

EFX Exists for Three Agents

BHASKAR RAY CHAUDHURY, Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken Graduate School of Computer Science, Germany

JUGAL GARG, University of Illinois at Urbana-Champaign, USA

KURT MEHLHORN, Max Planck Institute for Informatics, Saarland Informatics Campus, Germany

We study the problem of distributing a set of indivisible items among agents with additive valuations in a *fair* manner. The fairness notion under consideration is Envy-freeness up to *any* item (EFX). Despite significant efforts by many researchers for several years, the existence of EFX allocations has not been settled beyond the simple case of two agents. In this paper, we show constructively that an EFX allocation always exists for three agents. Furthermore, we falsify the conjecture by Caragiannis et al. [9] by showing an instance with three agents for which there is a partial EFX allocation (some items are not allocated) with higher Nash welfare than that of any complete EFX allocation.

CCS Concepts: • Theory of computation → Algorithmic game theory.

Additional Key Words and Phrases: Discrete Fair Division, EFX Allocations, Nash Welfare.

ACM Reference Format:

Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. 2020. EFX Exists for Three Agents. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC '20), July 13–17, 2020, Virtual Event, Hungary*. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3391403.3399511>

1 INTRODUCTION

Discrete fair division of resources is a fundamental problem in many multi-agent settings. Here, the goal is to distribute a set M of m *indivisible* items among n agents in a *fair* manner. Each agent i has a valuation function $v_i: 2^M \rightarrow \mathbb{R}_{\geq 0}$ that quantifies the amount of utility agent i derives from each subset of items. In case of *additive* valuation functions, $v_i(S) := \sum_{j \in S} v_i(\{j\})$, $\forall S \subseteq M$. Let $X = \langle X_1, X_2, \dots, X_n \rangle$ denote a partition of M into n bundles such that X_i is allocated to agent i . Among various choices, *envy-freeness* is the most natural fairness concept, where no agent i envies another agent j 's bundle, i.e., for all agents i, j with $i \neq j$ we have $v_i(X_i) \geq v_i(X_j)$. However, an envy-free allocation does not always exist, e.g., consider allocating a single valuable item among $n \geq 2$ agents. This necessitates the study of relaxed notions of envy-freeness:

Envy-freeness up to one item (EF1): This relaxation was introduced by Budish [7]. An allocation X is said to be EF1 if no agent i envies another agent j after the removal of *some* item in j 's bundle, i.e., $v_i(X_i) \geq v_i(X_j \setminus g)$ for *some* $g \in X_j$. So we allow i to envy j , but the envy must disappear after the removal of some valuable item (according to agent i) from j 's bundle. Note that there is no actual removal: This is simply to assess how agent i values his own bundle when compared to j 's bundle. It is well known that an EF1 allocation always exists, and it can be obtained in polynomial time using the famous envy-cycles procedure by Lipton et al. [26]. However, an EF1 allocation may be unsatisfactory: Intuitively, EF1 insists that envy disappears after the removal of the *most*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '20, July 13–17, 2020, Virtual Event, Hungary

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7975-5/20/07...\$15.00

<https://doi.org/10.1145/3391403.3399511>

valuable item according to the envying agent from the envied agent's bundle—however, in many cases, the most valuable item might be the primary reason for very large envy to exist in the first place. Therefore, stronger notions of fairness are desirable in many circumstances.

Envy-freeness up to any item (EFX): This relaxation was introduced by Caragiannis et al. [10]. An allocation X is said to be EFX if no agent i envies another agent j after the removal of *any* item in j 's bundle, i.e., $v_i(X_i) \geq v_i(X_j \setminus g)$ for all $g \in X_j$. Unlike EF1, in an EFX allocation, the envy between any pair of agents disappears after the removal of the *least valuable item* (according to agent i) from j 's bundle. Note that every EFX allocation is an EF1 allocation, but not *vice-versa*. Consider a simple example of two agents with additive valuations and three items $\{a, b, c\}$ from [13], where the agents valuation for individual items are as follows,

	g_1	g_2	g_3
<i>Agent 1</i>	1	1	2
<i>Agent 2</i>	1	1	2

Observe that g_3 is twice as valuable than g_1 or g_2 for both agents. An allocation where one agent gets $\{g_1\}$ and the other gets $\{g_2, g_3\}$ is EF1 but not EFX. The only possible EFX allocation is where one agent gets $\{g_3\}$ and the other gets $\{g_1, g_2\}$, which is clearly fairer than the given EF1 allocation. This example also shows how EFX helps to rule out some unsatisfactory EF1 allocations. Caragiannis et al. [9] remark that

“Arguably, EFX is the best fairness analog of envy-freeness of indivisible items.”

While an EF1 allocation is always guaranteed to exist, very little is known about the existence of EFX allocations. Caragiannis et al. [10] state that

“Despite significant effort, we were not able to settle the question of whether an EFX allocation always exists (assuming all items must be allocated), and leave it as an enigmatic open question.”

Plaut and Roughgarden [27] show two scenarios for which EFX allocations are guaranteed to exist: (i) All agents have identical valuations (i.e., $v_1 = v_2 = \dots = v_n$), and (ii) Two agents (i.e., $n = 2$). Unfortunately, starting from three agents, even for the well studied class of *additive valuations*, it is open whether EFX allocations exist. Plaut and Roughgarden [27] also remark that:

“The problem seems highly non-trivial even for three players with different additive valuations.”

Furthermore, it is also suspected in [27] that EFX allocations may not exist in general settings:

“We suspect that at least for general valuations, there exist instances where no EFX allocation exists, and it may be easier to find a counterexample in that setting.”

Contrary to this suspicion, we show that

THEOREM. *EFX allocations always exist for three agents with additive valuations.*

EFX with charity: Quite recently there have been studies [9, 13] that consider relaxations of EFX, called “EFX with charity”. Here we look for partial EFX allocations, where not all items need to be allocated (some of them remain unallocated). There is a trivial such allocation where no item is allocated to any agent. Therefore, the goal is to determine allocations with some *qualitative* or *quantitative* bound on the set of unallocated items. For instance, Chaudhury et al. [13] show how to determine a partial EFX allocation X and a pool of unallocated items P such that no agent envies the pool (i.e. for any agent i , we have $v_i(X_i) \geq v_i(P)$), and P has less than n items (i.e., $|P| < n$), even in the case of *general valuations*. In case of additive valuations, Caragiannis et al. [9] show the

existence of a partial EFX allocation $X = \langle X_1, X_2, \dots, X_n \rangle$, where every agent gets at least half the value of his bundle in the allocation that maximizes the *Nash welfare* i.e., the geometric mean of agents' valuations (suggesting that unallocated items are not too valuable).

The Nash welfare of a fair allocation is often considered as a measure of its *efficiency* [9]: Intuitively, it captures how much *average welfare* the allocation achieves while still remaining fair. The result of Caragiannis et al. [9] imply that there are *efficient* partial EFX allocations (partial EFX allocations with a 2-approximation of the maximum possible Nash welfare). Indeed, it is a natural question to ask whether there are complete EFX allocations (all items are allocated) with good efficiency. To this end, Caragiannis et al. [9] conjecture:

CONJECTURE 1. *“In particular, we suspect that adding an item to an allocation problem (that provably has an EFX allocation) yields another problem that also has an EFX allocation with at least as high Nash welfare as the initial one.”¹*

If this conjecture is true, it implies the existence of an efficient complete EFX allocation. We show (in Section 5) that

The above conjecture is false.

To disprove the conjecture we exhibit an instance where there exists a partial EFX allocation with higher Nash welfare than the Nash welfare of any complete EFX allocation. This also highlights an inherent barrier in the current techniques to determining EFX allocations: Several of the existing algorithms for approximate EFX allocations ([27]) and EFX allocations with charity ([13]) start with an inefficient partial EFX allocation and make it more efficient iteratively by cleverly allocating some of the unallocated items and unallocating some of the allocated items. However, our instance in Section 5 shows that such approaches will not help if our goal is to determine a complete EFX allocation.

A large chunk of our work in this paper develops better tools to overcome this particular barrier, and we consider the tools introduced here as the most innovative technical contribution of our work. We also feel that these tools and the instance may help resolving the major open problem of the existence of EFX allocations for more than three agents and more general valuations (positively or negatively).

1.1 Our Contributions

Our major contribution in this paper is to prove that an EFX allocation always exists when there are three agents with additive valuations. The proof is algorithmic. To discuss our techniques, we first briefly highlight how we overcome two barriers in the current techniques.

Splitting bundles: We first sketch the simple algorithm of Plaut and Roughgarden [27] that determines an EFX allocation when all agents have the same valuation function, say v . Let us restrict our attention to the special case where there is no zero marginals, i.e., for any $S \subseteq M$ and $g \notin S$ we have $v(S \cup g) > v(S)$. Also, note that since agents have the same valuation function, if $v(X_i) < v(X_j \setminus g)$ for two agents i and j for some $g \in X_j$ then we have $v(X_{i_{min}}) < v(X_j \setminus g)$ where i_{min} is the agent with the lowest valuation. The algorithm in [27] starts off with an arbitrary allocation (not necessarily EFX), and as long as there are agents i and j such that $v(X_i) < v(X_j \setminus g)$ for some $g \in X_j$, the algorithm takes the item g away from j (j 's new bundle is $X_j \setminus g$) and adds it to i_{min} 's bundle (i_{min} 's new bundle is $X_{i_{min}} \cup g$). Also, note that after re-allocation the only changed bundles are that of i_{min} and j , and both of them have valuations still higher than i_{min} 's initial valuation: $v(X_{i_{min}} \cup g) > v(X_{i_{min}})$ and $v(X_j \setminus g) > v(X_j)$. Observe that such an operation increases

¹This was posed as a monotonicity conjecture in their presentation at EC'19.

the valuation of an agent with the lowest valuation. Thus, after finitely many applications of this re-allocation we must arrive at an EFX allocation. Note that this crucially uses the fact that the agents have identical valuations. In the general case, the valuation of agent j may drop significantly after removing g and j 's current valuation may be even less than i_{\min} 's initial valuation. *Therefore, it is important to understand how agents value item(s) that we move across the bundles.* To this end, we carefully split every bundle into *upper* and *lower* half bundles (see (1) in Section 2). We systematically quantify the agents' relative valuations for these upper and lower half bundles. In most cases, we can move these bundles from one agent to another, improving the valuation of some of the agents while still guaranteeing the EFX property. This idea is detailed in Sections 3 and 4.

A new potential function: We need to show that there is progress after every swap of half bundles. The typical method here is to show improvement of the valuation vector on the Pareto front (see [13] and [27]). However, there are limitations to this approach: In particular, we show an instance and a partial EFX allocation such that the valuation vector of any complete EFX allocation does not Pareto dominate the valuation vector of the existing partial EFX allocation. To overcome this barrier, we first pick an arbitrary agent a at the beginning and show that whenever we are unable to improve the valuation vector on the Pareto front, we can strictly increase a 's valuation. In other words, the valuation of a particular agent a never decreases throughout re-allocations, and it improves after finitely many re-allocations, showing convergence. A more elaborate discussion on this technique is presented in Section 2.

1.2 Further Related Work

Fair division has received significant attention since the seminal work of Steinhaus [28] in the 1940s, where he introduced the cake cutting problem among $n > 2$ agents. Perhaps the two most crucial notions of fairness properties that can be guaranteed in case of divisible items are *envy-freeness* and *proportionality*. In a proportional allocation, each agent gets at least a $1/n$ share of all the items. In case of indivisible items, as mentioned earlier, none of these two notions can be guaranteed. While EF1 and EFX are fairness notions that relax envy-freeness, the most popular notion of fairness that relaxes proportionality for indivisible items is *maximin share* (MMS), which was introduced by Budish [7]. While MMS allocations do not always exist [24], but there has been extensive work to come up with approximate MMS allocations [1, 4, 6, 7, 19–21, 24].

While much research effort goes into finding fair allocations, there has also been a lot of interest in guaranteeing *efficient* fair allocations. A standard notion of efficiency is *Pareto-optimality*². Caragiannis et al. [10] showed that any allocation that has the maximum Nash welfare is guaranteed to be Pareto-optimal (efficient) and EF1 (fair). Therefore, the Nash welfare of an allocation is also considered as a measure of efficiency and fairness of an allocation. However, finding an allocation with the maximum Nash welfare is APX-hard [25], and its approximation has received a lot of attention recently, e.g., [2, 3, 5, 11, 14–17]. Barman et al. [5] give a pseudopolynomial algorithm to find an allocation that is both EF1 and Pareto-optimal. Other works try to approximate MMS with Pareto-optimality [18] or explore relaxations of EFX with high Nash welfare [9].

Applications: There are several real-world scenarios where resources need to be divided fairly and efficiently, e.g., splitting rent among tenants, dividing inheritance property in a family, splitting taxi fares among riders, and many more. One example of fair division techniques used in practice is Spliddit (<http://www.spliddit.org>). Since its launch in 2014, Spliddit has attracted several thousands of users [10]. For more details on Spliddit, we refer the reader to [22, 27]. Another example is

²An allocation $X = \langle X_1, \dots, X_n \rangle$ is Pareto-optimal if there is no allocation $Y = \langle Y_1, \dots, Y_n \rangle$ where $v_i(Y_i) \geq v_i(X_i)$ for all $i \in [n]$ and $v_j(Y_j) > v_j(X_j)$ for some j .

Course Allocate, which is used by the Wharton School at the University of Pennsylvania to fairly allocate 350 courses to 1700 MBA students [8, 27]. Kurokawa et al. [23] used *leximin* fairness to allocate unused classrooms in public schools to charter schools in California. The best part of the allocations determined in all these applications is that they yield results that not only *seem* fair on most instances, but also come with mathematical guarantees.

2 PRELIMINARIES AND TECHNICAL OVERVIEW

An instance I of fair allocation problem is a triple $\langle [3], M, \mathcal{V} \rangle$, where we have three agents 1, 2, and 3, a set M of m indivisible items (or goods), and a set of valuation functions $\mathcal{V} = \{v_1, v_2, v_3\}$, where each $v_i: 2^M \rightarrow \mathbb{R}_{\geq 0}$ captures the utility agent i has for all the different subsets of goods that can be allocated. We assume that the valuation functions are *additive* ($v_i(S) = \sum_{g \in S} v_i(\{g\})$) and *normalized* ($v_i(\emptyset) = 0$). For ease of notation, we write $v_i(g)$ for $v_i(\{g\})$. Further, we write $S \oplus_i T$ for $v_i(S) \oplus v_i(T)$ with $\oplus \in \{\leq, \geq, <, >\}$. Given an allocation $X = \langle X_1, X_2, \dots, X_n \rangle$ we say that i *strongly envies* a bundle $S \subseteq M$ if $X_i <_i S \setminus g$ for some $g \in S$, and we say that i *weakly envies* S if $X_i <_i S$ but $X_i \geq_i S \setminus g$ for all $g \in S$. From this perspective an allocation is an EFX allocation if and only if no agent strongly envies another agent.

Non-degenerate instances: We call an instance $I = \langle [3], M, \mathcal{V} \rangle$ non-degenerate if and only if no agent values two different sets equally, i.e., $\forall i \in [3]$ we have $v_i(S) \neq v_i(T)$ for all $S \neq T$. We first show that it suffices to deal with non-degenerate instances. Let $M = \{g_1, g_2, \dots, g_m\}$. We perturb any instance I to $I(\varepsilon) = \langle [3], M, \mathcal{V}(\varepsilon) \rangle$, where for every $v_i \in \mathcal{V}$ we define $v'_i \in \mathcal{V}(\varepsilon)$, as

$$v'_i(g_j) = v_i(g_j) + \varepsilon 2^j.$$

LEMMA 2. Let $\delta = \min_{i \in [3]} \min_{S, T: v_i(S) \neq v_i(T)} |v_i(S) - v_i(T)|$ and let $\varepsilon > 0$ be such that $\varepsilon \cdot 2^{m+1} < \delta$. Then

- (1) For any agent i and $S, T \subseteq M$ such that $v_i(S) > v_i(T)$, we have $v'_i(S) > v'_i(T)$.
- (2) $I(\varepsilon)$ is a non-degenerate instance. Furthermore, if $X = \langle X_1, X_2, X_3 \rangle$ is an EFX allocation for $I(\varepsilon)$ then X is also an EFX allocation for I .

PROOF. For the first statement of the lemma, observe that

$$\begin{aligned} v'_i(S) - v'_i(T) &= v_i(S) - v_i(T) + \varepsilon \left(\sum_{g_j \in S \setminus T} 2^j - \sum_{g_j \in T \setminus S} 2^j \right) \\ &\geq \delta - \varepsilon \sum_{g_j \in T \setminus S} 2^j \\ &\geq \delta - \varepsilon \cdot (2^{m+1} - 1) \\ &> 0 \end{aligned}$$

For the second statement of the lemma, consider any two sets $S, T \subseteq M$ such that $S \neq T$. Now, for any $i \in [3]$, if $v_i(S) \neq v_i(T)$, we have $v'_i(S) \neq v'_i(T)$ by the first statement of the lemma. If $v_i(S) = v_i(T)$, we have $v'_i(S) - v'_i(T) = \varepsilon(\sum_{g_j \in S \setminus T} 2^j - \sum_{g_j \in T \setminus S} 2^j) \neq 0$ (as $S \neq T$). Therefore, $I(\varepsilon)$ is non-degenerate.

For the final claim, let us assume that X is an EFX allocation in $I(\varepsilon)$ and not an EFX allocation in I . Then there exist i, j , and $g \in X_j$ such that $v_i(X_j \setminus g) > v_i(X_i)$. In that case, we have $v'_i(X_j \setminus g) > v'_i(X_i)$ by the first statement of the lemma, implying that X is not an EFX allocation in $I(\varepsilon)$ as well, which is a contradiction. \square

From now on we only deal with non-degenerate instances. In non-degenerate instances, all goods have positive value for all agents.

Overall approach: An allocation X' *Pareto dominates* an allocation X if $v_i(X_i) \leq v_i(X'_i)$ for all i with strict inequality for at least one i . The existing algorithms for “EFX with charity” [13] or “approximate EFX allocations” [27] construct a sequence of EFX allocations in which each allocation Pareto dominates its predecessor. However we exhibit in Section 5 a partial EFX allocation that is not Pareto dominated by any complete EFX allocation. Thus we need a more flexible approach in the search for a complete EFX allocation.

We name the agents a , b , and c arbitrarily and consider the lexicographic ordering of the triples

$$\phi(X) = (v_a(X_a), v_b(X_b), v_c(X_c)),$$

i.e., $\phi(X) <_{\text{lex}} \phi(X')$ (X' *dominates* X) if (i) $v_a(X_a) < v_a(X'_a)$ or (ii) $v_a(X_a) = v_a(X'_a)$ and $v_b(X_b) < v_b(X'_b)$ or (iii) $v_a(X_a) = v_a(X'_a)$ and $v_b(X_b) = v_b(X'_b)$ and $v_c(X_c) < v_c(X'_c)$. We construct a sequence of allocations in which each allocation dominates its predecessor. Of course, if X' Pareto dominates X , then it also dominates X , so we can use all the update rules in [13].

Our goal then is to iteratively construct a sequence of EFX allocations such that each EFX allocation dominates its predecessor.

Most envious agent: We use the notion of a most envious agent, introduced in [13]. Consider an allocation X , and a set $S \subseteq M$ that is envied by at least one agent. For an agent i such that $S >_i X_i$, we “measure the envy” that agent i has for S by $\kappa_X(i, S)$, where $\kappa_X(i, S)$ is the size of a smallest subset of S that i still envies, i.e., $\kappa_X(i, S)$ is the smallest cardinality of a subset S' of S such that $S' >_i X_i$. Thus, the smaller the value of $\kappa_X(i, S)$, the greater the envy of agent i for the set S . So let $\kappa_X(S) = \min_{i \in [3]} \kappa_X(i, S)$. Naturally, we define the set of the *most envious agents* $A_X(S)$ for a set S as the set of agents with smallest values of $\kappa_X(i, S)$, i.e.,

$$A_X(S) = \{i \mid S >_i X_i \text{ and } \kappa_X(i, S) = \kappa_X(S)\}.$$

The following simple observation about the most envious agents of specific kinds of bundles will be useful.

OBSERVATION 3. *Given any allocation X , and an unallocated good g , for any $i \in [3]$, $A_X(X_i \cup g) \neq \emptyset$.*

PROOF. It suffices to prove that there exists at least one agent who strictly prefers $X_i \cup g$ over his own bundle in allocation X . This is guaranteed since we are dealing with non-degenerate instances, in which $X_i \cup g >_i X_i$. \square

Champions and Champion Graph M_X : Let X be the partial EFX allocation at any stage in our algorithm, and let g be an unallocated good. We say that i *champions* j (w.r.t g) if i is a most envious agent for $X_j \cup g$, i.e., $i \in A_X(X_j \cup g)$. We define the *champion graph* M_X , where each vertex corresponds to an agent and there is a directed edge $(i, j) \in M_X$ if and only if i champions j .

OBSERVATION 4. *The champion graph M_X is cyclic.*

PROOF. By Observation 3, we have that the set of champions of any agent is never empty. Therefore, every vertex in M_X has at least one incoming edge. Thus M_X is cyclic. \square

If i champions j , we define G_{ij} as a largest cardinality subset of $X_j \cup g$ such that $(X_j \cup g) \setminus G_{ij} >_i X_i$. Since the valuations are additive, note that such a subset can be identified efficiently as the set K of the k least valuable goods for i in $X_j \cup g$ such that $(X_j \cup g) \setminus K >_i X_i$ and k is maximum. Now we make some small observations.

OBSERVATION 5. *Assume i champions j .*

- (1) *We have $((X_j \cup g) \setminus G_{ij}) \setminus h \leq_k X_k$ for all $h \in (X_j \cup g) \setminus G_{ij}$ and all agents k including i .*
- (2) *If agent k does not champion j , we have $(X_j \cup g) \setminus G_{ij} \leq_k X_k$.*

PROOF. Note that by definition, G_{ij} is a largest cardinality subset of $X_j \cup g$ such that i values $(X_j \cup g) \setminus G_{ij}$ more than X_i . This implies that $(X_j \cup g) \setminus G_{ij}$ is a smallest cardinality subset of $X_j \cup g$ that i values more than X_i . Thus $|(X_j \cup g) \setminus G_{ij}| = \kappa_X(i, X_j \cup g)$. Since i champions j , we have that $i \in A_X(X_j \cup g)$ and thus $\kappa_X(i, X_j \cup g) = \kappa_X(X_j \cup g)$. Now, no agent k values a subset of $X_j \cup g$ of size less than $\kappa_X(k, X_j \cup g)$ more than X_k . Note that $((X_j \cup g) \setminus G_{ij}) \setminus h$ has size $\kappa_X(X_j \cup g) - 1 < \kappa_X(k, X_j \cup g)$ and, thus, $((X_j \cup g) \setminus G_{ij}) \setminus h \leq_k X_k$.

Now if k did not champion j then $\kappa_X(k, X_j \cup g) < \kappa_X(X_j \cup g)$. Thus, $|(X_j \cup g) \setminus G_{ij}| = \kappa_X(X_j \cup g) < \kappa_X(k, X_j \cup g)$. Since k values any subset of $X_j \cup g$ of size less than $\kappa_X(k, X_j \cup g)$ at most X_k , we have $(X_j \cup g) \setminus G_{ij} \leq_k X_k$. \square

We next mention two cases where it is known how to obtain a Pareto dominating EFX allocation from an existing EFX allocation. For an allocation X , we define the *envy graph* E_X , whose vertices represent agents, and in which there is a directed edge from i to j if i envies j , i.e., $X_j >_i X_i$. We can assume without loss of generality (w.l.o.g.) that E_X is acyclic.

FACT 6 ([26]). *Let $X = \langle X_1, X_2, X_3 \rangle$ be an EFX allocation. Then there exists another EFX allocation $Y = \langle Y_1, Y_2, Y_3 \rangle$, where for all $i \in [3]$, $Y_i = X_j$ for some $j \in [3]$, such that E_Y is acyclic and $\phi(Y) \succeq_{lex} \phi(X)$ (because Y Pareto dominates X).*

OBSERVATION 7 ([13]). *Consider an EFX allocation X . Let s be any agent and let g be an unallocated good. If i champions s and i is reachable from s in E_X , then there is an EFX allocation Y Pareto dominating X . Additionally, agent s is strictly better off in Y , i.e., $Y_s >_s X_s$.*

PROOF. We have that i is reachable from s in E_X . Let $t_1 \rightarrow t_2 \rightarrow \dots \rightarrow t_k$ be the path from $t_1 = s$ to $t_k = i$ in E_X . We determine a new allocation Y as follows:

$$\begin{aligned} Y_{t_j} &= X_{t_{j+1}} && \text{for } j \in [k-1] \\ Y_i &= (X_s \setminus G_{is}) \cup g \\ Y_\ell &= X_\ell && \text{for all other } \ell \end{aligned}$$

Note that every agent along the path has strictly improved his valuation: Agents t_1 to t_{k-1} got bundles they envied in E_X and agent i championed s and got $(X_s \setminus G_{is} \cup g)$, which is more valuable to i than X_i (by definition of G_{is}). Also, every other agent retained their previous bundles and thus their valuations are not lower than before. Thus $\phi(Y) >_{lex} \phi(X)$ and also $Y_s >_s X_s$ (s was an agent along the path). It only remains to argue that Y is EFX. To this end, consider any two agents j and j' . We wish to show that j does not strongly envy j' in Y .

Case $j' \neq i$: Note that $Y_{j'} = X_\ell$ for some $\ell \in [3]$ (j' either received a bundle of another agent when we shifted the bundles along the path or retained the previous bundle). Also, note that $Y_j \geq_j X_j$ (no agent is worse off in Y). Therefore, $Y_j \geq_j X_j \geq_j X_\ell \setminus h =_j Y_{j'} \setminus h$ for all $h \in Y_{j'}$ (j did not strongly envy ℓ in X as X was EFX).

Case $j' = i$: We have $Y_{j'} = (X_s \setminus G_{is}) \cup g$. Since i championed s , by Observation 5 (part 1) we have that $((X_s \setminus G_{is}) \cup g) \setminus h \leq_j X_j$. Like earlier, $Y_j \geq_j X_j$ (no agent is worse off in Y). Thus j does not strongly envy i . \square

Observation 7 implies that if there is some unallocated good and (i) if the envy graph E_X has a single source³ or (ii) any agent champions himself then there is a strictly Pareto dominating EFX allocation.

³A source is a vertex in E_X with in-degree zero.

COROLLARY 8. *Let X be an EFX allocation, and g be an unallocated good. If E_X has a single source s , or M_X has a 1-cycle involving agent s , then there is an EFX allocation Y that Pareto dominates X in which $Y_s >_s X_s$.*

PROOF. If E_X has a single source s , the champion of s (which always exist, by Observation 3) is reachable from s . If M_X has a 1-cycle involving agent s then again the champion of s (which is s itself) is reachable from s . In both cases, since the champion of s is reachable from s in the envy graph E_X , there is a Pareto dominating allocation Y such that $Y_s >_s X_s$ by Observation 7. \square

Hence, starting from Section 3, we only discuss the cases where the envy-graph has more than one source and there are no self-champions.

We start with some simple yet crucial observations.

OBSERVATION 9. *If i champions j and $X_i \geq_i X_j$, then $g \notin G_{ij}$, $G_{ij} \subseteq X_j$, and $G_{ij} <_i g$.*

PROOF. We have $i \in A_X(X_j \cup g)$. Since $g \notin X_j$, $G_{ij} \subseteq X_j \cup g$, and valuations are additive and we have that $v_i((X_j \cup g) \setminus G_{ij}) = v_i(X_j) + v_i(g) - v_i(G_{ij})$. Again since $i \in A_X(X_j \cup g)$, by the definition of G_{ij} , $(X_j \cup g) \setminus G_{ij} >_i X_i$, and hence, $v_i(X_i) < v_i(X_j) + v_i(g) - v_i(G_{ij})$. Now we have $X_i \geq_i X_j$, implying that $G_{ij} <_i g$, and therefore, $g \notin G_{ij}$. \square

Observation 9 tells us that if i champions j , and i does not envy j , then $G_{ij} \subseteq X_j$. Therefore, we can split the bundle of agent j into two parts G_{ij} and $X_j \setminus G_{ij}$. We refer to G_{ij} as the *lower-half bundle* of j , and to $X_j \setminus G_{ij}$ as the *upper-half bundle* of j , and visualize the bundle of agent j as

$$X_j = \begin{array}{c|c} X_j \setminus G_{ij} \\ \hline G_{ij} \\ \hline (j) \end{array} \quad \text{if } i \text{ champions } j \text{ and } i \text{ does not envy } j. \quad (1)$$

We collect some more facts about the values of lower and upper half bundles.

OBSERVATION 10. *If i champions j and j does not champion himself (self-champion), then we have $G_{ij} \neq \emptyset$ and $G_{ij} \geq_j g$.*

PROOF. Since j does not self-champion, by Observation 5 (part 2), we have that $(X_j \cup g) \setminus G_{ij} \leq_j X_j$. Since $g \notin X_j$ and $G_{ij} \subseteq X_j \cup g$ we have $v_j((X_j \cup g) \setminus G_{ij}) = v_j(X_j) + v_j(g) - v_j(G_{ij}) \leq v_j(X_j)$, implying that $G_{ij} \geq_j g$. Since the value of g for j is non-zero, G_{ij} is non-empty. \square

OBSERVATION 11. *Let i champion j , and $X_i \geq_i X_j$. Let i' champion k and $X_{i'} \geq_{i'} X_k$. If i does not champion k , then $X_j \setminus G_{ij} >_i X_k \setminus G_{i'k}$.*

PROOF. Since $i \in A_X(X_j \cup g)$ and $X_i \geq_i X_j$, by Observation 9, we have $g \notin G_{ij}$. Thus, $G_{ij} \subseteq X_j$. By the same reasoning, $g \notin G_{i'k}$ and $G_{i'k} \subseteq X_k$. Therefore, $(X_j \cup g) \setminus G_{ij} = (X_j \setminus G_{ij}) \cup g$, and $(X_k \cup g) \setminus G_{i'k} = (X_k \setminus G_{i'k}) \cup g$. By the definition of G_{ij} , we have $(X_j \setminus G_{ij}) \cup g >_i X_i$. Since $i \notin A_X(X_k \cup g)$, we have $X_i \geq_i (X_k \setminus G_{i'k}) \cup g$ by Observation 5 (part 2). Combining the two inequalities, we have $(X_j \setminus G_{ij}) \cup g >_i (X_k \setminus G_{i'k}) \cup g$, which implies $X_j \setminus G_{ij} >_i X_k \setminus G_{i'k}$. \square

In the upcoming sections, we show how to derive a dominating EFX allocation from an existing EFX allocation. Corollary 8 already deals with the cases that E_X has a single source or M_X has a 1-cycle. *We proceed under the following general assumptions: E_X is cycle-free and has at least two sources and there is no 1-cycle in M_X .* We distinguish the remaining cases by the number of sources in E_X .

3 EXISTENCE OF EFX: THREE SOURCES IN E_X

If E_X has three sources, the allocation X is envy-free, i.e., $X_i \geq_i X_j$ for all i and j . We make a case distinction by whether or not M_X contains a 2-cycle.

3.1 2-cycle in M_X

Assume without loss of generality that agent 2 champions agent 1 and agent 1 champions agent 2. Since $X_1 \geq_1 X_2$ and $X_2 \geq_2 X_1$, the bundles X_1 and X_2 decompose according to (1). Since neither 1 nor 2 self-champion (as M_X has no 1-cycle), by Observation 11, we have $X_2 \setminus G_{12} >_1 X_1 \setminus G_{21}$ and $X_1 \setminus G_{21} >_2 X_2 \setminus G_{12}$. We swap the upper-halves of X_1 and X_2 to obtain

$$X' = \begin{array}{c} \boxed{X_2 \setminus G_{12}} \\ \hline \boxed{G_{21}} \end{array} \quad (1) \quad \begin{array}{c} \boxed{X_1 \setminus G_{21}} \\ \hline \boxed{G_{12}} \end{array} \quad (2) \quad \boxed{X_3} \quad (3)$$

Note that agent 3 has the same valuation as before, while 1 and 2 are strictly better off. If X' is EFX we are done. So assume otherwise. We first determine the potential strong envy edges.

- *From 1:* We replaced the more valuable (according to 1) $X_2 \setminus G_{12}$ in X_2 with the less valuable $X_1 \setminus G_{21}$ and left X_3 unchanged. Thus 1 is strictly better off and according to him, the valuations of the bundles of 2 and 3 in X' is at most the valuation of their bundles in X . As 1 did not envy 2 and 3 before in X , 1 does not envy 2 and 3 in X' .
- *From 2:* A symmetrical argument shows that 2 does not envy 1 and 3.
- *From 3:* For agent 3, the sum of the valuations of agents 1 and 2 has not changed by the swap and 3 envied neither 1 nor 2 before the swap. Thus 3 envies at most one of the agents 1 and 2 after the swap. Assume without loss of generality that he envies agent 2. We then replace the lower-half bundle of agent 2 (G_{12}) with g to obtain

$$X'' = \begin{array}{c} \boxed{X_2 \setminus G_{12}} \\ \hline \boxed{G_{21}} \end{array} \quad (1) \quad \begin{array}{c} \boxed{X_1 \setminus G_{21}} \\ \hline \boxed{g} \end{array} \quad (2) \quad \boxed{X_3} \quad (3)$$

In X'' , agent 2 is still strictly better off than in X since by the definition of G_{21} , we have $(X_1 \setminus G_{21}) \cup g >_2 X_2$. Thus, X'' Pareto dominates X . We still need to show that X'' is EFX. To this end, observe that as we have not changed the bundles of agents 1 and 3, there is no strong envy between them. So we only need to exclude strong envy edges to and from agent 2.

- *Nobody strongly envies agent 2:* Note that 2 championed 1. Thus, $((X_1 \setminus G_{21}) \cup g) \setminus h \leq_1 X_1$ and $((X_1 \setminus G_{21}) \cup g) \setminus h \leq_3 X_3$ for all $h \in (X_1 \setminus G_{21}) \cup g$ by Observation 5 (part 1). Since both 1 and 3 are not worse off than before, they do not strongly envy 2.
- *Agent 2 does not envy anyone:* We have that $(X_1 \setminus G_{21}) \cup g >_2 X_2$. Also according to 2, the valuation of the current bundles of 1 and 3 is at most their previous one, and 2 did not envy them before (when he had X_2). Hence, 2 does not envy 1 and 3.

We have thus shown that X'' is EFX and Pareto dominates X . Actually, the strategy described above handles a more general situation. It yields a Pareto dominating EFX allocation as long as 3 envies neither 1 nor 2 initially, even if 1 and 2 envied (*not strongly envied*) 3 initially:

REMARK 12. *Let X be an EFX allocation, and let g be an unallocated good. If M_X has a 2-cycle, say involving agents 1 and 2, and agent 3 envies neither 1 nor 2, then there exists an EFX allocation Y Pareto dominating X .*

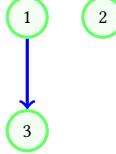


Fig. 1. Envy Graph for two sources when $(2,3) \notin E_X$: Green nodes correspond to the agents. Blue edges are the edges in E_X .

Remark 12 will be helpful when we deal with certain instances where E_X has two sources later in Section 4.

We now consider the case when M_X has no two cycle. Since M_X is cyclic and we neither have a 1-cycle nor a 2-cycle, we must have a 3-cycle. We can show by a very similar but more involved approach that there exists an EFX allocation Y that Pareto dominates X . The complete proof can be found in Subsection 3.2 in the full version of the paper [12]. Thus, we arrive at the following lemma,

LEMMA 13. *Let X be a partial EFX allocation and g be an unallocated good. If E_X has three sources, then there is an EFX allocation Y Pareto dominating X .*

4 EXISTENCE OF EFX: TWO SOURCES IN E_X

Let us assume that agents 1 and 2 are the sources, and let $(1,3) \in E_X$. We have two configurations for E_X now, depending on whether or not $(2,3) \in E_X$. If $(2,3) \in E_X$, it is relatively straightforward to determine a new EFX allocation Pareto dominating X . Agent 3 is reachable from both 1 and 2 in E_X , and hence, if 3 champions either 1 or 2, we have a Pareto dominating EFX allocation by Observation 7. If 3 champions neither 1 nor 2, 1 and 2 must be champions of each other (Recall that no agent self-champions). Also note that 3 envies neither 1 nor 2. Therefore, by Remark 12, we have a Pareto dominating EFX allocation.

From now on, we assume that $(2,3) \notin E_X$.

The envy graph of the scenario is now as shown in Figure 1. Next, we discuss the possible configurations of the champion graph M_X . We show that most configurations are easily handled. If 3 champions 1, then by Observation 7, there is a Pareto dominating EFX allocation. If 3 does not champion 1, and since 1 does not self-champion, agent 2 champions 1. If now 1 champions 2, we have a 2-cycle in M_X involving 1 and 2, and 3 envies neither of them. Therefore by Remark 12, there is a Pareto dominating EFX allocation. Thus, we may assume that 1 does not champion 2. Since 2 does not self-champion, agent 3 champions 2. *There are only three possible configurations for M_X now, depending on who champions 3 (only 1, only 2, both 1 and 2 as 3 does not self-champion) (see Figure 2).*

We now show how to deal with these configurations of M_X . In Section 3, we showed how to move from the current allocation X to an allocation that Pareto dominates X . In Section 5, we show that this is impossible in this particular configuration of E_X and M_X . More specifically, we exhibit an EFX allocation X that is not Pareto dominated by any complete EFX allocation. We also show that there is no complete EFX allocation with higher Nash welfare than X , thereby falsifying a conjecture of Caragiannis et al. [9].

Recall that our potential is $\phi(X) = (v_a(X_a), v_b(X_b), v_c(X_c))$. We move to an allocation in which agent a is strictly better off. We distinguish the cases: $a = 1$, $a = 2$, and $a = 3$.

Also, recall that we are in the scenario where 2 champions 1 and 2 does not envy 1. Similarly 3 champions 2 and 3 does not envy 2. Therefore, by Observation 9, we have that $g \notin G_{21}$ and $g \notin G_{32}$,

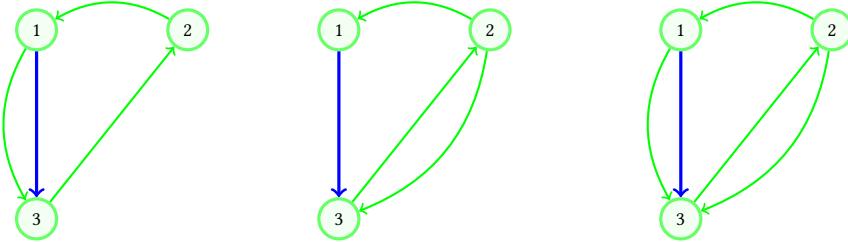


Fig. 2. The possible states of M_X that require further discussion: Green nodes correspond to the agents. Blue edges are the edges in E_X and green edges are the edges in M_X . There is a unique configuration of E_X and three different configurations of M_X , depending on who champions 3: only 1 (left), only 2 (middle), both 1 and 2 (right).

and hence, the bundles X_1 and X_2 decompose according to (1). Also, since 2 champions 1 and 1 does not self-champion, by Observation 10, we have that $G_{21} \neq \emptyset$, and a similar argument also shows that $G_{32} \neq \emptyset$.

4.1 Agent a is agent 1 or 3

We start from the allocation

$$X = \begin{array}{|c|} \hline X_1 \setminus G_{21} \\ \hline G_{21} \\ \hline \end{array} \quad (1) \quad \begin{array}{|c|} \hline X_2 \setminus G_{32} \\ \hline G_{32} \\ \hline \end{array} \quad (2) \quad \begin{array}{|c|} \hline X_3 \\ \hline \end{array} \quad (3)$$

Our goal is to determine an EFX allocation in which 1 and 3 are strictly better off (2 may be worse off). To this end, we consider

$$X' = \begin{array}{|c|} \hline X_3 \\ \hline \end{array} \quad (1) \quad \begin{array}{|c|} \hline X_1 \setminus G_{21} \\ \hline G_{32} \\ \hline \end{array} \quad (2) \quad \begin{array}{|c|} \hline X_2 \setminus G_{32} \\ \hline g \\ \hline \end{array} \quad (3)$$

In X' , every agent is better off than in X : 1 is better off because $X_3 >_1 X_1$ (1 envied 3 in E_X). We now show that 2 is better off: 2 championed 1 and 3 championed 2. Also, 2 did not self-champion, 2 did not envy 1 and 3 did not envy 2. Therefore, by Observation 11, (setting $i = k = 2, j = 1, i' = 3$), we have that $X_1 \setminus G_{21} >_2 X_2 \setminus G_{32}$. Hence, $(X_1 \setminus G_{21}) \cup G_{32} >_2 (X_2 \setminus G_{32}) \cup G_{32} = X_2$. Thus 2 is also better off. Agent 3 is better off as 3 championed 2, and by the definition of G_{32} , we have $(X_2 \setminus G_{32} \cup g) >_3 X_3$. Thus X' Pareto dominates X . If X' is EFX, we are done. So assume otherwise. We show that the only possible strong envy edge will be from 1 to 2.

- *Nobody envies 1*: Note that 1 has X_3 and neither 2 nor 3 envied X_3 earlier (3 had X_3 and 2 did not envy 3). Since both 2 and 3 are better off than before, they do not envy 1.
- *Nobody strongly envies 3*: 1 does not strongly envy 3 and 2 does not envy 3: 3 championed 2 and 1 did not. Therefore, by Observation 5 (part 1) we have $((X_2 \setminus G_{32}) \cup g) \setminus h \leq_1 X_1$ for all $h \in (X_2 \setminus G_{32}) \cup g$. Since 1 is better off than in X , it does not strongly envy 3. Agent 2 does not envy 3 since its prefers both of its parts over the corresponding part of agent 3. This was argued above for the top part and follows from Observation 10
- *3 does not envy 2*: 3 championed 2 and 3 did not envy 2 earlier. Therefore by Observation 9 we have that $G_{32} <_3 g$. Therefore $(X_1 \setminus G_{21}) \cup G_{32} <_3 (X_1 \setminus G_{21}) \cup g$. Since 2 championed 1

and 3 did not, by Observation 5 (part 2), we have $((X_1 \setminus G_{21}) \cup g) \leq_3 X_3$. Since 3 is better off than in X , 3 does not envy 2.

Thus, the only strong envy edge is from 1 to 2. The current state of the envy-graph is depicted below:



Let Z be a smallest cardinality subset of $(X_1 \setminus G_{21}) \cup G_{32}$ that 2 values more than $\max_2((X_2 \setminus G_{32}) \cup g, X_3)$, where $\max_2((X_2 \setminus G_{32}) \cup g, X_3)$ is defined as the more valuable bundle out of $(X_2 \setminus G_{32}) \cup g$ and X_3 according to 2. Note that $\max_2((X_2 \setminus G_{32}) \cup g, X_3) \leq_2 (X_1 \setminus G_{21}) \cup G_{32}$ since 2 does not envy neither 1 nor 3 in X' . Since the instance is non-degenerate, the inequality is strict, and hence Z exists. We now consider two allocations depending on 1's value for Z .

Case $Z \leq_1 X_3$: We replace 2's current bundle with Z and obtain

$$X'' = \begin{array}{ccc} \boxed{X_3} & \boxed{Z} & \boxed{\begin{array}{c} X_2 \setminus G_{32} \\ \hline g \end{array}} \\ (1) & (2) & (3) \end{array}$$

Agents 1 and 3 have the same bundles as in X' and hence are strictly better off than in X . Thus, X'' dominates X , as $a = 1$ or $a = 3$ and we improve a strictly. We next show that X'' is EFX. Since the only bundle we have changed is that of 2, and there were no strong envy edges between 1 and 3 earlier, it suffices to show that there are no strong envy edges to and from 2.

- *Nobody envies 2:* 3 did not envy the set $(X_1 \setminus G_{21}) \cup G_{32}$. As $Z \subseteq (X_1 \setminus G_{21}) \cup G_{32}$, agent 3 does not envy Z either. 1 does not envy Z because we are in the case where $Z \leq_1 X_3$.
- *2 does not envy anyone:* This follows from the definition of Z itself since $Z >_2 \max_2((X_2 \setminus G_{32}) \cup g, X_3)$.

Case $Z >_1 X_3$: In this case, we consider

$$X'' = \begin{array}{ccc} \boxed{Z} & \boxed{\max_2((X_2 \setminus G_{32}) \cup g, X_3)} & \boxed{\min_2((X_2 \setminus G_{32}) \cup g, X_3)} \\ (1) & (2) & (3) \end{array}$$

Agent 1 is still strictly better off than in X as we are in the case $Z >_1 X_3 >_1 X_1$, and agent 3 is not worse off than before as both X_3 and $(X_2 \setminus G_{32}) \cup g$ are at least as valuable to him as his previous bundle X_3 . We first show that X'' is EFX.

- *1 does not envy anyone:* We are in the case where $Z >_1 X_3$ and 1 did not envy $(X_2 \setminus G_{32}) \cup g$ when he had X_3 itself (and now 1 is better off than with X_3). Thus, 1 does not envy anyone.
- *2 does not strongly envy anyone:* Since 2 chooses the better bundle out of X_3 and $(X_2 \setminus G_{32}) \cup g$, 2 does not envy 3. Agent 2 does not strongly envy 1 since by the definition of Z , we have $Z \setminus h \leq_2 \max_2((X_2 \setminus G_{32}) \cup g, X_3)$ for all $h \in Z$. However, note that 2 envies 1. Thus, 2 does not envy 3 and does not strongly envy 1 (but envies 1).
- *3 does not strongly envy anyone:* 3 did not envy the set $(X_1 \setminus G_{21}) \cup G_{32}$,⁴ and $X_3 \leq X'_3$ as we argued above. Thus, 3 will not envy Z either as $Z \subseteq (X_1 \setminus G_{21}) \cup G_{32}$. We next show that 3 does not strongly envy 2, observe that $(X_2 \setminus G_{32}) \cup g >_3 X_3$. Therefore, if $\min_2((X_2 \setminus G_{32}) \cup g, X_3) = (X_2 \setminus G_{32}) \cup g$, we are done. So assume $\min_2((X_2 \setminus G_{32}) \cup g, X_3) = X_3$.

⁴We repeat the argument made earlier: 3 championed 2 and 3 did not envy 2 earlier. Therefore, by Observation 9 we have that $G_{32} <_3 g$. Hence, $(X_1 \setminus G_{21}) \cup G_{32} <_3 (X_1 \setminus G_{21}) \cup g$. Since 2 championed 1 and 3 did not, by Observation 5 (part 2), we have $((X_1 \setminus G_{21}) \cup g) \leq_3 X_3$.

Since 3 championed 2 and from Observation 5 (part 1), we have that $((X_2 \setminus G_{32}) \cup g) \setminus h \leq_3 X_3$ for all $h \in (X_2 \setminus G_{32}) \cup g$: Thus 3 does not strongly envy 2.

Now if $a = 1$, we are done, as X'' is EFX and agent 1 strictly improved. So assume $a = 3$. If $\min_2((X_2 \setminus G_{32}) \cup g, X_3) = (X_2 \setminus G_{32}) \cup g$, then agent 3 is strictly better off and we are done. This leaves the case that agent 3 gets X_3 , and hence

$$X'' = \begin{array}{ccc} \boxed{Z} & \boxed{X_2 \setminus G_{32}} & \boxed{X_3} \\ (1) & (2) & (3) \end{array}$$

The envy graph $E_{X''}$ with respect to allocation X'' is a path (shown below): 1 does not envy anyone, 2 envies 1 (not strongly) and does not envy 3, and 3 envies 2.



Also, note that we have some unallocated goods, e.g., the goods in G_{21} . Recall that we argued $G_{21} \neq \emptyset$ in the paragraph just before Section 4.1. Consider any good $g' \in G_{21}$. Since 3 is the only source in $E_{X''}$, by Corollary 8, there is an EFX allocation X''' Pareto dominating X'' , where $X'''_3 >_3 X''_3 = X_3$. Thus, we have an EFX allocation X''' that dominates X (as agent 3 is strictly better off and $a = 3$).

4.2 Agent a is agent 2

Recall that we argued just before the beginning of Section 4.1 that $g \notin G_{21}$ and $g \notin G_{32}$. Thus, the current EFX allocation X is

$$X = \begin{array}{ccc} \boxed{X_1 \setminus G_{21}} & \boxed{X_2 \setminus G_{32}} & \boxed{X_3} \\ (1) & (2) & (3) \end{array}$$

Our aim is to determine an EFX allocation, in which agent 2 has a bundle more valuable than X_2 . First, observe that $(X_1 \setminus G_{21}) \cup g$ is such a bundle. As 2 championed 1, we have $(X_1 \setminus G_{21}) \cup g >_2 X_2$ by the definition of G_{21} . We also observe that both agents 1 and 3 value X_3 as least as much as X_2 and $(X_1 \setminus G_{21}) \cup g$.

OBSERVATION 14. $X_3 >_i \max_i(X_2, ((X_1 \setminus G_{21}) \cup g))$ for $i \in \{1, 3\}$.

PROOF. We argue \geq_i ; strict inequality then follows from non-degeneracy.

Nobody envies 2 in X . Thus, $X_2 \leq_3 X_3$, and $X_2 \leq_1 X_1 <_1 X_3$ (the strict inequality holds as 1 envies 3 in X).

2 is the unique champion of 1 in X (both 1 and 3 do not champion 1). Therefore, by Observation 5 (part 2), we have $(X_1 \setminus G_{21}) \cup g \leq_3 X_3$ and $(X_1 \setminus G_{21}) \cup g \leq_1 X_1 <_1 X_3$ (the strict inequality holds as 1 envies 3 in X). \square

For $i \in \{1, 3\}$, let κ_i be the size of a smallest subset Z_i of X_3 such that $Z_i >_i \max_i((X_1 \setminus G_{21}) \cup g, X_2)$. We use the relative size of κ_1 and κ_3 to differentiate between agents 1 and 3. We use w (winner) to denote the agent with the smaller value of κ_i , i.e., $w = 1$ if $\kappa_1 \leq \kappa_3$ and $w = 3$ if $\kappa_1 > \kappa_3$. We use ℓ (loser) for the other agent. Consider

$$X' = \begin{array}{ccc} \boxed{X_3} & \boxed{\max_\ell(X_2, (X_1 \setminus G_{21}) \cup g)} & \boxed{\min_\ell(X_2, (X_1 \setminus G_{21}) \cup g)} \\ (w) & (\ell) & (2) \end{array}$$

In X' , the only possible strong envy edge is from ℓ to w . By Observation 14, w envies neither ℓ nor 2. Note that 2 championed 1 and therefore, $(X_1 \setminus G_{21}) \cup g >_2 X_2$, but by Observation 5 (part 1), we have $((X_1 \setminus G_{21}) \cup g) \setminus h \leq_2 X_2$ for all $h \in (X_1 \setminus G_{21}) \cup g$. Thus, 2 gets a bundle worth at least X_2 and does not strongly envy ℓ . 2 also does not envy w (as he did not envy X_3 when he had X_2). ℓ does not envy 2 as he chooses the better bundle out of X_2 and $X_1 \setminus G_{21} \cup g$. Thus, the only possible strong envy edge is from ℓ to w . How we proceed then depends on whether or not ℓ strongly envies w .

ℓ does not strongly envy w : Then X' is EFX. If $\min_{\ell}(X_2, (X_1 \setminus G_{21}) \cup g) = (X_1 \setminus G_{21}) \cup g$, we are done as X' dominates X (2 is strictly better off and $a = 2$). So assume otherwise. Then

$$X' = \begin{array}{ccc} \boxed{X_3} & \boxed{X_1 \setminus G_{21} \cup g} & \boxed{X_2} \\ (w) & (\ell) & (2) \end{array}$$

By Observation 14, ℓ envies w . Since 2 only envies ℓ , ℓ only envies w , and w does not envy anyone, the envy graph $E_{X'}$ is a path with source 2.



Also, note that there are unallocated goods, namely the goods in G_{21} (we argued just before the beginning of Section 4.1 that $G_{21} \neq \emptyset$). Therefore, by Corollary 8, there is an EFX allocation X'' , in which 2 is strictly better off. Thus, X'' dominates X as 2 is strictly better off and $a = 2$.

ℓ strongly envies w : We keep removing the least valuable good *according to w* from w 's bundle, until ℓ does not strongly envy w anymore. Let Z be the bundle obtained in this way. Consider

$$X' = \begin{array}{ccc} \boxed{Z} & \boxed{\max_{\ell}(X_2, (X_1 \setminus G_{21}) \cup g)} & \boxed{\min_{\ell}(X_2, (X_1 \setminus G_{21}) \cup g)} \\ (w) & (\ell) & (2) \end{array}$$

CLAIM 15. w does not envy 2 and ℓ .

PROOF. Recall that κ_w is the smallest cardinality of a subset of X_3 that w still values more than $\max_w(X_2, (X_1 \setminus G_{21}) \cup g)$; κ_w was defined just after Observation 14. Such a set can be obtained by removing w 's $|X_3| - \kappa_w$ least valuable goods from X_3 . Observe that Z is obtained by removing $|X_3| - |Z|$ of w 's least valuable goods from X_3 . If $|Z| \geq \kappa_w$, w will envy neither 2 nor ℓ . If $|Z| < \kappa_w \leq \kappa_{\ell}$ (recall that $\kappa_w \leq \kappa_{\ell}$), let h be the last good removed. Then ℓ strongly envies $Z \cup h$ (otherwise we would not have removed h), meaning that there exists an $h' \in Z \cup h$ such that $(Z \cup h) \setminus h' >_{\ell} \max_{\ell}(X_2, (X_1 \setminus G_{21}) \cup g)$. Thus, there is a subset of X_3 of size $|(Z \cup h) \setminus h'| < \kappa_w + 1 - 1 = \kappa_w$ that ℓ values more than $\max_{\ell}(X_2, (X_1 \setminus G_{21}) \cup g)$, a contradiction to $\kappa_w \leq \kappa_{\ell}$. \square

The allocation X' is EFX: w envies neither 2 nor ℓ , ℓ does not strongly envy w , ℓ does not envy 2, and 2 envies neither ℓ nor w . If $\min_{\ell}(X_2, (X_1 \setminus G_{21}) \cup g)$ is $X_1 \setminus G_{21} \cup g$, then we are done as X' dominates X (2 is strictly better off and $a = 2$). So assume otherwise. Then

$$X' = \begin{array}{ccc} \boxed{Z} & \boxed{X_1 \setminus G_{21} \cup g} & \boxed{X_2} \\ (w) & (\ell) & (2) \end{array}$$

In X' , w envies nobody (by Claim 15), 2 envies ℓ , and ℓ may or may not envy w . We distinguish cases according to whether or not ℓ envies w .



Case ℓ envies w : Then, the current envy graph is a path with 2 as the source.



Since there are unallocated goods, namely the goods in G_{21} (we argued just before the beginning of Section 4.1 that $G_{21} \neq \emptyset$), by Corollary 8, there is an EFX allocation X'' in which agent 2 is strictly better off. The allocation X'' dominates X (as 2 is strictly better off and $a = 2$).

Case ℓ does not envy w : Then the current envy graph has two sources, namely w and 2, and one envy edge from 2 to ℓ .



There are at least two unallocated goods, the goods in G_{21} (we argued just before the beginning of Section 4.1 that $G_{21} \neq \emptyset$) and the goods in $X_3 \setminus Z$ (note that this set is not empty; we definitely have removed at least one good from X_3 as ℓ strongly envied it in X'). Now consider the allocation X' and some $g' \in G_{21}$. If the champion of 2 is 2 itself or ℓ (definition of champion based on allocation X' and the unallocated good g'), by Observation 7 there is an EFX allocation Y where the source, namely 2, is strictly better off and hence Y will dominate X . So assume that the champion of 2 is w , i.e., $w \in A_{X'}(X'_2 \cup g')$. Let $g'' \in X_3 \setminus Z$ be the last element that we removed from X_3 when we constructed Z from X_3 . Then ℓ strongly envies $Z \cup g''$ and, according to w , g'' is the least valuable good in $Z \cup g''$. We observe that ℓ is the unique champion of w (definition of champion based on allocation X' and the unallocated good g''), i.e., $A_{X'}(X'_w \cup g'') = \{\ell\}$.

OBSERVATION 16. Let $g'' \in X_3 \setminus Z$ be the last element that we removed from X_3 when we constructed Z from X_3 . We have $A_{X'}(X'_w \cup g'') = \{\ell\}$.

PROOF. We have $X'_w = Z$. First we show that $2 \notin A_{X'}(Z \cup g'')$. Note that $Z \cup g'' \subseteq X_3$. Since $X_2 \geq_2 X_3$ (as 2 did not envy 3 in X), 2 will not envy $Z \cup g''$ either.

By the construction of Z , g'' is w 's least valuable good in $Z \cup g''$. Thus, the removal of any good from $Z \cup g''$ will result in a bundle whose value for w is no more than the value of Z for w . Therefore, $\kappa_{X'}(w, Z \cup g'') = |Z \cup g''|^5$. Note that ℓ strongly envies $Z \cup g''$. Hence, there exists $h \in Z \cup g''$ such that $(Z \cup g'') \setminus h >_\ell X'_\ell$. Therefore, $\kappa_{X'}(\ell, Z \cup g'') \leq |(Z \cup g'') \setminus h| = |Z \cup g''| - 1 < \kappa_X(w, Z \cup g'')$. Thus, w does not self-champion and hence $A_{X'}(Z \cup g'') = \{\ell\}$. \square

Consider

$$X'' = \begin{array}{c} (X'_2 \cup g') \setminus G_{w2} \\ (w) \end{array} \quad \begin{array}{c} (X'_w \cup g'') \setminus G_{\ell w} \\ (\ell) \end{array} \quad \begin{array}{c} X'_\ell \\ (2) \end{array}$$

or equivalently

$$X'' = \begin{array}{c} (X_2 \cup g') \setminus G_{w2} \\ (w) \end{array} \quad \begin{array}{c} (Z \cup g'') \setminus G_{\ell w} \\ (\ell) \end{array} \quad \begin{array}{c} (X_1 \setminus G_{21}) \cup g \\ (2) \end{array}$$

Note that every agent is strictly better off than in X' . w championed 2, and by the definition of G_{w2} , we have $(X'_2 \cup g') \setminus G_{w2} >_w X'_w$. Similarly, ℓ championed w , and by the definition of

⁵Recall that $\kappa_X(i, S)$ is the size of the smallest subset of S which is more valuable to i than X_i .

$G_{\ell w}$, we have $(X'_w \cup g'') \setminus G_{\ell w} >_{\ell} X'_\ell$. 2 is better off as 2 envied ℓ in X' i.e. $X'_2 <_2 X'_\ell$. Now we have an allocation X'' in which agent 2 is strictly better off than it was in X . Thus, X'' dominates X (as $a = 2$). It suffices to show that X'' is EFX now. To this end, observe that,

- *Nobody strongly envies w*: w championed 2. Thus, by Observation 5 (part 1), we have that $((X'_2 \cup g') \setminus G_{w2}) \setminus h \leq_2 X'_2$ and $((X'_2 \cup g') \setminus G_{w2}) \setminus h \leq_\ell X'_\ell$ for all $h \in ((X'_2 \cup g') \setminus G_{w2})$. Since both 2 and ℓ are better off than before (in X'), they do not strongly envy w .
- *Nobody strongly envies ℓ* : The argument is very similar to the previous case. ℓ championed 2. Thus, by Observation 5 (part 1), we have that $((X'_w \cup g'') \setminus G_{\ell w}) \setminus h \leq_2 X'_2$ and $((X'_w \cup g'') \setminus G_{\ell w}) \setminus h \leq_w X'_w$ for all $h \in ((X'_w \cup g'') \setminus G_{\ell w})$. Since both 2 and w are better off than before (than they were in X'), they do not strongly envy w .
- *Nobody strongly envies 2*: Both w and ℓ did not envy X'_ℓ (ℓ had X'_ℓ and w did not envy ℓ) when they had X'_w and X'_ℓ itself. Both w and ℓ are strictly better off than they were in X' . Therefore, they also do not envy 2.

We conclude that there is an EFX allocation dominating X in the case, $a = 2$ as well.

This allows us to summarize our main result for this section as follows,

LEMMA 17. *Let X be a partial EFX allocation, and let g be an unallocated good, where the envy graph E_X has two sources. Then there is an EFX allocation Y dominating X .*

Having covered all the cases, we arrive at our main result:

THEOREM 18. *For any instance $I = \langle [3], M, \mathcal{V} \rangle$ where all $v_i \in \mathcal{V}$ are additive, an EFX allocation always exists.*

PROOF. We start off with an empty allocation ($X_i = \emptyset$ for all $i \in [3]$), which is trivially EFX. As long as X is not a complete EFX allocation, there is an allocation Y that dominates X : If E_X has a single source or M_X has a 1-cycle, there is a dominating EFX allocation Y by Corollary 8. Lemmas 13 and 17 establish the existence of Y when E_X has multiple sources and M_X does not have a 1-cycle. Since ϕ is bounded from above, the process must stop. When it stops, we have arrived at a complete EFX allocation. \square

5 BARRIERS IN CURRENT TECHNIQUES

In this section, we highlight some barriers to the current techniques for computing EFX allocations. We give an instance with three agents and seven goods such that there is a partial EFX allocation for six of the goods that is not Pareto dominated by any complete EFX allocation for the full set of goods. We also generalize this example and give an instance with a partial EFX allocation which has a Nash welfare larger than the Nash welfare of any complete EFX allocation. These examples make it unlikely that there is an iterative algorithm towards a complete EFX allocation that improves the current EFX allocation in each iteration either in the sense of Pareto domination or in the sense of Nash welfare (like the algorithms in [27] and [13]). The second example (see Section 5 in the full version of the paper [12]) also falsifies the EFX monotonicity conjecture (see Conjecture 1 in Section ??) by Caragiannis et al. [9].

THEOREM 19. *For the instance given in Table 1, the partial allocation $X = \langle X_1, X_2, X_3 \rangle$, where*

$$X_1 = \{g_2, g_3, g_4\} \quad X_2 = \{g_1, g_5\} \quad X_3 = \{g_6\},$$

is an EFX allocation of the first six goods. No complete EFX allocation Pareto dominates X .

PROOF. Note that $v_1(X_1) = 16$, $v_2(X_2) = 15$, and $v_3(X_3) = 10$. We will show that there is no complete EFX allocation X' with $v_1(X'_1) \geq 16$, $v_2(X'_2) \geq 15$ and $v_3(X'_3) \geq 10$. To this end, we systematically consider potential bundles X'_1 that can keep a_1 's valuation at or above 16.

	g_1	g_2	g_3	g_4	g_5	g_6	g_7
a_1	8	2	12	2	0	17	1
a_2	5	0	9	4	10	0	3
a_3	0	0	0	0	9	10	2

Table 1. An instance where no complete EFX allocation dominates the EFX allocation X for the first six goods defined in the text. The valuations are assumed to be additive and the entry in row i and column j is the value of good j for agent i .

Let us first assume $g_6 \in X'_1$, and hence, $v_1(X'_1) \geq 17$. Now, to ensure $v_3(X'_3) \geq 10$, we need to allocate g_5 and g_7 to a_3 . We are left with goods g_1, g_2, g_3 and g_4 . In order to ensure $v_2(X'_2) \geq 15$, we definitely need to allocate g_1, g_3 and g_4 to a_2 . Now even if we allocate the remaining good g_2 to a_1 , we will have $v_1(X'_1) = v_1(\{g_2, g_6\}) = 19 < 20 = v_1(\{g_1, g_3\}) \leq v_1(X'_2 \setminus g_4)$. Therefore, a_1 will strongly envy a_2 . Thus $g_6 \notin X'_1$.

If $g_6 \notin X'_1$ and $v_1(X'_1) \geq 16$, X'_1 must contain g_3 (the total valuation for a_1 of all the goods other than g_3 and g_7 is less than 16). We need to consider several subcases.

Assume $g_1 \in X'_1$ first. Since X'_1 already contains g_1 and g_3 , the goods that can be allocated to a_2 and a_3 are g_2, g_4, g_5, g_6 , and g_7 . In order to ensure $v_2(X'_2) \geq 15$ we need to allocate g_4, g_5 , and g_7 to a_2 . Even if we allocate all the remaining goods (g_2 and g_6) to a_3 , we have $v_3(X'_3) = v_3(\{g_3, g_6\}) = 10 < 11 = v_3(\{g_5, g_7\}) \leq v_3(X'_2 \setminus g_4)$. Therefore, a_3 will strongly envy a_2 .

Thus $g_1 \notin X'_1$. Since neither g_1 nor g_6 belongs to X'_1 , the only way to ensure $v_1(X'_1) \geq 16$ is to at least allocate g_2, g_3 , and g_4 to a_1 (we can allocate more). Similarly, given that the goods not allocated yet are g_1, g_5, g_6 , and g_7 , the only way to ensure $v_1(X'_1) \geq 15$ is to allocate at least g_1 and g_5 to a_2 . Similarly, the only way to ensure $v_3(X'_3) \geq 10$ now is to allocate at least g_6 to a_3 . We next show that adding g_7 to any one of the existing bundles will cause a violation of the EFX property.

- Adding g_7 to X'_1 : a_2 strongly envies a_1 as $v_2(X'_2) = 15 < 16 = v_2(\{g_3, g_4, g_7\}) = v_2(X'_1 \setminus g_2)$.
- Adding g_7 to X'_2 : a_3 strongly envies a_2 as $v_3(X'_3) = 10 < 11 = v_3(\{g_5, g_7\}) = v_3(X'_2 \setminus g_1)$.
- Adding g_7 to X'_3 : a_1 strongly envies a_3 as $v_1(X'_1) = 16 < 17 = v_1(g_6) = v_1(X'_3 \setminus g_7)$.

Thus, there exists no complete EFX allocations Pareto dominating X . \square

6 CONCLUSION

In this paper, we have shown that EFX allocations always exist when we have three agents with additive valuations. Our proof is constructive and leads to a pseudo-polynomial algorithm. We have identified some crucial barriers in the current techniques and have overcome them with novel techniques. We feel that this is a step towards resolving the bigger question whether EFX allocations always exist when we have an arbitrary number of agents.

Our proofs crucially use additivity and do not work for more general valuation functions like submodular or subadditive. Therefore, an ideal next step would be to investigate EFX allocations with three agents, but more general valuations.

We also showed some barriers to finding *efficient* EFX allocations (EFX allocations with high Nash welfare). While efficient approximate EFX allocations or efficient EFX allocations with bounded charity exist, it is unclear how much efficiency we can guarantee for complete EFX allocations—i.e., what trade-off with efficiency is required to guarantee fairness.

ACKNOWLEDGMENTS

We would like to thank Hannaneh Akrami, Corinna Coupette, Kavitha Telikepalli and Alkmini Sgouritsa for helpful discussions. We thank Corinna Coupette also for a careful reading of the manuscript. This work is partially supported by NSF Grants CCF-1755619 (CRII) and CCF-1942321 (CAREER).

REFERENCES

- [1] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. 2017. Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms* 13, 4 (2017), 52:1–52:28.
- [2] Nima Anari, Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. 2017. Nash Social Welfare, Matrix Permanent, and Stable Polynomials. In *8th Innovations in Theoretical Computer Science Conference (ITCS)*. 1–12.
- [3] Nima Anari, Tung Mai, Shayan Oveis Gharan, and Vijay V. Vazirani. 2018. Nash Social Welfare for Indivisible Items under Separable, Piecewise-Linear Concave Utilities. In *Proceedings of the 29th Symposium on Discrete Algorithms (SODA)*. 2274–2290.
- [4] Siddharth Barman and Sanath Kumar Krishnamurthy. 2017. Approximation algorithms for maximin fair division. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*. 647–664.
- [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.
- [6] Sylvain Bouveret and Michel Lemaître. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *Autonomous Agents and Multi-Agent Systems (AAMAS) 30*, 2. 259–290.
- [7] Eric Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- [8] Eric Budish, Gérard P. Cachon, Judd B. Kessler, and Abraham Othman. 2017. Course Match: A Large-Scale Implementation of Approximate Competitive Equilibrium from Equal Incomes for Combinatorial Allocation. *Operations Research* 65, 2 (2017), 314–336.
- [9] Ioannis Caragiannis, Nick Gravin, and Xin Huang. 2019. Envy-Freeness Up to Any Item with High Nash Welfare: The Virtue of Donating Items. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*. 527–545.
- [10] 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.
- [11] Bhaskar Ray Chaudhury, Yun Kuen Cheung, Jugal Garg, Naveen Garg, Martin Hoefer, and Kurt Mehlhorn. 2018. On Fair Division for Indivisible Items. In *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS*. 25:1–25:17.
- [12] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. 2020. EFX Exists for Three Agents. *CoRR* abs/2002.05119 (2020).
- [13] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. 2020. A Little Charity Guarantees Almost Envy-Freeness. In *Proceedings of the 31st Symposium on Discrete Algorithms (SODA)*. 2658–2672.
- [14] Richard Cole, Nikhil Devanur, Vasilis Gkatzelis, Kamal Jain, Tung Mai, Vijay Vazirani, and Sadra Yazdanbod. 2017. Convex Program Duality, Fisher Markets, and Nash Social Welfare. In *Proc. 18th Conf. Economics and Computation (EC)*.
- [15] Richard Cole and Vasilis Gkatzelis. 2018. Approximating the Nash Social Welfare with Indivisible Items. *SIAM J. Comput.* 47, 3 (2018), 1211–1236.
- [16] Jugal Garg, Martin Hoefer, and Kurt Mehlhorn. 2018. Approximating the Nash Social Welfare with Budget-Additive Valuations. In *Proceedings of the 29th Symposium on Discrete Algorithms (SODA)*. 2326–2340.
- [17] Jugal Garg, Pooja Kulkarni, and Rucha Kulkarni. 2020. Approximating Nash Social Welfare under Submodular Valuations through (Un)Matchings. In *Proceedings of the 31st Symposium on Discrete Algorithms (SODA)*. 2673–2687.
- [18] Jugal Garg and Peter McGlaughlin. 2019. Improving Nash Social Welfare Approximations. In *IJCAI. ijcai.org*, 294–300.
- [19] Jugal Garg, Peter McGlaughlin, and Setareh Taki. 2019. Approximating Maximin Share Allocations. In *Proceedings of the 2nd Symposium on Simplicity in Algorithms (SOSA)*, Vol. 69. 20:1–20:11.
- [20] Jugal Garg and Setareh Taki. 2019. An Improved Approximation Algorithm for Maximin Shares. *CoRR* abs/1903.00029 (2019).
- [21] Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2018. Fair Allocation of Indivisible Goods: Improvements and Generalizations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. 539–556.

- [22] Jonathan R. Goldman and Ariel D. Procaccia. 2014. Spliddit: unleashing fair division algorithms. In *SIGecom Exchanges* 13(2). 41–46.
- [23] David Kurokawa, Ariel D. Procaccia, and Nisarg Shah. 2018. Leximin Allocations in the Real World. *ACM Trans. Economics and Comput.* 6, 3-4 (2018), 11:1–11:24.
- [24] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2018. Fair enough: Guaranteeing approximate maximin shares. *Journal of ACM* 65, 2 (2018), 8:1–27.
- [25] Euiwoong Lee. 2017. APX-hardness of Maximizing Nash Social Welfare with Indivisible Items. *Inf. Process. Lett.* 122 (2017), 17–20.
- [26] 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.
- [27] Benjamin Plaut and Tim Roughgarden. 2018. Almost Envy-Freeness with General Valuations. In *Proceedings of the 29th Symposium on Discrete Algorithms (SODA)*. 2584–2603.
- [28] Hugo Steinhaus. 1948. The Problem of Fair Division. *Econometrica* 16, 1 (1948), 101–104.