Almost Envy-Free Allocations of Indivisible Goods or Chores with Entitlements

Max Springer[1], MohammadTaghi Hajiaghayi[2], Hadi Yami[3]

[1]University of Maryland, Department of Mathematics, College Park MD
[2]University of Maryland, Department of Computer Science, College Park MD
[3]Microsoft Corporation, Redmond WA
{mss423, hajiagha}@umd.edu, hayami@microsoft.com

Abstract

We here address the problem of fairly allocating indivisible goods or chores to n agents with *weights* that define their entitlement to the set of indivisible resources. Stemming from well-studied fairness concepts such as envy-freeness up to one good (EF1) and envy-freeness up to any good (EFX) for agents with *equal* entitlements, we present, in this study, the first set of impossibility results alongside algorithmic guarantees for fairness among agents with *unequal* entitlements.

Within this paper, we expand the concept of envy-freeness up to any good or chore to the weighted context (WEFX and XWEF respectively), demonstrating that these allocations are not guaranteed to exist for two or three agents. Despite these negative results, we develop a WEFX procedure for two agents with *integer* weights, and furthermore, we devise an approximate WEFX procedure for two agents with *normalized* weights. We further present a polynomial-time algorithm that guarantees a weighted envy-free allocation up to one chore (1WEF) for any number of agents with additive cost functions. Our work underscores the heightened complexity of the weighted fair division problem when compared to its unweighted counterpart.

Introduction

The notion of a fair allocation of resources has been a fundamental problem in the field of economics and, more recently, computer science (Bouveret, Chevaleyre, and Maudet 2016; Brams and Taylor 1995; Peterson and Su 2002, 2009; Pikhurko 2000; Procaccia 2009; Robertson and Webb 1998). This idea of fairness takes on many forms in a wide array of application areas such as land division, divorce settlements, or natural resource distributions.

One of the most well-established mathematical definitions of fairness is *envy-freeness*. In an envy-free allocation each agent prefers her share more than any other: for the case of goods, she wants the most lucrative set of items whereas for chores, we seek to minimize the cost distributed across agents. However, envy-freeness is not necessarily obtainable when dealing with a set of indivisible goods or chores, and such an allocation is often difficult to compute. Hence, several relaxations of envy-freeness have been introduced.

Two such relaxations that we examine in the present work are *envy-free up to one good* (EF1) and *envy-free up to any good* (EFX). An allocation of indivisible items is EF1 if any possible envy of an agent for the share of another agent can be resolved by removing some good from the envied share. (Lipton et al. 2004) and (Budish 2010) provide polynomial time algorithms for an EF1 allocation for any number of agents. An allocation is said to be EFX if no agent envies another agent after the removal of *any* item from the other agent's bundle. Theoretically, this notion is strictly stronger than EF1 and as a result, despite significant effort, the existence of EFX allocations is still unknown. (Plaut and Roughgarden 2018) described the first *approximate* EFX results with later improvements by (Amanatidis, Markakis, and Ntokos 2020) and (Farhadi et al. 2021). With respect to chore division, the literature is more scarce: the first discrete and bounded envy-free protocol for any number of agents was not proposed until 2018 (Dehghani et al. 2018).

In both the good and chore division problems, the literature noted above is restricted to the case where each agent's perspective is weighted equally. However, more recently, the generalization has been proposed where agents have a valuation (or cost in the case of chores) associated with each item as well as an "entitlement" or weighting (Aziz, Chan, and Li 2019; Babaioff, Ezra, and Feige 2021; Chakraborty et al. 2020; Farhadi et al. 2019). This problem setting captures significantly more characteristics of real-world fair division issues than that of the equal entitlement case.

To further emphasize this natural problem, consider the recent example that arose during the height of the COVID-19 pandemic. When a vaccine was in production, producers were faced with the question of how to fairly distribute doses in a manner that respects all parties while also mitigating further spread of the virus. In this context, it is not enough to just give each nation an equal share of the supply – we must now take into account population densities and each nation's medical infrastructure among a plethora of other factors. This naturally gives rise to agent weightings that, as we will show in the paper, result in a considerable complication of the fair allocation problem.

Our Contributions

Although the existence of EFX allocations is a major open problem in the field, we demonstrate that WEFX, the gener-

	1WEF	WEFX (id.)	WEFX (int.)	WEFX (add.)	XWEF	α-WEFX
$n = 2$	✓(Thm. 19)	✓(Thm. 6)	✓(Thm. 10)	✗(Thm. 8)	✗(Thm. 17)	✓(Thm. 11)
$n = 3$	✓(Thm. 19)	✓(Thm. 6)	?	✗(Thm. 9)	✗(Thm. 18)	?
$n > 3$	✓(Thm. 19)	✓(Thm. 6)	?	?	?	?

Table 1: A summary of our existence results. "✓" indicates the type of allocation specified by the column is guaranteed to exist in the setting specified by the row, while "✗" indicates that we give a counterexample and "?" indicates an open question. "id.", "int.", and "add." serve as shorthand for identical valuation, integer valued weights, and additive valuation assumptions respectively.

alized version of EFX for agents with entitlements, cannot be guaranteed in general with counterexamples that are at best 0.786-WEFX for $n = 2$ agents (or 0.795 for $n = 3$), giving a novel upper bound on the problem. We further provide analogous impossibility results for the chore setting on $n = 2$ or 3 agents.

Nevertheless, we showcase that, for the case of two agents, a variant of the "I-cut-you-choose" procedure ensures the WEFX property when dealing with integer-valued weights. Additionally, we introduce a novel procedure for an approximate WEFX solution under normalized weights, employing an algorithm akin to the extensively studied "moving-knife" technique. This approach attains an approximate factor of $\frac{w}{2\sqrt[3]{m}}$, where w represents the highest agent weighting and m is the number of items. Synthesis of the above results demonstrates that the weighted version of the envy-free up to any good (chore) problem is *considerably* more challenging than its unweighted counterpart, even in the most simplistic case of $n = 2$. These results also nicely address an open direction posed in (Chakraborty et al. 2020) on this notion of weighted EFX allocations.

Finally, we extend the work of (Chakraborty et al. 2020) to give a weighted-picking sequence procedure for allocating indivisible chores to weighted agents that is *envy-free up to one chore (1WEF)* and runs in polynomial time. This positive result further reaffirms the gap in complexity between our two relaxed notions of envy-freeness. We here note that independent and in parallel to our work, (Wu, Zhang, and Zhou 2023) also proved such allocations exist and can be computed efficiently using a similar procedure. A summary of our existence and impossibility results is presented in Table 1.

Further Related Work

An extensive line of work exists concerning fair division of indivisible items, and for a complete overview we defer the reader to the surveys of (Chevaleyre, Endriss, and Maudet 2017) and (Markakis 2017). We here focus only a subset of papers that are most relevant to the present work.

Prior works pertaining specifically to the fair allocation of indivisible items to *asymmetric* agents have been predominantly focused on fairness notions not based on envy. (Farhadi et al. 2019) introduced the weighted extension of the *maximin share* (MMS) as studied by (Barman and Krishnamurthy 2020; Budish 2010; Feige, Sapir, and Tauber 2021; Garg, McGlaughlin, and Taki 2019). Moreover, (Aziz, Chan, and Li 2019) explored the weighted MMS (WMMS) concept for fair division of chores (negatively valued goods) where an agent's weight is intuitively their share of the overall workload to be completed. (Babaioff, Nisan, and Talgam-Cohen 2019) examine the competitive equilibrium of agents with asymmetric budgets. More recently, they worked to redefine and further investigate the MMS allocation for the case of agents with arbitrary entitlements (Babaioff, Ezra, and Feige 2021). The authors first note that a WMMS allocation, as presented in (Farhadi et al. 2019), does not align with the intuitive sense of ensuring that highly entitled agents receive a stronger preferential treatment in the allocation process. As such, they define the *AnyPrice Share* (APS) allocation for both the unweighted and weighted contexts and provide a $\frac{3}{5}$ approximate solution.

The recent work of (Aziz, Moulin, and Sandomirskiy 2020) exhibited a polynomial-time algorithm for the computation of an allocation that satisfies both Pareto optimality and *weighted proportionality up to one item* (WPROP1) of goods and chores for agents with asymmetric weights. A large body of work examines unequal agent weightings for divisible items through the notion of *proportionality* (Barbanel 1996; Brams and Taylor 1995; Cseh and Fleiner 2020; Segal-Halevi and Suksompong 2020). While in the unweighted setting, EF1 allocations are also PROP1, the weighted setting does not admit such a luxury (Chakraborty et al. 2020).

Most closely related to our paper is the recent work by Chakraborty et al. (Chakraborty et al. 2020) which considers the allocation problem for indivisible goods with the generalized notion of weighted agents, providing a polynomial time WEF1 alogirthm. We build upon this result and the algorithm for the case of chores, while also providing the first results for both WEFX and XWEF with various existential results for each – answering the question of whether these generalized notions of fair allocations can be computed.

Preliminaries and Basic Definitions

Fair Allocation Problem. An instance of a *weighted* fair allocation problem consists of a set $\mathcal{N}$ of n agents where each $i \in \mathcal{N}$ has weight $w_i > 0$. Let $\mathcal{M}$ be a set of m goods (or chores) with, potentially heterogenous, valuation functions $v_i : 2^m \to \mathbb{R}_{\geq 0}$ (or cost functions $c_i : 2^m \to \mathbb{R}_{\geq 0}$) that are assumed to be additive unless otherwise noted. An allocation of $\mathcal{M}$ is a partition of the set into n disjoint "bundles", $\mathcal{A} = (\mathcal{A}_1, ..., \mathcal{A}_n)$, such that $\bigcup_{i\in[n]} \mathcal{A}_i = \mathcal{M}$ and $\mathcal{A}_i \cap \mathcal{A}_j = \emptyset$ for any two $i, j \in [n]$. For readability, we will henceforth let $\mathcal{A} = (\mathcal{A}_1, ..., \mathcal{A}_n)$ denote an allocation of goods and $\mathcal{B} = (\mathcal{B}_1, ..., \mathcal{B}_n)$ chores.

Fairness Criteria. We seek to find an allocation over the set of goods (or chores) that is *envy-free*: given an instance of a fair division problem and an allocation $\mathcal{A}$, an agent i envies agent j if they strictly prefer the set $\mathcal{A}_j$ over their own bundle $\mathcal{A}_i$ up to a scaling by their weight.

Definition 1. *An allocation is weighted envy-free (WEF) if no agent envies another, i.e. for any pair $i, j \in \mathcal{N}$ of agents we have*

$$\frac{v_i(\mathcal{A}_i)}{w_i} \geq \frac{v_i(\mathcal{A}_j)}{w_j} \text{ for goods, } \frac{c_i(\mathcal{B}_i)}{w_i} \leq \frac{c_i(\mathcal{B}_j)}{w_j} \text{ for chores.}$$

However, envy-freeness is too strong a criteria to meet for both indivisible goods and chores. As such, we define two relaxations of this notion, namely *weighted envy-freeness up to one good or chore* (WEF1 and 1WEF respectively) and *weighted envy-freeness up to any good or chore* (WEFX and XWEF respectively).

Definition 2. *An allocation of goods $\mathcal{A}$ (or chores $\mathcal{B}$) is said to be (i) weighted envy-free up to one good (WEF1) if for any pair of agents $i, j \in \mathcal{N}$ if*

$$\frac{v_i(\mathcal{A}_i)}{w_i} \geq min_{a \in \mathcal{A}_j} \frac{v_i(\mathcal{A}_j \setminus \{a\})}{w_j}$$

(ii) weighted envy-free up to one chore (1WEF) if for any pair of agents $i, j \in \mathcal{N}$ if

$$min_{b \in \mathcal{B}_i} \frac{c_i(\mathcal{B}_i \setminus \{b\})}{w_i} \leq \frac{c_i(\mathcal{B}_j)}{w_j}$$

Definition 3. *An allocation of goods $\mathcal{A}$ (or chores $\mathcal{B}$) is said to be (i) weighted envy-free up to any good (WEFX) if for any pair of agents $i, j \in \mathcal{N}$ if*

$$\frac{v_i(\mathcal{A}_i)}{w_i} \geq \max_{a \in \mathcal{A}_j} \frac{v_i(\mathcal{A}_j \setminus \{a\})}{w_j}$$

(ii) weighted envy-free up to any chore (XWEF) if for any pair of agents $i, j \in \mathcal{N}$ if

$$\max_{b \in \mathcal{B}_i} \frac{c_i(\mathcal{B}_i \setminus \{b\})}{w_i} \leq \frac{c_i(\mathcal{B}_j)}{w_j}$$

While the envy-freeness up to one and any item are very closely related, they define relaxed notions of fairness that have a large discrepancy in terms of complexity. The current literature has demonstrated that a WEF1 allocation always exists for agents with additive valuation functions (Chakraborty et al. 2020), a result that we here extend to the chore division setting with a polynomial time algorithm. However, the problems of the existence of WEFX allocations for goods as well as XWEF for chores remain open. In this paper, we show that WEFX allocations are guaranteed to exist under certain, restrictive, problem assumptions. However, in general, we show that WEFX and XWEF allocations are not guaranteed to exist for agents with additive cost functions, but overcome this inherent barrier with a polynomial time algorithm that yields an *approximate* WEFX allocation for two agents.

Definition 4. *For constants $\alpha \leq 1$ and $\beta \geq 1$, an allocation of goods $\mathcal{A}$ (or chores $\mathcal{B}$) is called (i) α-approximate envy-free up to any good (α-WEFX) if for any $i, j \in \mathcal{N}$*

$$\frac{v_i(\mathcal{A}_i)}{w_i} \geq \alpha \cdot \max_{a \in \mathcal{A}_j} \frac{v_i(\mathcal{A}_j \setminus \{a\})}{w_j}$$

(ii) β-approximate envy-free up to any chore (β-XWEF) if for any $i, j \in \mathcal{N}$

$$\max_{b \in \mathcal{B}_i} \frac{c_i(\mathcal{B}_i \setminus \{b\})}{w_i} \leq \beta \cdot \frac{c_i(\mathcal{B}_j)}{w_j}$$

Weighted Division of Indivisible Goods

We here investigate the issue of fairly distributing a collection of indivisible goods among weighted agents, considering different problem assumptions. While achieving envy-freeness through a complete allocation is the ideal outcome, it is not always easily computed or assured by an algorithm. Consequently, we often turn to the less strict criteria of envy-freeness up to any item. The overall existence of such criteria remains an unresolved challenge in the realm of fair allocation for both unweighted and weighted systems.

In this context, we present certain assumptions that prove adequate for ensuring the existence of WEFX allocations. Moreover, when these assumptions are relaxed, we exhibit the non-existence of such allocations.

Identical Valuations

To begin, we simplify the scenario by initially considering a situation where the valuation functions are the same across all agents. In this setting, we observe that the classic *envy-cycle elimination* method by Lipton produces the desired WEFX allocation with only a minor adjustment. Specifically, the "envy-graph" is modified to have an edge from agent i to agent j if and only if i has *weighted* envy towards j. This gives us the following positive results.

Theorem 5. *The weighted version of (Lipton et al. 2004)'s envy-cycle elimination algorithm produces a complete WEFX allocation for agents with (additive) identical valuation functions.*

Moreover, we can slightly generalize this result to encompass the setting in which agents share a preferential ordering over the set of goods.

Theorem 6. *The weighted version of the envy-cycle elimination algorithm produces a complete WEFX allocation for agents with identical <u>ordinal</u> preferences.*

The proofs of both of these theorems are deferred to the appendix due to space constraints.

The simple model of identical valuations however affords results that are considerably more positive than the general case as exhibited by the below result.

Proposition 7 ((Chakraborty et al. 2020))**.** *The weighted version of the envy-cycle elimination algorithm may not produce a complete WEF1 allocation, even in a problem instance with two agents and additive valuations.*

General Impossibility

For $n = 2$ and 3, it is known that EFX allocations are guaranteed to exist under additive valuation functions (Plaut and Roughgarden 2018; Chaudhury, Garg, and Mehlhorn 2020). However, we here present novel existential results that show the same result does *not* hold in the weighted setting – a considerable deviation between the two problem contexts.

We first sketch the construction for a constant factor approximation upper bound to the WEFX problem on two agents, presenting a complete analysis of the result in the appendix.

Theorem 8. *For $n = 2$ agents with additive valuation functions and weights w_1 and w_2 such that $w_1 + w_2 = 1$, there exists an instance where a complete WEFX allocation of goods does not exist while a 0.786-WEFX allocation does.*

Proof Sketch. Consider an instance of two agents with weights $w_1 = \alpha$ and $w_2 = 1 - \alpha$, and suppose there are $m = 4$ goods (denoted a_i for $1 \leq i \leq 4$) with the following valuation profiles:

	a_1	a_2	a_3	a_4
v_1	0	1	φ	φ
v_2	0	0	1	1

where $\varphi := (1 + \sqrt{5})/2$. We seek to demonstrate that for all the possible allocations of these four items, the largest approximate factor obtainable is 0.786. More specifically, for each allocation we will derive the approximation ratio as a function of the item values and α. After collecting these factors, we will adversarially select α to ensure that all are at most the desired approximate guarantee. The full proof of this result is deferred to the appendix due to space constraints. □

Furthermore, we expand this straightforward construction to the setting of $n = 3$ agents. This extension serves to underscore the significant increase in complexity inherent in pursuing the WEFX objective, particularly when contrasted with its unweighted counterpart. The theorem's proof employs the same analytical approach as previously described and is thus deferred to the appendix.

Theorem 9. *For $n = 3$ agents with additive valuation functions and weights w_1, w_2 and w_3 such that $w_1 + w_2 + w_3 = 1$, there exists an instance where a complete WEFX allocation of goods does not exist while a 0.795-WEFX allocation does.*

As a consequence of these impossibility results, we proceed by designing procedures for the two-agent problem. These procedures yield a WEFX allocation or an approximation to this objective under varied problem assumptions.

Two Agent Procedures

Our first algorithm computes a WEFX allocation for two agents with *integer valued* weights. Formally, we are given two (additive) valuation functions, v_1 and v_2, for a set $\mathcal{M}$ of m items wherein the corresponding agents have weights of 1 and $W \in \mathbb{Z}$. The algorithm involves a modified "I-cut-you-choose" approach, where the agent with the higher weight greedily divides the goods into $W + 1$ bundles. Agent 1 then gets the opportunity to select their preferred bundle from among these. The procedure is presented as pseudocode in Algorithm 1.

Algorithm 1: Integer Weight WEFX Algorithm

Initialize $X_1 \leftarrow \emptyset, X_2 \leftarrow \emptyset$
if $m \leq W$ **then**
  $X_1 \leftarrow \arg\max_{g \in \mathcal{M}} v_1(g)$
  $X_2 \leftarrow \mathcal{M} \setminus X_1$
else
  Initialize $(P_1, ..., P_{W+1}) \leftarrow (\emptyset, ..., \emptyset)$
  while $\mathcal{M} \neq \emptyset$ **do**
    $k = \arg\min_{i \in [W+1]} v_2(P_i)$
    $P_k \leftarrow P_k \cup \{g \in \arg\max_{h \in \mathcal{M}} v_2(h)\}$
  end while
  $X_1 = \arg\max_{k \in [W+1]} v_1(P_k)$
  $X_2 = \mathcal{M} \setminus X_1$
end if
return (X_1, X_2)

Theorem 10. *For $n = 2$ agents with weights 1 and W, along with additive valuation functions, Algorithm 1 computes a WEFX allocation.*

Proof sketch. Due to space constraints, we here give a high-level intuition of the algorithms correctness and defer the full analysis to the appendix.

Given a set of $W + 1$ bundles, agent 1's favorite bundle must necessarily be of value at least the average for the partition which necessarily guarantees the WEFX property. The crux of the analysis is thus verifying that, regardless of the first agent's selection, the second agent remains satisfied. By noting that the greedy partitioning procedure conducted by agent 2 reduces to computing an EFX allocation on $n = W + 1$ identical agents, we can verify that the stronger notion of WEFX is necessarily satisfied. □

Despite this positive result, the inherent upper bound demonstrated in Theorem 8 indicates that adjusting the assumption of agent weightings to be *normalized* instead compels us to design an approximation algorithm. We proceed to present Algorithm 2 which guarantees an *approximate* WEFX allocation for two agents. The algorithm is a modification of the famous moving-knife procedure. Intuitively, the procedure works to minimally satisfy the higher priority agent before reassigning some items to the other agent in an effort to more equitably balance the two and mitigate any envious relationship. This natural algorithm adapted from the unweighted problem setting yields the first approximate WEFX procedure to date. Assuming w is the weight of the higher priority agent, $\alpha = \frac{w}{2\sqrt[3]{m}}$ is our objective approximation factor of WEFX allocation. Our main technical result is stated formally as follows.

Theorem 11. *Algorithm 2 guarantees an $\frac{w}{2\sqrt[3]{m}}$-WEFX allocation for two agents.*

Algorithm 2: Approximate WEFX Algorithm

Input: Two agents with valuation functions v_1 and v_2 with weights w and $1-w$ respectively and set of items $\mathcal{M}$
Output: Allocation $\mathcal{A} = \{\mathcal{A}_1, \mathcal{A}_2\}$
Normalize the valuations, $v_1(\mathcal{M}) = v_2(\mathcal{M}) = 1$
Let $\mathcal{A}_1 = \{g \in \mathcal{M} | v_1(g) \geq v_2(g)\}$, $\mathcal{A}_2 = \mathcal{M} \setminus \mathcal{A}_1$
Sort $\mathcal{A}_1$ in descending, let g_i be the i-th largest item
Let $k = \arg\max_m v_1(\{g_1, \ldots, g_{m-1}\}) \leq w$
Case I: Set $\mathcal{A}_1 = \{g_1\}$, $\mathcal{A}_2 = \mathcal{M} \setminus \mathcal{A}_1$
if $\mathcal{A}$ is α-WEFX **then**
  return $\mathcal{A}$.
end if
Case II: Let $\mathcal{A}_1 = \{g_1, ..., g_{k-1}\}$, $\mathcal{A}_2 = \mathcal{M} \setminus \mathcal{A}_1$
if $\mathcal{A}$ is α-WEFX **then**
  return $\mathcal{A}$.
end if
Case III: Let $\mathcal{A}_1 = \{g_1, ..., g_k\}$, $\mathcal{A}_2 = \mathcal{M} \setminus \mathcal{A}_1$
if $\mathcal{A}$ is α-WEFX **then**
  return $\mathcal{A}$.
end if
Case IV: Let $\mathcal{A}_2 = \arg\max_{g \in B_1} v_2(g)$, $\mathcal{A}_1 = \mathcal{M} \setminus \mathcal{A}_2$

In this algorithm, we first split the items into two bundles $\mathcal{A}_1$ and $\mathcal{A}_2$ based on their normalized valuations for the agents (ie. without loss of generality assume the sum of each agents valuation over all the items is 1). If the first agent has higher value for an item compared to the second agent, the item goes to the first bundle, otherwise it goes to the second bundle. Without loss of generality, we can assume that $v_1(\mathcal{A}_1) \geq w$. Specifically, if $v_1(\mathcal{A}_1) < w$ then we can re-index the two agents so that the inequality holds (since the inequality must hold for the other agent by normalization assumptions on the weights and valuations). We then sort the items in $\mathcal{A}_1$ in descending order based on the valuation function of the first agent, and assume that g_i is the i^{th} highest value item in $\mathcal{A}_1$ for the first agent. Also, assume that k is the maximum number where $v_1(\{g_1, \ldots, g_{k-1}\}) \leq w$ holds. Then, we consider four different allocations as explained in the algorithm, and prove that at least one of these allocations must guarantee α-WEFX.

We proceed to prove the result by analyzing each case independently to obtain inequalities that would necessarily hold and lastly show that all cannot hold at the same time, thus yielding our result. To start, in Case I, we check if we can satisfy the first agent by allocating only one item to her. Hence, the second agent cannot envy this agent after removing the minimum item from $\mathcal{A}_1$. If this does not guarantee an α-WEFX allocation, we can prove the following lemma:

Lemma 12. *If Case I fails, we have:*

$$1/k < \alpha \left(\frac{1 - w/k}{1 - w} \right) \tag{1}$$

Proof. According to the definition of k, $v_1(\{g_1, \ldots, g_k\} \geq w$. Hence, we have $v_1(g_1) > w/k$. This further implies that $v_1(\mathcal{A}_2) < 1 - w/k$. If $\mathcal{A}$ is α-WEFX, we return this allocation. Otherwise, the first agent envies the second even after adding the approximation factor, thus we have:

$$\frac{w/k}{w} < \alpha \left(\frac{1 - w/k}{1 - w} \right)$$

Simplification yields (1). ☐

Let $v_1(\{g_1, \ldots, g_{k-1}\} = w - \beta$. Lemma 13 proves that β is bounded by w/k.

Lemma 13. *If $v_1(\{g_1, \ldots, g_{k-1}\} = w - \beta$, then $0 \leq \beta < w/(k-1)$.*

Proof. By definition, $\beta \geq 0$ necessarily holds. Assume towards contradiction that $\beta \geq w/(k-1)$. This yields $v_1(g_k) > w/(k-1)$. Since the items of $\mathcal{A}_1$ are sorted, we further have that $v_1(g_i) > w/(k-1)$ for every $1 \leq i < k$. Hence, $v_1(\{g_1, \ldots, g_{k-1}\} > w$, a contradiction. ☐

If Case II occurs, an α-WEFX allocation is guaranteed, and the problem is solved. Otherwise, we can prove the following lemma:

Lemma 14. *If Case II fails, we have:*

$$\frac{w - \beta}{w} < \alpha \left(\frac{1 - w + \beta}{1 - w} \right) \tag{2}$$

Proof. In Case II, we allocate the first $k-1$ items of $\mathcal{A}_1$ to the first agent, and we have $v_1(\mathcal{A}_1) = w - \beta$. Since these items come from $\mathcal{A}_1$, by definition we have $v_2(\mathcal{A}_1) \leq w - \beta$. Hence, we have $v_2(\mathcal{A}_2) \geq 1 - w + \beta$. Therefore, if we normalize the allocations based on the weights w and $1-w$, the second agent cannot envy the first agent. If $\mathcal{A}$ is α-WEFX, this is the final allocation, otherwise, the first agent envies the second even after adding the approximation factor, and therefore Inequality (2) holds. ☐

If Case III occurs, an α-WEFX allocation is guaranteed, and the problem is solved. Otherwise, we have the following result:

Lemma 15. *If Case III fails, we have:*

$$\frac{1 - w - w/(k-1) + \beta}{1 - w} < \alpha \left(\frac{w - \beta}{w} \right) \tag{3}$$

Proof. In Case III, we allocate the first k items of $\mathcal{A}_1$ to the first agent, and have $v_1(\mathcal{A}_1) > w$. In this case, if we normalize the allocations based on the weights w and $1-w$, the first agent does not envy the second agent. Furthermore, we have $v_1(\{g_1, \ldots, g_{k-1}\} = w - \beta$, and as we discussed in the proof of Lemma 13, this implies $v_1(g_k) \leq w/(k-1)$. Hence, we can say that $v_1(\mathcal{A}_1) \leq w - \beta + w/(k-1)$. Since the value of every item in $\mathcal{A}_1$ for the first agent is *at least* its value for the second agent, we can say that $v_2(\mathcal{A}_1) \leq w - \beta + w/(k-1)$ and therefore $v_2(\mathcal{A}_2) \geq 1 - w + \beta - w/(k-1)$.

Since for every item in $\mathcal{A}_1$ the value of the first agent is at least the value of the second agent, and given the fact that $v_1(\mathcal{A}_1)$ after removing the minimum item from $\mathcal{A}_1$ (according to the first agent's perspective) is equal to $w - \beta$, we thus argue that $v_2(\mathcal{A}_1)$ is at most $w - \beta$ after removing the minimum item from $v_1(\mathcal{A}_1)$ (according to the second agent's perspective).

If $\mathcal{A}$ is α-WEFX, this is the final allocation, otherwise, the second agent envies the first even after adding the approximation factor, and as a result, Inequality (3) should hold. $\square$

If none of these three cases occur, we must have Case IV. If the allocation is α-WEFX, the problem is solved. Otherwise, the following inequality bound must hold:

Lemma 16. *If Case IV occurs but the allocation is not α-WEFX, we have:*

$$\frac{w/k}{1-w} < \alpha\left(\frac{1-w/k}{w}\right) \tag{4}$$

Proof. In Case IV, we allocate the item with the highest value among all the items in $\mathcal{A}_1$ to the second agent. Since Case III failed, we have $v_2(\{g_1,\dots,g_k\}) > w$. Hence, the value of at least one of the items in $\{g_1,\dots,g_k\}$ is at least w/k for the second agent. Therefore, in the Case IV allocation, $v_2(\mathcal{A}_2) \geq w/k$ and $v_2(\mathcal{A}_1) \leq 1 - w/k$.

If $\mathcal{A}$ is α-WEFX, this is the final allocation, otherwise, the second agent envies the first agent even after adding the approximation factor, and as a result, Inequality (4) should hold. (Since we only allocate one item to the second agent, the first agent does not envy the second after removing the minimum item from $\mathcal{A}_2$.) $\square$

We now have the necessary tools to prove our main theorem which states that Algorithm 2 guarantees an α-WEFX allocation. We show this by proving that Inequalities (1-4) cannot hold at the same time.

Proof of Theorem 11. If Algorithm 2 does not guarantee an α-WEFX allocation, then all the Inequalities (1 - 4) should hold. After simplifying Inequality (1), we have

$$1 - w < \alpha(k - w)$$

Since $0 < \alpha < 1$, we necessarily have $k > 1$. Hence, we can safely use $k-1$ as a denominator in Inequality (3). Also after merging Inequalities (2) and (3), we have:

$$1 - w - w/(k-1) + \beta < \alpha^2(1 - w + \beta)$$

Since we have $0 < \alpha < 1$, the following inequality should hold:

$$1 - w - w/(k-1) < \alpha^2(1-w)$$

Since $k \geq 2$, we can use the above inequality and write:

$$1 - w - 2w/k < \alpha^2(1-w) \tag{5}$$

After simplifying Inequality (4), we have:

$$(w^2/k)(1/\alpha - 1) + w/k + w < 1 \tag{6}$$

We rewrite Inequality 5 using 6 as follows:

$$(w^2/k)(1/\alpha - 1) + w/k + w - w - 2w/k < \alpha^2(1-w)$$

After simplifying the above inequality, we have:

$$w^2 < \alpha^3 k(1-w) + \alpha w(w+1)$$

Since $0 < w < 1$, we can rewrite the above inequality as follows:

$$w^2 < \alpha^3 k + 2\alpha w$$

If we use $\alpha = \frac{w}{2\sqrt[3]{m}}$, we have:

$$w^2 < \frac{w^3 k}{8m} + \frac{w^2}{\sqrt[3]{m}}$$

Since $0 < w < 1$ and $2 \leq k \leq m$, we have:

$$w^2 < \frac{w^2}{8} + \frac{w^2}{\sqrt[3]{2}}$$

which is a contradiction. Hence, all the Inequalities (1 - 4) cannot hold at the same time, proving the main result. $\square$

Weighted Division of Indivisible Chores

We now turn attention to the problem of chore division. Specifically, we present symmetric upper bound results for envy-freeness up to any chore to those of the goods setting before demonstrating that a variant on the WEF1 algorithm of (Chakraborty et al. 2020) yields weighted envy-freeness up to one chore.

General Impossibility

As is the case for the allocation of goods, we cannot in general guarantee that an allocation which is envy-free up to any chore exists for even small problem instances. By slight modification, the construction used for Theorem 8 yields the following non-existence result.

Theorem 17. *For $n = 2$ agents with additive cost functions and weights w_1 and w_2 such that $w_1 + w_2 = 1$, there exists an instance where a complete XWEF allocation of chores does not exist while a 1.272-XWEF allocation does.*

Note that the approximation factor here of 1.272 is equivalent to the reciprocal of the approximation factor for WEFX, as the two problems are nearly symmetric to one another. Lastly, this counterexample construction can be extended to the $n = 3$ setting:

Theorem 18. *For $n = 3$ agents with additive cost functions and weights w_1, w_2 and w_3 such that $w_1 + w_2 + w_3 = 1$, there exists an instance where a complete XWEF allocation of chores does not exist while a 1.214-XWEF allocation does.*

1WEF Algorithm

We lastly present our polynomial time algorithm for weighted envy-freeness up to one chore (1WEF) to compliment the WEF1 algorithm of (Chakraborty et al. 2020). To achieve this, we must devise a proper sequential picking protocol that allows agents to pick their most preferred chores in a predetermined ordering in a similar vein to the traditional *round-robin* procedure from the symmetric agent literature (Budish 2010).

In our general case with unequal weights, we devise a *weight-dependent* picking sequence, in addition to an arbitrary ordering final loop, to yield the desired 1WEF allocation for any number of agents and arbitrary weights. Although the proof of our algorithm is intricate and precise, the algorithm itself is intuitive and computationally efficient.

Algorithm 3: Chore-Division Round-Robin

Input: $\mathcal{L}, (w_1, w_2, ..., w_n), c_i$ for each $i \in \mathcal{N}$
Output: returns 1WEF allocation for the n agents.
Remaining chores $\widehat{\mathcal{L}} \leftarrow \mathcal{L}$
Bundles $B \leftarrow \emptyset, \forall i \in \mathcal{N}$
$t_i \leftarrow 1, \forall i \in \mathcal{N}$
Allocate chores:
while $|\widehat{\mathcal{L}}| > |\mathcal{N}|$ **do**
    $i^* \leftarrow \arg\min_{i \in \mathcal{N}} \frac{t_i}{w_i}$, break ties lexographically
    $l^* \leftarrow \arg\min_{l \in \widehat{\mathcal{L}}} c_{i^*}(l)$
    $B_{i^*} \leftarrow B_{i^*} \cup \{l^*\}$
    $\widehat{\mathcal{L}} \leftarrow \widehat{\mathcal{L}} \setminus l^*$
    $t_{i^*} \leftarrow t_{i^*} + 1$
end while
Allocate Remaining Chores:
for $i \in \mathcal{N}$ **do**
    $l^* \leftarrow \arg\min_{l \in \widehat{\mathcal{L}}} c_i(l)$
    $B_i \leftarrow B_i \cup \{l^*\}$
    $\widehat{\mathcal{L}} \leftarrow \widehat{\mathcal{L}} \setminus l^*$
end for

Theorem 19. *For any number of agents with additive cost functions and arbitrary positive real weights, there exists an algorithm that computes a 1WEF complete allocation in polynomial time.*

The crux of this theorem relies on the carefully constructed picking-sequence of Algorithm 3 so that each agent receives at least one item of higher cost (in the final loop) and the remaining items are allocated based on a weight-adjusted picking frequency for each agent. We claim that every agent is 1WEF up to the chore selected in the final for loop at every iteration of the while loop. We begin by presenting two insightful lemmas concerning the algorithm itself, their proofs are deferred to the appendix due to space constraints.

Lemma 20. *Consider an agent i chosen by Algorithm 3 to pick a chore at some iteration t, and suppose it is not their first pick. Let t_i and t_j denote the number of times agents i and j picked a chore respectively prior to the current iteration. Then $\frac{t_j}{t_i} \geq \frac{w_j}{w_i}$.*

Lemma 20 is sufficient for ensuring that agent i remains weighted envy-free up to the final chore allocated at each iteration of the loops execution. We verify this guarantee this fact using the following:

Lemma 21. *Suppose that, for every iteration t in which agent i picks an item prior to their final pick, the number of times that agents i and j have picked chores (t_i and t_j respectively) satisfy $\frac{t_j}{t_i} \geq \frac{w_j}{w_i}$. Then, in every partial allocation (plus the final chore of the for loop) up to and including i's latest pick, agent j is weighted envy-free up to the item allocated in the final loop of Algorithm 3.*

Using Lemmas 20 and 21, we now have all the necessary facts to prove Theorem 19. Intuitively, we maintain an 1WEF invariant up to the chores that are allocated last, which will necessarily incur the most cost for each agent. This is effectively a reversal of the weighted picking sequence procedure of (Chakraborty et al. 2020) that maintains the invariant of a WEF1 allocation up the *first* allocated item (of largest value to each agent). Though simplistic and intuitive, this extension is non-trivial and relies on a careful analysis of the initial weighted-picking sequence to ensure that our inductive invariant is maintained. Due to space constraints, we defer the reader to Appendix for this intricate analysis, and here present a proof of the algorithm's correctness using the two results depicted above.

Proof of Theorem 19. In the instance that $n > l$, we have that only the for loop of Algorithm 3 runs. Thus each agent is allocated *at most* one chore, thus any pair of agents is 1WEF, and more so XWEF.

Now in the more interesting case of $l > n$, we invoke Lemmas 20 and 21. The combination of the two lemmas gives us that on any iteration when agent i is picking a chore, they will remain 1WEF up to the final chore allocated in the second loop. As this property holds until a complete allocation is achieved, we have that the final result is 1WEF.

For time complexity of the algorithm, we note that the while loop runs exactly $l - n$ iterations. Within this loop, identification of the minimal $\frac{t_i}{w_i}$ takes at $O(n)$ time followed by the picking of the minimal cost chore, which takes $O(l)$ time. Since the first loop only runs in the event $l > n$, the loop runs in $O(l^2)$ time. Lastly, we have the for loop which runs in $O(n)$ iterations, each taking $O(1)$ time. Therefore, the algorithm runs in time $O(l^2)$. □

Conclusions and Future Work

Our work highlights the increased complexity of the fair allocation problem on both goods and chores when agents are assumed to be asymmetric. While strong negative existential results on WEFX (and XWEF) for the case of $n = 2$ and 3 showcase a considerable deviation from the unweighted problem setting, our simple adaptation on the "I-cut-you-choose" and moving-knife procedures provide a novel first step on both exact and approximate guarantees for these challenging problems. Furthermore, the presented intricate analysis of our $\frac{w}{2\sqrt[3]{m}}$-WEFX suggests that the current methods for approximate EFX guarantees will not necessarily translate to this generalized context. For example, Proposition 7 reveals the inadequacy of the envy-cycle procedure for general valuations. Furthermore, our efforts to ensure a 1WEF allocation highlight that the standard round-robin procedure requires specific adjustments to encompass the asymmetric agent configuration. Consequently, the amalgamation of these two approaches, though effective in yielding a $(\varphi - 1)$-EFX allocation (Amanatidis et al. 2021), does not naturally expand to address this extended problem.

We leave the question of the existence of approximate WEFX allocations open as well as the connection between WEFX and the other studied fairness notions for the weighted context for future work. We hope the present work inspires more study on the generalized notion of weighted fair division problems as they will require a plethora of novel techniques to tackle.

Acknowledgements

This work is partially supported by DARPA QuICC NSF AF:Small #2218678, and NSF AF:Small #2114269. Max Springer was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 1840340. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2021. Maximum Nash welfare and other stories about EFX. *Theoretical Computer Science*, 863: 69–85.

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

Aziz, H.; Chan, H.; and Li, B. 2019. Weighted maxmin fair share allocation of indivisible chores. *arXiv preprint arXiv:1906.07602*.

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. 2021. Fair-share allocations for agents with arbitrary entitlements. *arXiv preprint arXiv:2103.04304*.

Babaioff, M.; Nisan, N.; and Talgam-Cohen, I. 2019. Fair Allocation through Competitive Equilibrium from Generic Incomes. In *FAT*, 180.

Barbanel, J. 1996. Game-theoretic algorithms for fair and strongly fair cake division with entitlements. In *Colloquium Mathematicae*, volume 69, 59–73.

Barman, S.; and Krishnamurthy, S. K. 2020. Approximation algorithms for maximin fair division. *ACM Transactions on Economics and Computation (TEAC)*, 8(1): 1–28.

Bouveret, S.; Chevaleyre, Y.; and Maudet, N. 2016. Fair Allocation of Indivisible Goods.

Brams, S. J.; and Taylor, A. D. 1995. An Envy-Free Cake Division Protocol. *The American Mathematical Monthly*, 102(1): 9–18.

Budish, E. 2010. The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. In *BQGT*.

Chakraborty, M.; Suksompong, W.; Zick, Y.; and Igarashi, A. 2020. Weighted Envy-freeness in Indivisible Item Allocation. *ACM Transactions on Economics and Computation (TEAC)*, 9: 1 – 39.

Chaudhury, B. R.; Garg, J.; and Mehlhorn, K. 2020. EFX exists for three agents. In *Proceedings of the 21st ACM Conference on Economics and Computation*, 1–19.

Chevaleyre, Y.; Endriss, U.; and Maudet, N. 2017. Distributed fair allocation of indivisible goods. *Artificial Intelligence*, 242: 1–22.

Cseh, Á.; and Fleiner, T. 2020. The complexity of cake cutting with unequal shares. *ACM Transactions on Algorithms (TALG)*, 16(3): 1–21.

Dehghani, S.; Farhadi, A.; HajiAghayi, M.; and Yami, H. 2018. Envy-free chore division for an arbitrary number of agents. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2564–2583. SIAM.

Farhadi, A.; Ghodsi, M.; Hajiaghayi, M. T.; 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.

Farhadi, A.; Hajiaghayi, M.; Latifian, M.; Seddighin, M.; and Yami, H. 2021. Almost Envy-freeness, Envy-rank, and Nash Social Welfare Matchings. In *AAAI*.

Feige, U.; Sapir, A.; and Tauber, L. 2021. A tight negative example for MMS fair allocations. arXiv:2104.04977.

Garg, J.; McGlaughlin, P. C.; and Taki, S. 2019. Approximating Maximin Share Allocations. In *SOSA*.

Lipton, R.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *EC '04*.

Markakis, E. 2017. Approximation algorithms and hardness results for fair division with indivisible goods. *Trends in Computational Social Choice*, 231–247.

Peterson, E.; and Su, F. 2002. Four-Person Envy-Free Chore Division. *Mathematics Magazine*, 75: 117 – 122.

Peterson, E.; and Su, F. 2009. N-person envy-free chore division. *arXiv: Combinatorics*.

Pikhurko, O. 2000. On Envy-Free Cake Division. *The American Mathematical Monthly*, 107: 736 – 738.

Plaut, B.; and Roughgarden, T. 2018. Almost Envy-Freeness with General Valuations. In *SODA*.

Procaccia, A. D. 2009. Thou Shalt Covet Thy Neighbor's Cake. In *IJCAI*.

Robertson, J.; and Webb, W. 1998. Cake-cutting algorithms - be fair if you can.

Segal-Halevi, E.; and Suksompong, W. 2020. How to cut a cake fairly: A generalization to groups. *The American Mathematical Monthly*, 128(1): 79–83.

Wu, X.; Zhang, C.; and Zhou, S. 2023. Weighted EF1 Allocations for Indivisible Chores. arXiv:2301.08090.