 \geq v_1(A_2)/w_2 = (m-1)/2 > 2$, since $m > 5$. However, for any $o \in A_2$, we have $v_1(A_1 \cup \{o\})/w_1 = v_1(\{o_1, o\}) = 2$. Hence the allocation cannot be WWEF1. □

One of the key ideas in the proof of Theorem 5.1 is what we can call the *transferability* property: if agent i has (weighted) envy towards agent j under additive valuations, then there is at least one item o in j's bundle for which agent i has positive (marginal) utility, i.e., the item o could be transferred from j to i to augment i's realized valuation. Unfortunately, this property no longer holds for non-additive valuations.

Proposition 8.2. *There exists an instance such that an agent i with a non-additive valuation function has weighted envy towards agent j under allocation A, but there is no item in j's bundle for which i has positive marginal utility, i.e., $\nexists o \in A_j$ such that $v_i(A_i \cup \{o\}) - v_i(A_i) > 0$.*

Proof. Consider the example in Proposition 8.1. Under any allocation with $|A_1| = m - 1$ and $|A_2| = 1$, agent 2 has weighted envy towards agent 1 since $v_2(A_2) = 1/2 < 1 = v_2(A_1)/w_1$. However, $v_2(A_2 \cup \{o\}) = 1 = v_2(A_2)$ for every $o \in A_1$. □

In addition to exploring weighted envy-based fairness notions for non-additive (e.g., submodular) valuations, other potential directions for future research include identifying conditions under which WEF allocations are likely to exist (Section 7)[3] and investigating weighted envy in the allocation of *chores* (items with negative valuations). It would also be interesting to consider weighted versions of other envy-freeness approximations, such as *EFX* [11, 23].

ACKNOWLEDGMENTS

Chakraborty and Zick are supported by the Singapore NRF Research Fellowship R-252-000-750-733. Igarashi is supported by the KAKENHI Grant-in-Aid for JSPS Fellows no. 18J00997 and JST, ACT-X. Suksompong is supported by the European Research Council (ERC) under grant no. 639945 (ACCORD), his visit to NUS is supported by MOE Grant R-252-000-625-133. We also thank the anonymous AAMAS reviewers for their constructive comments.

[3] This has been done for envy-freeness in the unweighted setting [16, 21].

REFERENCES

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

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

[3] Haris Aziz, Hervé Moulin, and Fedor Sandomirskiy. 2019. A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation. *arXiv preprint arXiv:1909.00740* (2019).

[4] Moshe Babaioff, Noam Nisan, and Inbal Talgam-Cohen. 2019. Fair allocation through competitive equilibrium from generic incomes. In *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency (ACM FAT*)*. 180.

[5] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018. Finding fair and efficient allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. 557–574. Extended version available as arXiv:1707.04731.

[6] Nawal Benabbou, Mithun Chakraborty, Edith Elkind, and Yair Zick. 2019. Fairness towards groups of agents in the allocation of indivisible items. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 95–101.

[7] Nawal Benabbou, Mithun Chakraborty, Vinh Ho Xuan, Jakub Sliwinski, and Yair Zick. 2018. Diversity constraints in public housing allocation. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 973–981.

[8] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. 2016. Fair allocation of indivisible goods. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, Chapter 12, 284–310.

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

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

[11] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2016. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*. 305–322.

[12] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. 2019. Weighted envy-freeness in indivisible item allocation. *arXiv preprint arXiv:1909.10502* (2019).

[13] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. 2017. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*. 629–646.

[14] Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan. 2019. Group fairness for the allocation of indivisible goods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 1853–1860.

[15] Ágnes Cseh and Tamás Fleiner. 2018. The complexity of cake cutting with unequal shares. In *Proceedings of the 11th International Symposium on Algorithmic Game Theory (SAGT)*. 19–30.

[16] John P. Dickerson, Jonathan Goldman, Jeremy Karp, Ariel D. Procaccia, and Tuomas Sandholm. 2014. The computational rise and fall of fairness. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*. 1405–1411.

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

[18] Duncan K. Foley. 1967. Resource allocation and the public sector. *Yale Economics Essays* 7, 1 (1967), 45–98.

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

[20] Pasin Manurangsi and Warut Suksompong. 2017. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences* 89 (2017), 100–108.

[21] Pasin Manurangsi and Warut Suksompong. 2019. When do envy-free allocations exist?. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 2109–2116.

[22] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. Oxford University Press.

[23] Benjamin Plaut and Tim Roughgarden. 2018. Almost envy-freeness with general valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2584–2603.

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

[25] Erel Segal-Halevi. 2019. Cake-cutting with different entitlements: How many cuts are needed? *J. Math. Anal. Appl.* 480, 1 (2019), 123382.

[26] Warut Suksompong. 2018. Approximate maximin shares for groups of agents. *Mathematical Social Sciences* 92 (2018), 40–47.

[27] Hal R. Varian. 1974. Equity, envy, and efficiency. *Journal of Economic Theory* 9 (1974), 63–91.

[28] Dao-Zhi Zeng. 2000. Approximate envy-free procedures. In *Game Practice: Contributions from Applied Game Theory*, Fioravante Patrone, Ignacio García-Jurado, and Stef Tijs (Eds.). Springer, Chapter 17, 259–271.