Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Improving Envy Freeness up to Any Good Guarantees Through Rainbow Cycle Number

Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta, Pranabendu Misra

informs
https://pubsonline.informs.org/journal/moor

MATHEMATICS OF OPERATIONS RESEARCH
Vol. 49, No. 4, November 2024, pp. 2323–2340
ISSN 0364-765X (print), ISSN 1526-5471 (online)

Improving Envy Freeness up to Any Good Guarantees Through Rainbow Cycle Number

Bhaskar Ray Chaudhury,[a] Jugal Garg,[a,*] Kurt Mehlhorn,[b] Ruta Mehta,[c] Pranabendu Misra[d]

[a] Department of Industrial and Enterprise Systems Engineering, University of Illinois Urbana–Champaign, Urbana, Illinois 61820; [b] Saarland Informatics Campus, Max Planck Institute for Informatics, 66123 Saarbrucken, Germany; [c] Department of Computer Science, University of Illinois Urbana-Champaign, Urbana, Illinois 61801; [d] Department of Computer Science, Chennai Mathematical Institute, Chennai 603103, India

*Corresponding author

Contact: braycha@illinois.edu, https://orcid.org/0000-0002-7808-1483 (BRC); jugal@illinois.edu, https://orcid.org/0000-0001-6439-7308 (JG); mehlhorn@mpi-inf.mpg.de, https://orcid.org/0000-0003-4020-4334 (KM); rutameht@illinois.edu, https://orcid.org/0000-0002-1549-2752 (RM); pranabendu@cmi.ac.in, https://orcid.org/0000-0002-7086-5590 (PM)

Received: September 22, 2021
Revised: November 29, 2022
Accepted: September 19, 2023
Published Online in Articles in Advance: November 22, 2023

MSC2020 Subject Classifications: Primary: 91B32, 68Q25, 68W40

https://doi.org/10.1287/moor.2021.0252

Abstract. We study the problem of fairly allocating a set of indivisible goods among n agents with additive valuations. Envy freeness up to *any* good (EFX) is arguably the most compelling fairness notion in this context. However, the existence of an EFX allocation has not been settled and is one of the most important problems in fair division. Toward resolving this question, many impressive results show the existence of its relaxations. In particular, it is known that 0.618-EFX allocations exist and that EFX allocation exists if we do not allocate at most (n-1) goods. Reducing the number of unallocated goods has emerged as a systematic way to tackle the main question. For example, follow-up works on three- and four-agents cases, respectively, allocated two more unallocated goods through an involved procedure. In this paper, we study the general case and achieve *sublinear* numbers of unallocated goods. Through a new approach, we show that for every $\varepsilon \in (0, 1/2]$, there always exists a $(1-\varepsilon)$-EFX allocation with *sublinear* number of unallocated goods and high Nash welfare. For this, we reduce the EFX problem to a novel problem in extremal graph theory. We define the notion of *rainbow cycle number* $R(\cdot)$ in directed graphs. For all $d \in \mathbb{N}$, $R(d)$ is the largest k such that there exists a k-partite graph $G = (\cup_{i\in[k]} V_i, E)$, in which each part has at most d vertices (i.e., $|V_i| \leq d$ for all $i \in [k]$); for any two parts V_i and V_j, each vertex in V_i has an incoming edge from some vertex in V_j and vice versa; and there exists no cycle in G that contains at most one vertex from each part. We show that any upper bound on $R(d)$ directly translates to a sublinear bound on the number of unallocated goods. We establish a polynomial upper bound on $R(d)$, yielding our main result. Furthermore, our approach is constructive, which also gives a polynomial-time algorithm for finding such an allocation.

Funding: J. Garg was supported by the Directorate for Computer and Information Science and Engineering [Grant CCF-1942321]. R. Mehta was supported by the Directorate for Computer and Information Science and Engineering [Grant CCF-1750436].

Keywords: discrete fair division • EFX allocations • rainbow cycle number

1. Introduction

Fair division of resources is a fundamental problem in many disciplines, including computer science, economics, and social choice theory. The objective is to distribute resources among agents in a *fair* (no agent is significantly unhappy with her allocation) and *efficient* (there is no other *fair allocation* that can achieve better total welfare) manner. Mentions of such problems date back to the Bible and ancient Greek mythology. Today, the issue of fair division arises in division of labor, inheritance, computing resources, divorce settlements, partnership dissolutions, splitting rent among tenants, splitting taxi fare among passengers, dividing household tasks, air traffic management, frequency allocation, and so on. In the internet age, the existence of several centralized platforms and more computational power has triggered substantial interest from the economics and computer science community to find computationally tractable protocols to allocate resources fairly; see Spliddit (www.spliddit.org) and Fair Outcomes (www.fairoutcomes.com) for more details on fair division protocols used in real-life scenarios.

1.1. Discrete Fair Division

In this paper, we focus on one of the most important open problems in discrete fair division. To this end, we first describe a typical setup of a problem instance. Given a set N of n agents and a set M of m *indivisible* goods, the goal is to determine a partition $X = \langle X_1, X_2, \dots, X_n \rangle$ of the good set M such that agent $i \in N$ receives the bundle X_i and the allocation is *fair*.

1.2. Envy Freeness up to Any Good Allocations

A quintessential notion of fairness is that of envy freeness; an allocation X is said to be envy free if and only if for every pair of agents i and j, we have $v_i(X_i) \geq v_i(X_j)$ (i.e., each agent i values her own bundle at least as much as she values the bundles of other agents). However, such allocations may not always exist; consider a simple example with two agents having a positive valuation toward a single good. The agent who gets this good is envied by the one who does not. Therefore, several relaxations of envy freeness have been proposed and studied over the last 15 years (Budish [16], Caragiannis et al. [18], Lipton et al. [34]). The most compelling relaxation is *envy freeness up to any good* (EFX) proposed by Caragiannis et al. [18], where no agent envies the other agent following the removal of *any* single good from the other's bundle; that is, an allocation X is said to be EFX if and only if for every pair of agents i and j, we have $v_i(X_i) \geq v_i(X_j \setminus \{g\})$ for all $g \in X_j$. It is also regarded as the best analogue of envy freeness in discrete fair division. Caragiannis et al. [17, p. 528] remarked that "Arguably, EFX is the best fairness analog of envy-freeness for indivisible items."

Unfortunately, it is not known whether EFX allocations always exist, even when there are only four agents with additive valuations despite significant efforts (e.g., see Caragiannis et al. [18], Moulin [37]). Indeed, only recently was this question resolved affirmatively for three agents with additive valuations by Chaudhury et al. [19]. In fact, Procaccia [40, p. 118] remarked: "This fundamental and deceptively accessible question is open. In my view, it is the successor of envy-free cake cutting as fair division's biggest problem."

There has been a substantial study on the existence of an EFX allocation in special cases and its relaxations. For instance, EFX allocations exist when agents' valuations are identical (Plaut and Roughgarden [39]), binary (Barman et al. [11], Darmann and Schauer [25]), and bivalued (Amanatidis et al. [5], Garg and Murhekar [28]). The two primary relaxations of EFX are approximate EFX allocations and partial EFX allocations.

- Approximate EFX allocation. An allocation $X = \langle X_1, X_2, \dots, X_n \rangle$ is an α-EFX allocation for some scalar $\alpha \in (0, 1]$ if for every pair of agents i and j, we have $v_i(X_i) \geq \alpha \cdot v_i(X_j \setminus \{g\})$ for all $g \in X_j$. Plaut and Roughgarden [39] showed the existence of 0.5-EFX allocations. Amanatadis et al. [3] show that a clever modification of the same approach leads to a 0.618-EFX allocation.
- Partial EFX allocation. An allocation $X = \langle X_1, X_2, \dots, X_n \rangle$ is called a partial EFX allocation if X is EFX and not all goods are necessarily allocated (i.e., $\cup_{i \in [n]} X_i \subseteq M$). There is always a trivial partial-EFX allocation where each X_i is empty. Therefore, a good partial EFX allocation is the one that has good *qualitative* and *quantitative* guarantees on the unallocated goods. Caragiannis et al. [17] showed that there exists a partial EFX allocation where every agent gets a bundle that she values at least as much as half of her value for the bundle she receives in a *Nash welfare*-maximizing allocation. Here, the Nash welfare of an allocation $NW(X) = (\prod_{i \in [n]} v_i(X_i))^{1/n}$ is another popular measure of fairness and economic efficiency. Following the same line of work, Chaudhury et al. [20] showed that there always exists a partial EFX allocation X and a set of unallocated goods P such that
- nobody envies the set of unallocated items: $v_i(X_i) \geq v_i(P)$ for all $i \in N$; and
- at most $n - 1$ goods are unallocated: $|P| \leq n - 1$.

There have been recent interesting studies on the relaxations of EFX allocations. Berger et al. [13] improved the number of unallocated goods to $(n - 2)$ when there are n agents with additive valuations and to one in case of four agents. Very recently, Mahara [35] showed how to reduce the number of unallocated goods to $(n - 2)$ even when agents have general monotone valuations. We remark that studying relaxations (of EFX allocations) is a systematic and promising direction to investigate the existence of an EFX allocation. It has been suspected in Plaut and Roughgarden [39, p. 1062] that EFX allocations may not exist in the general setting: "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."

However, finding counterexamples, at least in the additive setting, seems to be a very challenging task; quite recently, Manurangsi and Suksompong [36] showed that when agents' valuations for individual items are drawn at random from a probability distribution, then EFX allocations exist with high probability. This demands a nonbrute-force approach to find counterexamples, if any. Thus, finding better relaxations (improving the approximation factor or reducing the number of unallocated goods in a partial EFX allocation) is a crucial step toward finding the right answer to this big open question. We achieve exactly this by our first main result.

Theorem 1. *For all $\varepsilon \in (0, 1/2]$, we can determine a partial allocation X and a set of unallocated goods P in polynomial time such that*

- *X is $(1-\varepsilon)$-EFX and*
- *$|P| \in \mathcal{O}((n/\varepsilon)^{4/5})$.*

We remark that reducing the number of unallocated goods could be quite challenging. Indeed, a corollary on the bounded-charity result in Chaudhury et al. [20] already establishes that there exists a partial EFX allocation with at most two goods unallocated when there are three agents. However, removing the last two goods to obtain an EFX allocation for three agents turns out to be a highly nontrivial task, and the proof by Chaudhury et al. [19] requires careful and cumbersome case analysis. Furthermore, in the appendix, we show that the technique in Chaudhury et al. [19] does not extend to four agents with additive valuations for finding a $(1-\varepsilon)$-EFX allocation.

In this paper, we develop a novel method that reduces the problem of determining good relaxations of EFX allocations to a combinatorial problem in graph theory. We call it the *rainbow cycle number* of an integer, defined as follows.

Definition 1. For any positive integer d, the rainbow cycle number or $R(d)$ is the largest k such that there exists a directed k-partite graph $G = (\cup_{i\in[k]} V_i, E)$ such that

1. $|V_i| \leq d$ for all $i \in [k]$;
2. for any two distinct parts V_i and V_j in G, every vertex in V_i has an incoming edge from a vertex in V_j; and
3. there exists no cycle in G that intersects each part at most once.

Let us deduce that $R(1) = 1$. It is clear that G can be a single vertex and satisfy all the conditions in Definition 1; thus, $R(1) \geq 1$. However, $R(1)$ cannot be larger than one as otherwise, we have two parts V_1 and V_2 in a graph G, where there is exactly one vertex each in V_1 and V_2. So, let $V_1 = \{a_1\}$ and $V_2 = \{a_2\}$. By condition (2) in Definition 1, we must have an edge from a_1 to a_2 and an edge from a_2 to a_1. This gives a two-cycle $a_1 \to a_2 \to a_1$. However, this cycle contains exactly one vertex from each V_1 and V_2, which contradicts condition (3) in Definition 1.

Similarly, using a more involved argument, we can also determine that $R(2) = 2$. However, it is not at all clear what values $R(d)$ takes or if it is finite for all integers d. A key technical result of this paper is a polynomial (in d) upper bound on $R(d)$.

Theorem 2. *For all $d \geq 1$, we have $R(d) \leq d^4 + d$. Furthermore, let G be a k-partite digraph with $k > d^4 + d$ parts of cardinality at most d each such that for every vertex v and any part W not containing v, there is an edge from W to v. Then, there exists a cycle in G visiting each part at most once, and it can be found in time polynomial in k.*

Observe that the definition of the rainbow cycle number ($R(\cdot)$) is independent of the agents, goods, and valuation functions. In the second key result of this paper, we establish a direct relation between the rainbow cycle number and the existence of better EFX relaxations. Finding a good upper bound on the rainbow cycle number can get us weaker relaxations of EFX allocations (we can asymptotically improve the number of unallocated goods). Formally, we have Theorem 3.

Theorem 3. *Let $\zeta(n/\varepsilon)$ be the largest integer d such that $d \cdot R(d) \leq n/\varepsilon$ for $\varepsilon \in (0, 1/2]$. Then, there is a $(1-\varepsilon)$-EFX allocation X and a set of unallocated goods P such that $|P| \leq 4n/(\varepsilon \cdot \zeta(2n/\varepsilon))$.*

Theorems 2 and 3 imply Theorem 1. We remark that, although we give a polynomial upper bound on $R(d)$, we believe that there is further room for improvement. As an illustration, we briefly show that $R(2) \leq 2$, which is significantly better than our upper bound for $d = 2$ obtained from Theorem 2. We prove this by contradiction. Let us assume otherwise, and let V_1, V_2, and V_3 be any three parts of G. We first look into the edges of the induced bipartite graph $G[V_1 \cup V_2]$. Without loss of generality, let us assume that vertex b_1 in V_2 has an incoming edge from vertex a_1 in V_1. By condition (2) in Definition 1, a_1 has an incoming edge from some vertex in V_2. However, this vertex cannot be b_1 as this will violate condition (3) in Definition 1. This implies that there must be another vertex in V_2, say b_2 that has an edge to a_1. Again, by a similar argument, b_2 cannot have an incoming edge from a_1 and therefore, has an incoming edge from another vertex in V_1, say that a_2 and a_2 have the incoming edge from b_1 and not b_2 (because there can be no other vertices in V_2). Thus, the induced bipartite graph $G[V_1 \cup V_2]$ is a four cycle as shown in Figure 1.

Note that the induced bipartite graph $G[V_2 \cup V_3]$ will be isomorphic to $G[V_1 \cup V_2]$. Thus, so far we have the following edges in $G[V_1 \cup V_2 \cup V_3]$ (Figure 2).

We now look at the edges between the parts V_1 and V_3. Because $G[V_1 \cup V_3]$ is isomorphic to $G[V_1 \cup V_2]$, it must also be a four cycle, and hence, in $G[V_1 \cup V_3]$, there is an edge either from a_1 to c_1 or from c_1 to a_1. If there is an edge from a_1 to c_1, then we have a three-cycle $a_1 \to c_1 \to b_2 \to a_1$, which visits each part of G at most once, and thus, this is a contradiction. Similarly, if there is an edge from c_1 to a_1, then also we have a three-cycle $a_1 \to b_1 \to c_1 \to a_1$, which visits each part of G at most once, and thus, this is also a contradiction.

Figure 1. (Color online) Induced bipartite graph $G[V_1 \cup V_2]$ when showing $R(2) \le 2$.

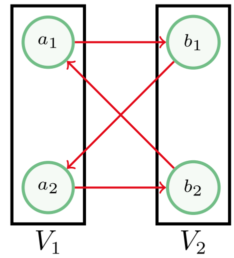

We suspect that $R(d) \in \mathcal{O}(d)$. We believe that finding better upper bounds on $R(d)$ is a natural combinatorial question, and better upper bounds to $R(d)$ imply the existence of better relaxations of EFX allocations. Therefore, investigating better upper bounds on the rainbow cycle number is of interest in its own right, and we leave this as an interesting open problem.

1.3. Finding $(1-\varepsilon)$-EFX Allocations with High Nash Welfare

Let us recall that *efficiency* is also an important and desirable property of the allocations in fair division. The efficiency of an allocation is a measure of the overall welfare the allocation achieves. This is important as an envy-free allocation could be otherwise unsatisfactory; consider a simple instance with two agents 1 and 2 and two goods g_1 and g_2. Let $v_1(g_1) = v_2(g_2) = 1$ and $v_1(g_2) = v_2(g_1) = 0$. Note that $X_1 \leftarrow \{g_2\}$ and $X_2 \leftarrow \{g_1\}$ are an EFX allocation as each bundle is a singleton, and following the removal of a single good results in an empty bundle, which is unenvied. However, there is clearly a better EFX allocation, where the individual and the total welfare are better, namely $X_1 \leftarrow \{g_1\}$ and $X_2 \leftarrow \{g_2\}$.

Nash welfare of an allocation X defined as the geometric mean of the valuations of the agents, $(\prod_{i\in[n]} v_i(X_i))^{1/n}$, is a popular measure of economic efficiency.[1] In fact, when agents have additive valuations, then the allocation with the highest Nash welfare is also EF1 (another popular fairness notion weaker than EFX). Unfortunately, maximizing Nash welfare is Approximable hard. However, there have been several approximation algorithms (Anari et al. [7], Barman et al. [10], Cole and Gkatzelis [22]) that give a constant factor approximation. The best approximation ratio is $e^{1/e} \approx 1.445$, given by Barman et al. [10].

Similar to the algorithm in Chaudhury et al. [20], we show that with minor modifications to our main algorithm, we can determine an allocation that satisfies the conditions in Theorem 1 and simultaneously achieves a $2e^{1/e} \approx 2.88$ approximation of the Nash welfare (i.e., in polynomial time, we can find efficient $(1-\varepsilon)$-EFX allocation with a sublinear number of unallocated goods).

Theorem 4. *For all $\varepsilon \in (0, 1/2]$, we can determine a partial $(1-\varepsilon)$-EFX allocation X and a set of unallocated goods P in polynomial time such that $|P| \in \mathcal{O}((n/\varepsilon)^{4/5})$ and $NW(X) \ge (1/2.88) \cdot NW(X^*)$, where X^* is the allocation with the highest Nash welfare.*

1.4. Further Related Work

Because the fair division literature is too vast, we restrict here to previous work that appears most relevant and refer the reader to recent surveys (Amanatidis et al. [6], Walsh [42]).

Fair division has received significant attention since the seminal work of Steinhaus [41] in the 1940s. Other than envy freeness, another fundamental fairness notion is that of *proportionality*. Recall that, in an envy-free allocation,

Figure 2. (Color online) Bipartite graph $G[V_1 \cup V_2 \cup V_3]$ when showing $R(2) \le 2$.

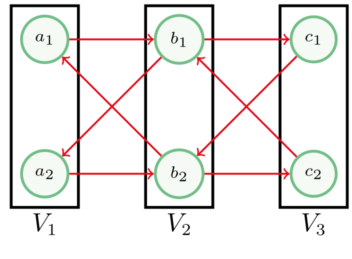

every agent values her own bundle at least as much as she values the bundle of any other agent. However, in a proportional allocation, each agent gets a bundle that is worth $1/n$ times her valuation on the entire set of goods. Because envy freeness and proportionality cannot always be guaranteed while dividing indivisible goods, various relaxations of the same have been studied. Alongside EFX, another popular relaxation of envy freeness is *envy freeness up to one good* (EF1), where no agent envies another agent following the removal of *some* good from the other agent's bundle. Although the existence of an EFX allocation is open, EF1 allocations are known to exist for any number of agents, even when agents have weakly monotone valuation functions (Lipton et al. [34]). Although 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 [16]. Although MMS allocations do not always exist (Kurokawa et al. [32]), there has been extensive work to come up with approximate MMS allocations (Amanatidis et al. [4], Barman and Krishnamurthy [9], Bouveret and Lemaître [14], Budish [16], Garg and Taki [29], Garg et al. [30], Ghodsi et al. [31], Kurokawa et al. [32]). Some works assume ordinal ranking over the goods as opposed to cardinal values (e.g., Aziz et al. [8], Brams et al. [15]).

Alongside fairness, the efficiency of an allocation is also a desirable property. Two common measures of efficiency are that of Pareto optimality and Nash welfare. Caragiannis et al. [18] showed that any allocation that has the maximum Nash welfare is guaranteed to be Pareto optimal (efficient) and EF1 (fair). Barman et al. [10] give a pseudopolynomial algorithm to find an allocation that is both EF1 and Pareto optimal. Other works explore relaxations of EFX with high Nash welfare (Caragiannis et al. [17], Chaudhury et al. [20]).

A one-page abstract of our work appeared in Chaudhury et al. [21].

The rest of the paper is organized as follows. In Section 2, we briefly highlight our main techniques used to prove our main results (Theorems 1–3). Then, in Section 3, we outline the basic concepts, notations, and techniques from existing literature on EFX allocations that will be useful to prove our main results. In Sections 4 and 5, we give the proofs of Theorem 3 and Theorem 2, respectively. In Section 6, we show how a minor modification of our main algorithm helps us achieve our main result (Theorem 1) with high Nash welfare (efficiency guarantees). Finally, in the appendix, we show why the technique from Chaudhury et al. [19] does not extend to a setting with four agents with additive valuations.

2. Our Techniques

In this section, we give a brief overview of our key ideas and techniques. We first sketch the key idea that relates the number of unallocated goods to the function $R(d)$ (Theorem 3), and then, we briefly show that $R(d)$ is finite.

2.1. Relation Between the Number of Unallocated Goods and the Rainbow Cycle Number

A very crucial concept that is often used while studying relaxations of envy freeness in discrete fair division is the *envy graph of an allocation*. Given an allocation $X = \langle X_1, X_2, \ldots, X_n \rangle$, the envy graph E_X has vertices corresponding to the agents, and there is an edge from agent i to agent j in E_X if agent i envies agent j ($v_i(X_i) < v_i(X_j)$). Without loss of generality, one assumes that the envy graph of an allocation is acyclic. If there is a cycle, then one can shift the bundles along the cycle, thereby giving every agent in the cycle a strictly better bundle, and the other agents retain their previous bundle. Such a procedure reduces the number of edges in the envy graph, and one can continue this until E_X is cycle free.

Most of the algorithms that have been used to prove the existence of relaxations of EFX allocations (Chaudhury et al. [19], Chaudhury et al. [20], Plaut and Roughgarden [39]) maintain a *relaxed EFX allocation*[2] X on the set of allocated goods, and as long as the envy graph E_X and the set of unallocated goods satisfy some "properties," they determine another relaxed EFX allocation X', in which $\phi(X') \geq \phi(X) + \delta$ for some $\delta \geq 1$, where ϕ is an integral upper-bounded function. In that case, we say that the relaxed EFX allocation X' *dominates* the relaxed EFX allocation X. Because ϕ is integral and upper bounded, such a procedure will finally converge to a relaxed EFX allocation where the envy graph E_X and the unallocated goods will not satisfy the said properties, and this will be the final allocation of the algorithms.

We now highlight another crucial concept used in these algorithms. The envy graph E_X does not provide any information on an agent's valuations of the bundles formed by adding unallocated goods to the current bundles of the allocation. This information is crucial when we want to create another dominating relaxed EFX allocation by allocating some of the unallocated goods and unallocating some of the already allocated goods. The algorithms in Chaudhury et al. [19] and Chaudhury et al. [20] make use of this information through other concepts. For instance, Chaudhury et al. [19] and Chaudhury et al. [20] define *champions*[3] and *champion graphs*. Given an allocation

Figure 3. (Color online) Illustration of a group champion graph.

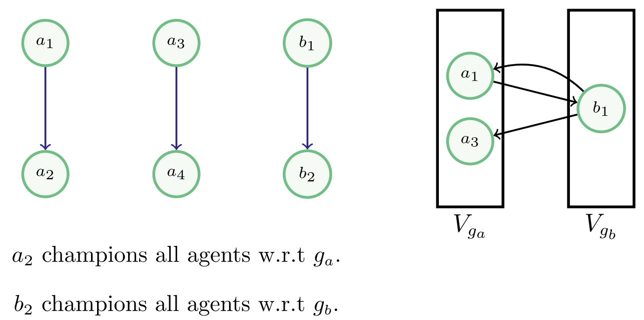

Notes. We have an instance with six agents $\cup_{i\in[4]}a_i$ and $\cup_{i\in[2]}b_i$ and two unallocated goods, namely g_a and g_b. The agents $\cup_{i\in[4]}a_i$ find g_a valuable, and the agents $\cup_{i\in[2]}b_i$ find g_b valuable. The envy graph E_X of the instance is shown in the left panel. E_X shows that $s(a_2) = a_1$, $s(a_4) = a_3$, and $s(b_2) = b_1$. Also, we have that agent a_2 champions all the agents with respect to (w.r.t.) g_a and that b_2 champions all the agents w.r.t. g_b. The group champion graph (right panel) has two parts: V_{g_a} corresponding to g_a and V_{g_b} corresponding to g_b. V_{g_a} contains the sources of all the agents who find g_a valuable, namely a_1 and a_3. Similarly, V_{g_b} contains b_1. There is an edge from a_1 to b_1 as a_2 (which is reachable from a_1 in E_X) champions b_1 w.r.t. g_a. Similarly, there is an edge from b_1 to a_1 and a_3 as b_2 (which is reachable from b_1 in E_X) champions a_3 w.r.t. g_b.

X and an unallocated good g, we say that an agent i is a champion for agent j w.r.t. g if there is a set $S \subseteq X_j \cup \{g\}$ such that $v_i(X_i) < (1-\varepsilon)\cdot v_i(S)$ and no agent (including i and j) envies S up to a factor of $(1-\varepsilon)$, following the removal of a single good (i.e., for all $\ell \in [n]$, we have $(1-\varepsilon)\cdot v_\ell(S \setminus \{h\}) \le v_\ell(X_\ell)$ for all $h \in S$).[4] A champion graph w.r.t. an unallocated good g has vertices corresponding to the agents (similar to the envy graph), and there is an edge from agent i to agent j if agent i champions agent j w.r.t. g. Depending on the configuration of the envy graph and the champion graphs (one for each unallocated good), the current $(1-\varepsilon)$-EFX allocation X is transformed into another $(1-\varepsilon)$-EFX allocation X' such that X' dominates X. However, when the number of agents is large, there are several different possible configurations of the champion graphs and the envy graph, and it is very hard and tedious to come up with better update rules. In this paper, we introduce the notion of a *group champion graph*, which is significantly more insightful and well structured than the champion graphs.

Given a $(1-\varepsilon)$-EFX allocation X and a set of unallocated goods M', we define the group champion graph. To this end, for each agent $a \in [n]$, we assign a unique source $s(a)$ in E_X such that a is reachable from $s(a)$ in E_X (if there are multiple sources from which a is reachable in E_X, then pick one source arbitrarily). The group champion graph of M' is a $|M'|$-partite graph $G = (\cup_{g\in M'} V_g, E)$, in which each part V_g contains a copy of the assigned sources of all the agents who find g "valuable." An agent a finds g valuable if $v_a(\{g\}) > \varepsilon \cdot v_a(X_a)$. There is an edge from vertex $s(a)$ in V_g to $s(a')$ in V_h if and only if a champions $s(a')$ w.r.t. g (see Figure 3 for an illustration). At a high level, the group champion graph encodes the most relevant information from all the champion graphs. We make this point more explicit by briefly explaining how group champion graphs help us prove Theorem 3.

We first observe that if there is an unallocated good g and an agent i such that the other agents do not envy $(X_i \cup \{g\}) \setminus g'$ up to a factor of $(1-\varepsilon)$ for all $g' \in X_i \cup \{g\}$, then we allocate g to i. Thus, we assume that for each unallocated good g and each agent i, there is an agent j that envies $(X_i \cup \{g\}) \setminus g'$ up to a factor of $(1-\varepsilon)$ for some $g' \in X_i \cup \{g\}$. In particular, this implies that every unallocated good is valuable to some agent because if there is a good g that is not valuable to any agent (i.e., $v_i(\{g\}) \le \varepsilon \cdot v_i(X_i)$ for all $i \in [n]$), then we can simply allocate g to a source s in E_X as no agent will envy the bundle $X_s \cup \{g\}$ up to a factor of $(1-\varepsilon)$; for all $i \in [n]$, we have that $v_i(X_i) \ge v_i(X_s)$ (as s is unenvied) and $\varepsilon \cdot v_i(X_i) \ge v_i(\{g\})$, implying that $(1+\varepsilon)v_i(X_i) \ge v_i(X_s \cup \{g\})$ and further implying that $v_i(X_i) \ge (1-\varepsilon)\cdot v_i(X_s \cup \{g\})$. Now, we classify the set of unallocated goods into two categories depending on how many agents find them valuable. We fix an integer $d < n$ and define "high-demand goods" and "low-demand goods." A high-demand good is valuable to more than d agents, and a low-demand good is valuable to at most d agents. We show in Section 4 that if the number of high-demand goods is more than $2n/(\varepsilon \cdot d)$, then we can determine a dominating $(1-\varepsilon)$-EFX allocation from the existing $(1-\varepsilon)$-EFX allocation. Thus, we may assume that the number of high-demand goods is at most $2n/(\varepsilon d)$. We now bound the number of low-demand goods. Let M'' be the set of low-demand goods. We construct the group champion graph $G = (\cup_{g\in M''} V_g, E)$ of M'', in which part V_g contains the assigned sources of the agents who find g valuable. Note that for all $g \in M''$, g is not valuable to more than d agents. Thus, $|V_g| \le d$ for all $g \in M''$. Now, consider any two parts V_g and V_h in G. By our assumption, for all $a \in V_h$, there is an agent who envies $(X_a \cup \{g\}) \setminus g'$ up to a factor of $(1-\varepsilon)$ for some good $g' \in X_a \cup \{g\}$, implying that for each a in V_h, there are agents who champion a w.r.t. g. We prove in Section 4 that because a is a source in E_X, the agents who champion a w.r.t. g must find g valuable. Therefore, for all $a \in V_h$, there is a source $s(a') \in V_g$, where a' champions a w.r.t. g. Thus, every vertex in V_h has an incoming edge from a vertex in V_g. In Section 4, we further show that whenever G has a cycle that visits each part at most once, then we can determine a $(1-\varepsilon)$-EFX allocation that dominates X. Therefore, we can assume that G has no cycle that visits each part at most once. Because G is an $|M''|$-partite graph that satisfies the conditions in Definition 1, we have that the number of low-demand goods is $|M''| \le R(d)$. Therefore, the total number of unallocated goods is $2n/(d\varepsilon) + R(d) \in \mathcal{O}(max(2n/(\varepsilon \cdot d), R(d)))$. By choosing the appropriate value for d, we arrive at the statement of Theorem 3.

We now elaborate that $R(d)$ is indeed upper bounded, which then establishes the existence of $(1-\varepsilon)$-EFX allocations with a sublinear number of unallocated goods.

2.2. Upper Bounds on the Rainbow Cycle Number

We briefly show that for any $d \in \mathbb{N}$, $R(d)$ is finite. Consider a k-partite graph $G = (\cup_{i\in[k]} V_i, E)$ in Definition 1. For all $i \in [k]$, let $V_i = \{(i,1),(i,2),\ldots,(i,|V_i|)\}$. For all $i<j$ and $i'<j'$, we say that the directed bipartite graphs $G[V_i \cup V_j]$ and $G[V_{i'} \cup V_{j'}]$ have the same configuration if and only if for each directed edge from vertex (i,a) to (j,b) (and equivalently, from (j,b') to (i,a')) in $G[V_i \cup V_j]$, there is an edge from (i',a) to (j',b) (and equivalently, from (j',b') to (i',a')) in $G[V_{i'} \cup V_{j'}]$ and vice versa. We first show that if there are $4d$ parts in G, say without loss of generality (w.l.o.g.) $V_1, V_2, \ldots, V_{4d}$, such that the induced directed bipartite graph $G[V_i \cup V_j]$ has the same configuration for all $1 \le i < j \le 4d$, then there exists a cycle in G that visits each part at most once.

Consider the parts V_1 and V_2 and the induced directed bipartite graph $G[V_1 \cup V_2]$. Because every vertex in one part has an incoming edge from a vertex in the other part, $G[V_1 \cup V_2]$ is cyclic. Let the simple cycle be $C = (1,i_1) \to (2,i_2) \to (1,i_3) \to \cdots \to (2,i_{2\beta}) \to (1,i_1)$ for some $\beta \le d$. Because all the induced bipartite graphs $G[V_i \cup V_j]$ have the same configuration for all $1 \le i < j \le 4d$, we can claim that for all $\ell \in [\beta]$, for each edge $(1,i_{2\ell-1}) \to (2,i_{2\ell})$ in C, there is an edge from $(2\ell-1, i_{2\ell-1})$ to $(4d-\ell, i_{2\ell})$ in $G[V_{2\ell-1}, V_{4d-\ell}]$ (note that $2\ell-1 < 4d-\ell$ as $\ell \le \beta \le d$). Similarly for all $\ell \in [\beta]$, for each edge $(2,i_{2\ell}) \to (1,i_{2\ell+1})$ in C ($2\beta+1$ is to be interpreted as one), there is an edge from $(4d-\ell, i_{2\ell})$ to $(2\ell+1, i_{2\ell+1})$ in $G[V_{2\ell+1}, V_{4d-\ell}]$ (again, note that $2\ell+1 < 4d-\ell$ as $\ell \le \beta \le d$). This implies that there is a cycle $C' = (1,i_1) \to (4d-1,i_2) \to (3,i_3) \to (4d-2,i_4) \to \cdots \to (4d-\beta, i_{2\beta}) \to (1,i_1)$ in G. Clearly, C visits each part of G at most once. Therefore, there cannot be $4d$ parts in G such that the induced directed bipartite graph $G[V_i \cup V_j]$ has the same configuration for all $1 \le i < j \le 4d$.

We now rephrase the question about an upper bound on $R(d)$. Let $\mathcal{D}$ be the set of all configurations of a directed bipartite graph, where the number of vertices in each part is at most d and every vertex has an incoming edge. We treat $\mathcal{D}$ as a set of *colors* and note that $|\mathcal{D}| \in 2^{\mathcal{O}(d^2)}$. Now, consider a complete graph K_n with vertex set $[n]$, where the vertex $\ell \in [n]$ corresponds to part V_n in G. For all $1 \le i < j \le n$, we color/label the edge (i,j) in K_n with a color from $\mathcal{D}$. The color on the edge (i,j) corresponds to the configuration of the directed bipartite graph $G[V_i \cup V_j]$. Clearly, $R(d)$ must be strictly smaller than the largest n such that every coloring of the edges of K_n with colors from $\mathcal{D}$ contains a monochromatic clique of size $4d$. This value of n corresponds to the *(multicolor) Ramsey number* (Diestel [26]) $\mathcal{R}(n_1, n_2, \ldots n_{|\mathcal{D}|})$, in which $n_i = 4d$ for all $i \in [|\mathcal{D}|]$. This number is finite, and the current best-known upper bounds on it are exponential in $|\mathcal{D}|$ and d (Conlon and Ferber [23], Diestel [26], Erdős and Szekeres [27], Lefmann [33]). Therefore, $R(d)$ is also bounded. However, this upper bound is very large and only provides a weak version of Theorem 1. This necessitates the study of finding "good" upper bounds on $R(d)$: in particular, upper bounds that are polynomial in d. We address this in Section 5 by showing that $R(d) \in \mathcal{O}(d^4)$.

3. Preliminaries and Tools

A fair division instance is given by the three tuple $\langle [n], M, \mathcal{V}\rangle$, where $[n]$ is the set of agents, M is the set of indivisible goods, and $\mathcal{V} = \{v_1(), v_2(), \ldots, v_n()\}$, where each $v_i : 2^M \to \mathbb{R}_{\ge 0}$ denotes the valuation function of agent i. We assume that agents have *additive valuations* (i.e., for all $i \in [n]$, we have $v_i(S) = \sum_{g\in S} v_i(\{g\})$ for all $S \subseteq M$). For the ease of notation, we write $v_i(g)$ instead of $v_i(\{g\})$ and similarly, $v_i(S \cup g)$ for $v_i(S \cup \{g\})$. We assume that $v_i(g)$ can be accessed in constant time for any i and g. For a fixed $0 < \epsilon < 1$ and an allocation $X = (X_1, \ldots, X_n)$, we say that an agent i

- *envies* a set S of goods if $v_i(X_i) < v_i(S)$;
- *heavily envies* a set S of goods if $v_i(X_i) < (1-\varepsilon) v_i(S)$;
- *strongly envies* a set S of goods if it heavily envies a proper subset of S; and
- is a *most envious agent* for a set S of goods if there exists a subset $Z \subseteq S$ such that i heavily envies Z and no agent strongly envies Z. The pair (i, Z) is called a *most-envious-agent-witness pair* for S. We emphasize that the most envious agent of the set S is not necessarily the agent with the highest envy for S, but it is the agent who envies a subset of S that no other agent strongly envies.

An agent envies (heavily envies, strongly envies) an agent j if it has these feelings for the set X_j. Clearly, strong envy implies heavy envy implies envy. An allocation X' *ε-strongly Pareto dominates* an allocation X or equivalently, $X' >_{PD(\varepsilon)} X$ if and only if $v_i(X'_i) \geq v_i(X_i)$ for all $i \in [n]$ and for some agent $j \in [n]$, we have $(1-\varepsilon)\cdot v_j(X'_j) \geq v_j(X_j)$.

At a high level, our algorithm is similar to previous algorithms used to prove the existence of relaxations of EFX allocations (Chaudhury et al. [19], Chaudhury et al. [20], Plaut and Roughgarden [39]). Our algorithm always maintains a $(1-\varepsilon)$-EFX allocation on the set of allocated goods, and as long as the current allocation and the set of unallocated goods P satisfy "some properties," it determines another $(1-\varepsilon)$-EFX allocation that ε-strongly Pareto dominates the previous $(1-\varepsilon)$-EFX allocation. Because the valuation of an agent for the entire good set is bounded, this procedure will eventually converge to a $(1-\varepsilon)$-EFX allocation, where the current allocation and the set of unallocated goods do not satisfy these properties. The bulk of the effort goes into determining the right properties under which one can come up with update rules that transform one $(1-\varepsilon)$-EFX allocation into a "better" $(1-\varepsilon)$-EFX allocation. We briefly recollect the update rules used in Chaudhury et al. [20] and Lipton et al. [34].

3.1. Envy Cycle Elimination (Lipton et al. [34])

The *envy graph* E_X of a $(1-\varepsilon)$-EFX allocation X has the agents as its vertex set, and there is an edge from vertex i to vertex j in E_X if agent i envies agent j (i.e., $v_i(X_i) < v_i(X_j)$). The paper by Lipton et al. [34] shows that whenever E_X has a cycle, then one can determine another $(1-\varepsilon)$-EFX allocation X' in which no agent has a worse bundle and $E_{X'}$ is acyclic. Formally, we have Lemma 1.

Lemma 1 (Lipton et al. [34]). *Consider a $(1-\varepsilon)$-EFX allocation X. If there is a cycle in E_X, then in polynomial time, we can determine a $(1-\varepsilon)$-EFX allocation X' such that $v_i(X'_i) \geq v_i(X_i)$ for all $i \in [n]$, and $E_{X'}$ is acyclic.*[5]

3.2. Update Rules in Chaudhury et al. [20]

We modify the update rules in Chaudhury et al. [20] slightly, as we are dealing with $(1-\varepsilon)$-EFX allocations and not EFX allocations. These rules are more involved and make essential use of the concept of *a most envious agent*.

Lemma 2. *Consider an allocation X and a set $S \subseteq M$. If there is an agent who heavily envies the bundle S, then we can determine a most-envious-agent-witness pair (t, Z) for S in $\mathcal{O}(n\cdot|S|^2)$ time. If there is an agent who strongly envies S, then t strongly envies S.*

Proof. Let i be an agent who heavily envies S. We construct a sequence (t_ℓ, Z_ℓ) as follows; initially, we set t_1 to i and Z_1 to S. Assume that $(t_{\ell-1}, Z_{\ell-1})$ is defined. If no agent (including $t_{\ell-1}$) strongly envies $Z_{\ell-1}$, then we stop. Otherwise, let i' be an agent such that $v_{i'}(X_{i'}) < (1-\varepsilon)\cdot v_{i'}(Z_{\ell-1}\setminus\{g\})$ for some $g \in Z_{\ell-1}$. We set t_ℓ to i' and Z_ℓ to $Z_{\ell-1}\setminus\{g\}$ and continue. We will eventually stop as with every next pair in the sequence, the size of the set Z_ℓ decreases by one. Say we stop at ℓ^*. Then, we have an agent t_{ℓ^*} that heavily envies the subset Z_{ℓ^*} of S. Moreover, no agent strongly envies Z_{ℓ^*}. Thus, (t_{ℓ^*}, Z_{ℓ^*}) is a most-envious-agent-witness pair.

If there is an agent who strongly envies S, then $\ell \geq 1$, and hence, t_{ℓ^*} heavily envies a proper subset of S. Thus, t_{ℓ^*} strongly envies S.

It is clear that we can determine the pair in $\mathcal{O}(n\cdot|S|^2)$ time; the maximum length of the sequence constructed is $|S|+1$ as the size of the set $Z_\ell = |S|+1-\ell$. We need time $\mathcal{O}(n|S|)$ to determine $v_i(S)$ for all i and can update any such value in time $\mathcal{O}(1)$ after the removal of an element. For each value of ℓ, it takes $\mathcal{O}(n\cdot|Z_\ell|) = \mathcal{O}(n\cdot|S|)$ time to find $(t_{\ell+1}, Z_{\ell+1})$. Thus, the total time needed is $\mathcal{O}(n\cdot|S|^2)$. □

For an allocation X and a set S of goods that is heavily envied by some agent, let (t, Z) be the pair returned by the procedure in Lemma 2. Now, for notational convenience only, we introduce a slightly different definition of *champions*. We call t *the champion* of S and Z the corresponding witness. We now state the update rules.

3.3. Update Rule U_1 (Chaudhury et al. [20])

The first rule is the simplest. It is applicable whenever we can allocate an unallocated good to an unenvied agent (a source in E_X) without creating any strong envy. In this case, we simply allocate this good to the corresponding source. This creates another $(1-\varepsilon)$-EFX allocation where no agent gets a worse bundle and the number of unallocated goods decreases.

Lemma 3 (U_1 (Chaudhury et al. [20])). *Consider a $(1-\varepsilon)$-EFX allocation X. If there is a source s in E_X and an unallocated good g such that no agent strongly envies $X_s \cup g$, then $X' = \langle X_1, X_2, \ldots, X_s \cup g, \ldots, X_n\rangle$ is a $(1-\varepsilon)$-EFX allocation and $v_i(X'_i) \geq v_i(X_i)$ for all $i \in [n]$.*

Note that there can be at most m consecutive applications of this rule as the number of unallocated goods decreases by one every time we apply this update rule. The remaining rules are applicable whenever there are either "valuable" goods unallocated or if "too many" goods are unallocated. We state the second update rule.

3.4. Update Rule U_2 (Chaudhury et al. [20])

This update rule is applicable if there is an agent $i \in [n]$ who heavily envies the set of unallocated goods P. In this case, let t be the champion of P and Z be the corresponding witness. In X', one assigns Z to t and changes the pool to $X_t \cup (P \setminus Z)$. The resulting allocation X' is EFX and ε-strongly Pareto dominates X.

Lemma 4. (U_2 (Chaudhury et al. [20])). *Consider a $(1-\varepsilon)$-EFX allocation X, and let P be the set of unallocated goods. If there is an agent $i \in [n]$ that heavily envies P, then in polynomial time, we can determine a $(1-\varepsilon)$-EFX allocation $X' >_{PD(\varepsilon)} X$.*

The third update rule is a refinement of envy-cycle elimination.

3.5. Update Rule U_3 (Chaudhury et al. [20])

This rule is applicable whenever the number of unallocated goods is at least the number of agents. Chaudhury et al. [20]) shows that when the number of unallocated goods is larger than the number of agents and when rule U_1 is no longer applicable, then in polynomial time, we can find a set of sources $s_1, s_2, \dots, s_\ell$ in E_X; a set of unallocated goods $g_1, g_2, \dots, g_\ell$; and a set of agents $t_1, t_2, \dots, t_\ell$ such that each t_i is reachable from s_i in E_X, the paths from s_i to t_i for all $i \in [\ell]$ are disjoint, and t_i is the champion of $X_{s_{i+1}} \cup g_{i+1}$ (indices are modulo ℓ). Then, one essentially proceeds as in cycle elimination. Let $Z_{i+1} \subseteq X_{s_{i+1}} \cup g_{i+1}$ be the witness corresponding to t_i. For each I, one assigns Z_{i+1} to t_i, and to each agent on the path from s_i to t_i except for t_i, one assigns the bundle owned by the successor on the path. The resulting allocation X' is EFX and ε-strongly Pareto dominates X.

Lemma 5 (U_3 (Chaudhury et al. [20])). *Consider a $(1-\varepsilon)$-EFX allocation X. If there exists a set of sources $s_1, s_2, \dots s_\ell$ in E_X; a set of unallocated goods $g_1, g_2, \dots, g_\ell$; and a set of agents $t_1, t_2, \dots, t_\ell$ such that each t_i is reachable from s_i in E_X, the paths from s_i to t_i for all $i \in [\ell]$ are disjoint, and t_i is the champion of $X_{s_{i+1}} \cup g_{i+1}$ (indices are modulo ℓ), then in polynomial time, we can determine a $(1-\varepsilon)$-EFX allocation $X' >_{PD(\varepsilon)} X$.*

4. Relating the Number of Unallocated Goods to the Rainbow Cycle Number

In this section, we give the proof of Theorem 3 (i.e., we show how any upper bound on $R(d)$ allows us to obtain a $(1-\varepsilon)$-EFX with sublinear many goods unallocated). More precisely, we show that given a $(1-\varepsilon)$-EFX allocation X, if E_X is acyclic, the update rules U_1 and U_2 are not applicable, and the number of unallocated goods is larger than $4n/(\varepsilon \cdot \zeta(2n/\varepsilon))$, then rule U_3 is applicable. Therefore, for most of this section, we proceed under the assumption that

$$E_X \text{ is acyclic and the update rules } U_1 \text{ (Lemma 3) and } U_2 \text{ (Lemma 4) are not applicable.} \quad (*)$$

We start with some definitions. We first make an observation about the agents who could potentially strongly envy $X_s \cup g$, where s is a source in E_X and g is an unallocated good.

Observation 1. Consider an unallocated good g and any source s in E_X. If agent i heavily envies $X_s \cup g$, then g is valuable to agent i.

Proof. We have $v_i(X_s) \le v_i(X_i)$ because s is a source of E_X and $v_i(X_i) < (1-\varepsilon)v_i(X_s \cup g)$ because i heavily envies $X_s \cup g$. Thus, $v_i(X_i) < (1-\varepsilon)(v_i(X_i) + v_i(g))$, and hence, $(1-\varepsilon)v_i(g) > \varepsilon v_i(X_i)$. □

Note that under assumption (*) for each unallocated good g and each source s in the envy graph, there is an agent who strongly envies $X_s \cup g$ (because the conditions of the update rule U_1 in Lemma 3 are not satisfied). Thus, each unallocated good is valuable to some agent. Now, we make a classification of the unallocated goods based on the number of agents who find them valuable. To be precise, given an allocation X, we classify the unallocated goods into two categories: *high-demand goods* H_X and *low-demand goods* L_X. A good g belongs to H_X if it is valuable to at least $d+1$ agents and to L_X if it is valuable to at most d agents. We will choose the exact value of d later (right now, just think of it as any integer less than n). Observe that the set of unallocated goods $P = H_X \cup L_X$. To prove our claim, it suffices to show that when $|H_X| + |L_X| > 4n/(\varepsilon \cdot \zeta(2n/\varepsilon))$, the rule U_3 is applicable. To this end, we first make a simple observation about $|H_X|$.

Observation 2. Under assumption (*), we have $|H_X| < 2n/(\varepsilon \cdot d)$.

Proof. For each good $g \in H_X$, let η_g be the number of agents who find g valuable. By definition of H_X, we have that $\eta_g > d$, and hence, $\sum_g \eta_g > |H_X| d$. We next upper bound $\sum_g \eta_g$ by $n \cdot (2/\varepsilon)$ by showing that at most $2/\varepsilon$ unallocated goods are valuable to any agent.

Consider any agent i. By assumption (*) rule, U_2 is not applicable, and hence, the value of the unallocated goods to i is at most $1/(1-\varepsilon)v_i(X_i)$. This is at most $2v_i(X_i)$ because $\varepsilon \le 1/2$. Any valuable good has a value at least $\varepsilon v_i(X_i)$ for i. Thus, the number of unallocated goods valuable to i is at most $2/\varepsilon$. □

We next bound $|L_X|$. In particular, we show that $|L_X| \le R(d)$. To this end, we introduce the notion of *group champion graph G*.

4.1. Group Champion Graph

Recall that we are operating under assumption (*), and hence, E_X is acyclic. Given E_X and the sources in E_X, we fix an arbitrary total order $\prec$ among the sources. To each agent a, we assign a source $s(a)$ such that a is reachable from $s(a)$ in the envy graph E_X (note that a can coincide with $s(a)$). If a is reachable from multiple sources, we pick $s(a)$ to be the source with the highest rank in the order $\prec$. However, once picked, $s(a)$ is fixed and remains unique throughout our algorithm and its analysis. Let $k := |L_X|$. For each $g \in L_X$, let Q_g be the set of all agents who find g valuable. By definition of L_X, we have $|Q_g| \le d$ for all $g \in L_X$. We now define a k-partite graph $G = (\cup_{g\in L_X} V_g, E)$, in which the part V_g corresponding to g consists of copies of the sources assigned to the agents in Q_g: formally, $V_g = \{(g, s(a)) | a \in Q_g\}$. For any goods g and h and agents $a \in Q_g$ and $b \in Q_h$, there is an edge from $(g, s(a))$ in V_g to $(h, s(b))$ in V_h if and only if a is the champion of $X_{s(b)} \cup g$. We now make a claim about the set of edges between V_g and V_h in G for any $g, h \in L_X$.

Lemma 6. *Under assumption* (*), *consider any* $g, h \in L_X$. *Then, each vertex in* V_h *has an incoming edge from a vertex in* V_g.

Proof. Consider any vertex $(h, s(b)) \in V_h$. By assumption (*), there is an agent who strongly envies the bundle $X_{s(b)} \cup g$. Otherwise, rule U_1 would be applicable. By Observation 1, all agents who strongly envy $X_{s(b)} \cup g$ consider g valuable and hence, belong to Q_g. Let a be the champion of $X_{s(b)} \cup g$. By Lemma 2, a strongly envies $X_{s(b)} \cup g$ and hence, belongs to Q_g. Thus, there is an edge from $(g, s(a))$ in V_g to $(h, s(b))$ in V_h (by the construction of G). □

Now, we claim that the existence of a cycle that visits each part of G at most once would imply the existence of a $(1-\varepsilon)$-EFX allocation that ε-strongly Pareto dominates the existing $(1-\varepsilon)$-EFX allocation.

Lemma 7. *Given a cycle C in G that contains at most one vertex from each* V_g, *for all* $g \in L_X$, *we can determine a* $(1-\varepsilon)$*-EFX allocation* $X' >_{PD(\varepsilon)} X$ *in polynomial time.*

Proof. Let $C = (g_{i+1}, s_i) \to (g_{i+2}, s_{i+1}) \to \cdots \to (g_{j+1}, s_j) \to (g_{i+1}, s_i)$ be a cycle in G that visits each part at most once. It will become clear why we index the g's starting at $i+1$. Consider the sequence $s_i, s_{i+1}, \ldots, s_j$. If all the sources in this sequence are not distinct, there exists a contiguous subsequence $s_{i'}, s_{i'+1}, \ldots, s_{j'}$ where all the sources are distinct and $s_{j'+1} = s_{i'}$ with $i \le i' < j' \le j$ (index $j+1$ is to be interpreted as i).

We now work with the sequence $s_{i'}, s_{i'+1}, \ldots, s_{j'}$ where all the sources are distinct and $s_{j'+1} = s_{i'}$. For all $\ell \in [i'+1, j'+1]$, the existence of the edge $(g_\ell, s_{\ell-1}) \to (g_{\ell+1}, s_\ell)$ implies the existence of an agent $t_{\ell-1}$ such that $t_{\ell-1}$ is the champion of $X_{s_\ell} \cup g_\ell$ and $s(t_{\ell-1}) = s_{\ell-1}$ (i.e., $t_{\ell-1}$ is reachable from $s_{\ell-1}$ in E_X). Furthermore, note that the paths from $s_{\ell-1}$ to $t_{\ell-1}$ for all $\ell \in [i'+1, j'+1]$ are disjoint. Assume otherwise, and let there be an intersection between paths from s_a to t_a and from s_b to t_b and w.l.o.g. $s_a \prec s_b$. Note that because the paths intersect, both t_a and t_b are reachable from s_b, and $s_a \prec s_b$, we have $s(t_a) \ne s_a$, which is a contradiction. Because the sources $s_{i'}, s_{i'+1}, \ldots, s_{j'}$ are distinct, the agents $a_{i'}, a_{i'+1}, \ldots, a_{j'}$ are also distinct (as each agent has a unique source assigned). Therefore, we have distinct sources $s_{i'}, \ldots, s_{j'}$ in E_X; distinct goods $g_{j'+1}, g_{i'+1}, \ldots, g_{j'}$; and distinct agents $t_{i'}, \ldots t_{j'}$ that satisfy the conditions under which the update rule U_3 (Lemma 5) is applicable. By applying U_3, we can get a $(1-\varepsilon)$-EFX allocation $X' >_{PD(\varepsilon)} X$.

We clarify a boundary case of this analysis. Note that in principle, the length of the contiguous subsequence could be also one (i.e., $i' = j'$). In this case, it means that there is an agent $t_{i'}$, reachable from $s_{i'}$ in E_X, who is the champion of $X_{s_{i'}} \cup g_{i'+1}$ (i.e., the most envious agent of $X_{s_{i'}} \cup g_{i'+1}$ is reachable from s_i, and thus, we apply rule U_3 and get a $(1-\varepsilon)$-EFX allocation $X' >_{PD(\varepsilon)} X$). □

With Lemma 7, we are now ready to give an upper bound on $|L_X|$. Observe that $|L_X|$ equals the number of parts in G. Now, the question is how many parts can G have such that it does not admit a cycle that visits each part at most once. This is where we upper bound $|L_X|$ with the rainbow cycle number.

Lemma 8. *Consider a* $(1-\varepsilon)$*-EFX allocation X. If* $|L_X| > R(d)$, *there is a* $(1-\varepsilon)$*-EFX allocation* $X' >_{PD(\varepsilon)} X$.

Proof. Recall that $|L_X| = k$, where k is the number of parts in G. Note that each part of G corresponds to the sources assigned to the agents who find a particular good in L_X valuable (Q_g for some $g \in L_X$). By definition of

L_X, there are at most d agents who find a good in L_X valuable. Thus, each part has at most d vertices. Again, by Lemma 6, between any two parts V_g and V_h of G, each vertex in V_h has an incoming edge from a vertex in V_g. Therefore, by Definition 1, we have that if $k > R(d)$, then there exists a cycle C in G that visits each part at most once. Once we have C, by Lemma 7, we can determine a $(1-\varepsilon)$-EFX allocation $X' >_{PD(\varepsilon)} X$. □

Given a $(1-\varepsilon)$-EFX allocation X such that $|L_X| > R(d)$, Lemma 8 only gives the existence of a $(1-\varepsilon)$-EFX allocation $X' >_{PD(\varepsilon)} X$. However, to determine X' in polynomial time, one needs to find a cycle C in G that visits each part at most once when $|L_X| > R(d)$ in polynomial time. Let us remark that this is a nontrivial problem in general, reminiscent of the well-known κ-Path and κ-Cycle problems, which are nondeterministic polynomial-time complete (Cygan et al. [24]). Here, the input is a (di-)graph G and an integer k, and the objective is to determine if there is a path (cycle) on at least k-distinct vertices of the graph. These problems can be solved in $2^{\mathcal{O}(k)} \cdot \text{poly}(\mathsf{n})$ time using techniques based on color coding, hash functions, and splitters (Alon et al. [2], Cygan et al. [24], Naor et al. [38]). In particular, we can reduce κ-Path to the following problem in polynomial time. Find a k-path in a *colorful* graph on n vertices, whose vertices have been colored with $\mathcal{O}(\text{poly}(k) \cdot \log n)$ colors, such that every vertex of the k-path has a distinct color. However, for our purposes, the construction of the cycle C in G is a part of the proof of Theorem 6 (described in Section 5); we show that in polynomial time, one can find a cycle in a $(d^4 + d)$-partite digraph, in which each part has at most d vertices and for any two parts V and V' in the digraph, every vertex in V' has an incoming edge from some vertex in V and vice versa. This implies that if $|L_X| > d^4 + d$, then in polynomial time, we can determine a cycle C in G that visits each part at most once and then determine a $(1-\varepsilon)$-EFX allocation $X' >_{PD(\varepsilon)} X$ by applying U_3. This also implies that $R(d) \le d^4 + d$. Therefore, we have Lemma 9.

Lemma 9. *Consider a $(1-\varepsilon)$-EFX allocation X. If $|L_X| > d^4 + d$, then in polynomial time, we can determine a $(1-\varepsilon)$-EFX allocation $X' >_{PD(\varepsilon)} X$.*

4.2. Putting it Together

We give the existence proof and indicate the appropriate changes required for the polynomial-time algorithm. We start with an empty allocation, which is trivially a $(1-\varepsilon)$-EFX. Then, the algorithm iteratively maintains a $(1-\varepsilon)$-EFX allocation X and a pool of unallocated goods. In each iteration, the algorithm first makes E_X acyclic in polynomial time by Lemma 1. Thereafter, the algorithm checks whether any one of the update rules U_1 and U_2 is applicable. If U_1 is applicable, then it determines a $(1-\varepsilon)$-EFX allocation X', where $v_i(X'_i) \ge v_i(X_i)$ for all $i \in [n]$ and the number of unallocated goods reduces. If U_2 is applicable, then it determines a $(1-\varepsilon)$-EFX allocation $X' >_{PD(\varepsilon)} X$. If neither U_1 nor U_2 is applicable, then it determines the sets H_X and L_X. By Lemma 2, we have $|H_X| \le 2n/(\varepsilon \cdot d)$. If $|L_X| \le R(d) \le d^4 + d$, then it returns the allocation X. Otherwise, it determines a cycle that visits each part of G at most once and then determines $(1-\varepsilon)$-EFX allocation $X' >_{PD(\varepsilon)} X$ by applying update rule U_3, as in Lemma 8. If $|L_X| > d^4 + d$, the cycle can be determined in polynomial time. Therefore, when the algorithm terminates, we have that $|H_X| \le 2n/(\varepsilon \cdot d)$ and $|L_X| \le R(d) \le d^4 + d$, implying that the total number of unallocated goods is $|H_X| + |L_X| \le 2 \cdot \max(2n/(\varepsilon \cdot d), R(d)) \le 2 \cdot \max(2n/(\varepsilon \cdot d), 2d^4)$.

We now state the explicit value of d first for the existence proof. We choose d as the largest integer such that $R(d) \le 2n/(\varepsilon d)$ (i.e., $d = \zeta(2n/\varepsilon)$). Recall that $\zeta(2n/\varepsilon)$ is defined as the largest integer d such that $d \cdot R(d) \le 2n/\varepsilon$. Therefore, the number of unallocated goods is at most $4n/(\varepsilon \cdot \zeta(2n/\varepsilon))$.

For the algorithmic result, we choose d as the smallest integer such that $2n/(\varepsilon \cdot d) \le 2d^4$. Then, $d = \lceil (n/\varepsilon)^{1/5} \rceil$, and the number of unallocated goods is at most $4\lceil (n/\varepsilon)^{1/5} \rceil^4 \in \mathcal{O}((n/\varepsilon)^{4/5})$ as $n/\varepsilon \ge 1$.

It only remains to show that the algorithm will terminate. We prove a polynomial bound on the number of iterations. The bound applies to the existence and the algorithmic version. To this end, note that in each iteration, after removing cycles from E_X, our algorithm determines a new $(1-\varepsilon)$-EFX allocation X' through one of the following procedures:

- applying U_1,
- applying U_2, or
- determining a cycle C that visits each part in G at most once and then applying U_3.

Note that the initial envy-cycle elimination and subsequent application of all of the procedures ensure that $v_i(X'_i) \ge v_i(X_i)$ for all $i \in [n]$ (Lemmas 1 and 3–5). Thus, throughout the algorithm, the valuation of an agent never decreases. Note that there cannot be more than *m consecutive applications* of U_1, as the number of unallocated goods decreases with each application of U_1. Every time we apply U_2 or U_3, we ensure that $X' >_{PD(\varepsilon)} X$, implying that the valuation of some agent improves by a factor of at least $(1+\varepsilon)$. Because each agent's valuation is bounded by $W = max_{i \in [n]} v_i(M)$ and the valuation of an agent never decreases throughout the algorithm, we can have at most $\text{poly}(n, m, \log W, 1/\varepsilon)$ many iterations that involve applications of U_2 and U_3. Therefore, the total number of iterations of our algorithm is $m \cdot$ (iterations involving application of U_2 or U_3), which is also $\text{poly}(n, m, \log W, 1/\varepsilon)$.

Notice that in the algorithmic case, each of the iterations can also be implemented in polynomial time; U_1 and U_2 can be implemented in polynomial time (Lemmas 3 and 4). When $|L_X| \geq 2d^4 \geq d^4 + d$, then in polynomial time, we can determine the cycle C and apply U_3 (Lemma 9). We can now state the main result of this section.

Theorem 5. *There exists a* $(1-\varepsilon)$*-EFX allocation* X *and a set of unallocated goods* P *such that* $|P| \leq 4n/(\varepsilon \cdot \zeta(2n/\varepsilon))$. *In polynomial time, we can find a* $(1-\varepsilon)$*-EFX allocation and a set* P *of unallocated goods such that* $|P| \in \mathcal{O}((n/\varepsilon)^{4/5})$.

Note that any upper bound on the rainbow cycle number will imply an upper bound on the number of unallocated goods.

5. Bounds on the Rainbow Cycle Number

In this section, we give the proof of Theorem 2. We briefly recall the setup. There is a k-partite digraph $G = (\cup_{i\in[k]} V_i, E_G)$ such that each part has at most d vertices. For every pair of distinct parts V_i and V_j, every vertex in V_j has an incoming edge from some vertex in V_i. There is no cycle in G that visits each part at most once. Our goal is to establish an upper bound on k.

We now introduce some helpful notations and concepts. For each $i \in [k]$, we represent the vertices in the part V_i as $(i, \text{vertex id})$ (i.e., $V_i = \{(i,1),(i,2),\ldots,(i,|V_i|)\}$). For any positive integer d and $a, b \in [d]$, we use $\sigma_d(a,b)$ to denote $(a-1)\cdot d + b$. Note that $1 \leq \sigma_d(a,b) \leq d^2$. The $\sigma_d(a,b)$ captures the lexicographic ordering among the pairs $\cup_{a\in[d]} \cup_{b\in[d]} (a,b)$. For any Boolean vector $u \in \{0,1\}^r$, we use $u[k]$ to refer to the kth coordinate of the vector u. We introduce the simple yet crucial notion of *representative set for a set of Boolean vectors*. Given a set D of r-dimensional Boolean vectors, the set $B \subseteq D$ is a *representative set* of D if and only if $\{\ell | a[\ell]\} = 1$ for some $a \in D = \{\ell | b[\ell]\} = 1$ for some $b \in B$. We first make an observation about the size of B.

Observation 3. Given any set D of r-dimensional Boolean vectors, there exists a representative set $B \subseteq D$ of size at most r.

Proof. For each coordinate $\ell \in [r]$, we do; if there is a vector $a \in D$ with $a[\ell] = 1$, we put one such vector into B. Clearly, $|B| \leq r$. □

We prove Theorem 2 by contradiction. To be precise, we show that if $k > d^4 + d$, then there exists a cycle in G that visits every part at most once. Moreover, this cycle can be found in time polynomial in k.

We construct the cycle in two steps. We first show the existence of a part $V_{\tilde{\ell}}$ such that there is a directed cycle that visits only the parts $V_{\tilde{\ell}}, V_1, V_2, \ldots, V_d$ and moreover, each of the parts $V_1, V_2, \ldots, V_d$ at most once. In the second step, we replace the vertices in $V_{\tilde{\ell}}$ in this cycle by vertices in distinct parts.

For each ordered pair $(i,j) \in [d] \times [d]$ and $\ell \in [k] \setminus [d]$, we define a d^2-dimensional vector $u_{i,j,\ell}$ as follows; for all $x \in [d]$ and $y \in [d]$, we set $u_{i,j,\ell}[\sigma_d(x,y)] = 1$ if and only if there exists a path $(i,x) \to (\ell,z) \to (j,y)$ in G for some $(\ell,z) \in V_\ell$ (i.e., if there exists a path from vertex (i, x) in V_i to vertex (j, y) in V_j through some vertex in V_ℓ). Otherwise, we set $u_{i,j,\ell}[\sigma_d(x,y)] = 0$.

Let $\mathcal{L} = [k] \setminus [d]$. For each ordered pair $(i,j) \in [d] \times [d]$, we construct the sets $B^{i,j}$ and $\mathcal{L}^{i,j}$ as follows. For each (i, j) taken in the increasing order of $\sigma_d(i,j)$, define $\mathcal{L}^{i,j} = \mathcal{L}$ and $B^{i,j}$ as a representative vector set of $\{u_{i,j,\ell} | \ell \in \mathcal{L}^{i,j}\}$ of size at most d^2. A set $B^{i,j}$ of this size exists because our vectors have dimension d^2. Then, we set $\mathcal{L} = \mathcal{L} \setminus \{\ell | u_{i,j,\ell} \in B^{i,j}\}$. At most, d^2 elements are removed from $\mathcal{L}$ in each iteration.

For clarity, we write $\mathcal{L}^f$ to denote the set $\mathcal{L}$ at the end of the construction. Observe that $|\mathcal{L}^f| \geq 1$. This holds because we start with a set of size larger than d^4 and remove at most d^2 elements in each of the d^2 iterations.

Observation 4. Consider distinct ordered pairs $(i,j) \in [d] \times [d]$ and $(i',j') \in [d] \times [d]$. The sets $\{\ell | u_{i,j,\ell} \in B^{i,j}\}$ and $\{\ell | u_{i',j',\ell} \in B^{i',j'}\}$ are disjoint.

Proof. Let us assume without loss of generality that $\sigma_d(i,j) < \sigma_d(i',j')$. Consider any ℓ such that $u_{i,j,\ell} \in B^{i,j}$. Then, ℓ is removed from $\mathcal{L}$ at the end of the iteration for the pair (i, j) and hence, does not belong to $\mathcal{L}$ at the beginning of the iteration for the pair (i',j'). Consequently, $u_{i',j',\ell} \notin B^{i',j'}$ (by definition of $B^{i',j'}$, if $u_{i',j',\ell} \in B^{i',j'}$, then $\ell \in \mathcal{L}^{i',j'}$). □

At the end of the construction, we arbitrarily pick a $\tilde{\ell} \in \mathcal{L}^f$ (this is possible as $\mathcal{L}^f \neq \emptyset$). Now, we make a small observation about the vector $u_{i,j,\tilde{\ell}}$ for all $i,j \in [d]$.

Observation 5. For all $i,j \in [d]$, if $u_{i,j,\tilde{\ell}}[q] = 1$ for some $q \in [d^2]$, then there exists a vector $u_{i,j,l'} \in B^{i,j}$ such that $u_{i,j,l'}[q] = 1$.

Proof. Observe that $\mathcal{L}^f \subseteq \mathcal{L}^{i,j}$. Therefore, $\tilde{l} \in \mathcal{L}^{i,j}$. By definition, $B^{i,j}$ is a representative vector set of $\{u_{i,j,\ell} | \ell \in \mathcal{L}^{i,j}\}$. Therefore, by the definition of representative set, there exists a vector $u_{i,j,\ell'} \in B^{i,j}$ such that $u_{i,j,\ell'}[q] = 1$. □

We are now ready for the construction of a cycle that visits each part at most once. We first show that there exists a cycle C in G that visits only the parts $V_{\tilde{\ell}}, V_1, \dots, V_d$ and each of the parts $V_1, \dots, V_d$ at most once (i.e., the only part it may visit more than once is $V_{\tilde{\ell}}$). See Figure 4 for an illustration.

Let $(\tilde{\ell}, w_d)$ be an arbitrary vertex in $V_{\tilde{\ell}}$. We construct a path

$$(\tilde{\ell}, w_0) \rightarrow (1, v_1) \rightarrow \cdots \rightarrow (i-1, v_{i-1}) \rightarrow (\tilde{\ell}, w_{i-1}) \rightarrow (i, v_i) \rightarrow (\tilde{\ell}, w_i) \rightarrow \cdots \rightarrow (d, v_d) \rightarrow (\tilde{\ell}, w_d)$$

by starting at $(\tilde{\ell}, w_d)$ and tracing backward. We start in $(\tilde{\ell}, w_d)$. Assume that we already traced back to $(\tilde{\ell}, w_i)$ with $i = d$ initially. By the construction of G, there must be an edge from some vertex (i, v_i) in V_i to $(\tilde{\ell}, w_i)$ in $V_{\tilde{\ell}}$, and there must be an edge from some vertex $(\tilde{\ell}, w_{i-1})$ in $V_{\tilde{\ell}}$ to (i, v_i) in V_i. Thus, there is the path $(\tilde{\ell}, w_{i-1}) \rightarrow (i, v_i) \rightarrow (\tilde{\ell}, w_i)$ in G. We keep continuing this procedure until we reach $(\tilde{\ell}, w_0)$.

Because the part $V_{\tilde{\ell}}$ can have at most d vertices, by the pigeonhole principle, there must be i and j with $0 \le i < j \le d$ such that $w_i = w_j$. Let C be the subpath from $(\tilde{\ell}, w_i)$ to $(\tilde{\ell}, w_j)$: that is,

$$C = (\tilde{\ell}, w_i) \rightarrow (i+1, v_{i+1}) \rightarrow (\tilde{\ell}, w_{i+1}) \rightarrow \cdots \rightarrow (\tilde{\ell}, w_{j-1}) \rightarrow (j, v_j) \rightarrow (\tilde{\ell}, w_j).$$

Observe that C visits all the parts of G except $V_{\tilde{\ell}}$ at most once. We now show that by using "bypass" parts, we can make the cycle simple. For clarity, we rewrite C as

$$C = (i+1, v_{i+1}) \rightarrow (\tilde{\ell}, w_{i+1}) \rightarrow \cdots \rightarrow (\tilde{\ell}, w_{j-1}) \rightarrow (j, v_j) \rightarrow (\tilde{\ell}, w_j) \rightarrow (i+1, v_{i+1}).$$

5.1. Making the Cycle Simple

For all $q \in [i+1, j]$, consider the subpath

$$(q, v_q) \rightarrow (\tilde{\ell}, w_q) \rightarrow (q+1, v_{q+1})$$

of C (index $j+1$ is to be interpreted as $i+1$). The existence of such a subpath in G implies that $u_{q,q+1,\tilde{\ell}}[\sigma_d(v_q, v_{q+1})] = 1$. By Observation 5, we know that there is a vector $u_{q,q+1,\ell_q} \in B^{q,q+1}$ such that $u_{q,q+1,\ell_q}[\sigma_d(v_q, v_{q+1})] = 1$. This

Figure 4. (Color online) Illustration of the first part of the construction.

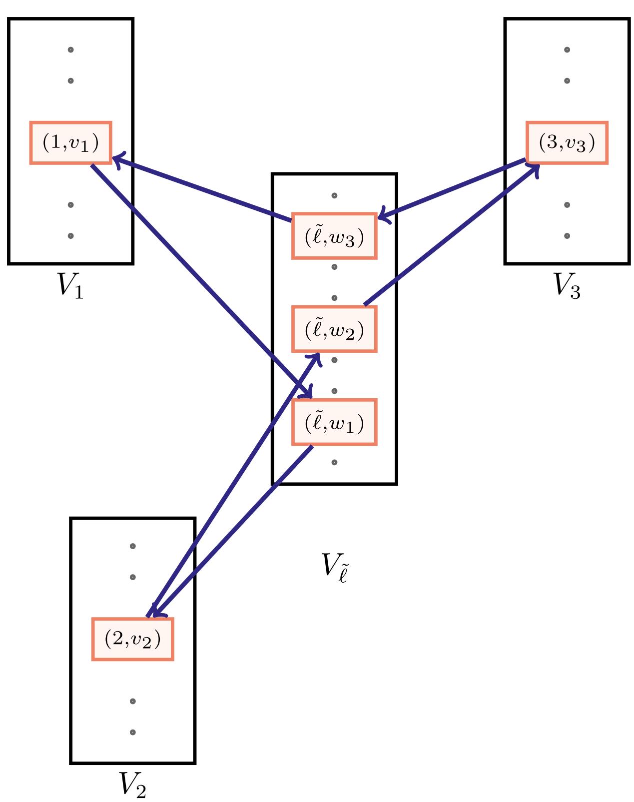

Notes. The cycle in the figure visits the parts V_1, V_2, and V_3 exactly once and the part $V_{\tilde{\ell}}$ three times. It is given by $(\tilde{\ell}, w_3) \rightarrow (1, v_1) \rightarrow (\tilde{\ell}, w_1) \rightarrow (2, v_2) \rightarrow (\tilde{\ell}, w_2) \rightarrow (3, v_3) \rightarrow (\tilde{\ell}, w_3)$.

implies that there exists a part V_{ℓ_q} and a vertex (ℓ_q, y_q) in part V_{ℓ_q} such that there is a subpath

$$(q, v_q) \to (\ell_q, y_q) \to (q+1, v_{q+1}).$$

By Observation 4, we have that $\ell_q \neq \ell_{q'}$ for all $q \neq q'$. Therefore, we have a simple cycle C' in G that visits each part in G at most once; namely,

$$C' = (i+1, v_{i+1}) \to (\ell_{i+1}, y_{i+1}) \to \cdots \to (\ell_{j-1}, y_{j-1}) \to (j, v_j) \to (\ell_j, y_j) \to (i+1, v_{i+1}).$$

See Figure 5 for an illustration of this entire procedure.

Therefore, if $k > d^4 + d$, then there exists a cycle in G that visits each part at most once. Moreover, this cycle can be found in time polynomial in k. With this, we arrive at the main result of this section.

Figure 5. (Color online) Illustration of the existence of a cycle that visits every part at most once.

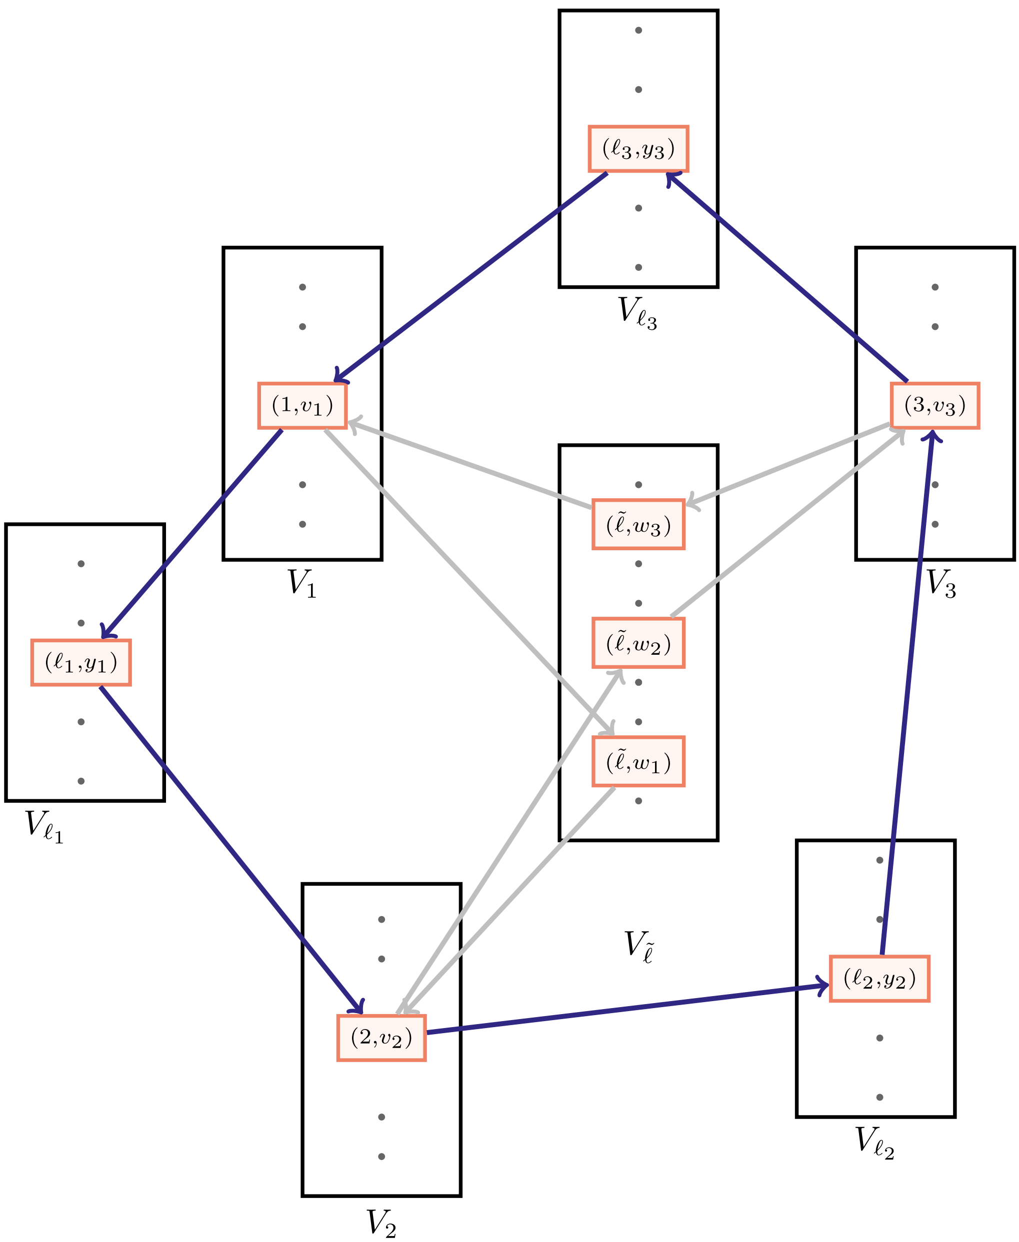

Notes. We take the instance in Figure 4, where there exists a cycle C that visits every part other than $V_{\tilde{\ell}}$ at most once. The edges of the cycle C are light in color. The figure shows how to obtain a cycle C' that visits every part at most once from C. The edges of C' are dark in color. For all $i \in [3]$, we replace the subpath in C of the form $(i, v_i) \to (\tilde{\ell}, w_i) \to (i+1, v_{i+1})$ ($3+1$ is to be interpreted as one) by $(i, v_i) \to (\ell_i, y_i) \to (i+1, v_{i+1})$ to get C'.

Figure 6. (Color online) Illustration of the construction of d-partite graph G that satisfies all the conditions in Definition 1 for $d = 2$ (left panel) and $d = 3$ (right panel).

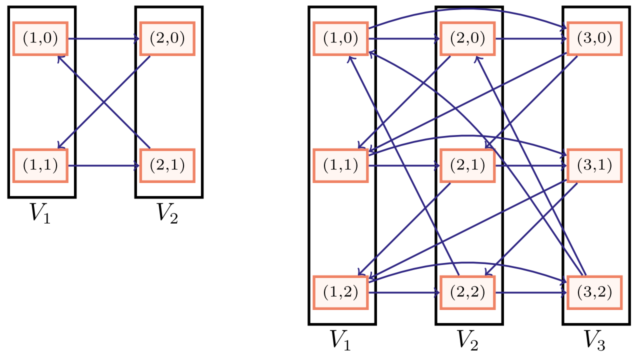

Theorem 6. *For all $d \geq 1$, we have $R(d) \leq d^4 + d$. Furthermore, let G be a k-partite digraph with $k > d^4 + d$ parts of cardinality at most d each such that for every vertex v and any part W not containing v, there is an edge from W to v. Then, there exists a cycle in G visiting each part at most once, and it can be found in time polynomial in k.*

An improved upper bound on $R(d)$ would imply a better bound on the number of unallocated goods. However, we show that an exponential improvement (e.g., $R(d) \in \text{poly}(\log(d))$) is not possible by showing a *linear* lower bound (i.e., $R(d) \geq d$)). However, this still leaves room for polynomial improvement, and we suspect that $R(d) \in \mathcal{O}(d)$. This would imply the existence of a $(1-\varepsilon)$-EFX allocation with $\mathcal{O}(\sqrt{n/\varepsilon})$ many goods unallocated. For a polynomial-time algorithm, the construction of a cycle as in Theorem 6 would have to be polynomial time. However, we remark that this is an initiation study for determining $(1-\varepsilon)$-EFX allocations with a sublinear number of unallocated goods, and we use concepts like the *group champion graph* that are natural extensions of the champion graph. We believe that this still leaves room for developing more sophisticated concepts and techniques that may reduce the number of unallocated goods to $o(\sqrt{n/\varepsilon})$.

5.2. Lower Bound on *R*(*d*)

We show that $R(d) \geq d$. We construct a d-partite graph $G = (\cup_{i\in[d]} V_i, E)$ such that each part V_i has d vertices; for all pairs of parts V_i and V_j, every vertex in V_j has an incoming edge from a vertex in V_i and vice versa; and there exists no cycle that visits each part at most once.

We now define the edges in G. Let $V_i = \{(i,0),(i,1),\ldots,(i,d-1)\}$. Consider any i and j such that $i<j$. For each $0 \leq \ell \leq d-1$, we have an edge from (i,ℓ) in V_i to (j,ℓ) in V_j, and there is an edge from (j,ℓ) in V_j to $(i,(\ell+1) \bmod d)$ in V_i (see Figure 6 for an illustration). One can easily verify that for all parts V_i and V_j, every vertex in part V_j has an incoming edge from part V_i and vice versa. It suffices to show that G admits no cycle that visits each part at most once.

Lemma 10. *There exists no cycle in G that visits each part at most once.*

Proof. We prove by contradiction. Assume that there is a cycle $C = (i_1,\ell_1) \to (i_2,\ell_2) \to \cdots \to (i_r,\ell_r) \to (i_1,\ell_1)$ that visits each part at most once (i.e., $i_x \neq i_y$ for all $x,y \in [r]$). From here on, all the indices are modulo r. Note that by the construction of the edges of G, for all $q \in [r]$, we have $\ell_{q+1} = \ell_q$ if $i_q < i_{q+1}$ and $\ell_{q+1} = (\ell_q + 1) \bmod d$ if $i_q > i_{q+1}$. Let $\#_1 = |\{q \in [r] | i_q > i_{q+1}\}|$ (recall that $r+1$ is one). The existence of the cycle C in G implies that $\ell_1 = (\ell_1 + \#_1) \bmod d$.

Because $i_x \neq i_y$ for all $x,y \in [r]$ and there exists the cycle C in G, there are indices q' and q'' such that $i_{q'} > i_{q'+1}$ and $i_{q''} < i_{q''+1}$, further implying that $1 \leq \#_1 \leq r-1$. Because G has d parts, we have $r \leq d$, implying that $1 \leq \#_1 \leq d-1$. However, this implies that $(\ell_1 + \#_1) \bmod d \neq \ell_1$, which is a contradiction. □

6. Finding *Efficient* $(1-\varepsilon)$-EFX Allocations with a Sublinear Number of Unallocated Goods

We note that like the algorithms in Chaudhury et al. [20] and Plaut and Roughgarden [39], our algorithm is flexible with the initialization (i.e., starting with any initial $(1-\varepsilon)$-EFX allocation X, it can determine a final $(1-\varepsilon)$-EFX allocation Y with at most $\mathcal{O}((n/\varepsilon)^{4/5})$ many goods unallocated and $v_i(Y_i) \geq v_i(X_i)$ for all $i \in [n]$). This is a consequence of the fact that the valuation of an agent never decreases throughout our algorithm. Therefore, our algorithm maintains the welfare of the initial allocation. Thus, if we choose the initial $(1-\varepsilon)$-EFX allocation carefully, we can also guarantee high Nash welfare for our final $(1-\varepsilon)$-EFX allocation with sublinear many goods unallocated. To this end, we use an important result from Caragiannis et al. [17] about determining partial EFX allocations with high Nash welfare in polynomial time.

Theorem 7 (Caragiannis et al. [17]). *In polynomial time, we can determine a partial EFX allocation X such that* $NW(X) \geq 1/(2.88) \cdot NW(X^*)$, *where* X^* *is the Nash welfare-maximizing allocation. In fact, the result in Caragiannis et al.* [17] *shows the existence of partial EFX allocations that achieve a* $1/2$ *approximation of the Nash welfare. However, in polynomial time, one can only find a partial EFX allocation with a* $1/2.88$ *approximation of the Nash welfare.*

Let X be the partial EFX allocation that achieves a 2.88 approximation of the Nash welfare. We run our algorithm starting with X as the initial allocation. The final $(1-\varepsilon)$-EFX allocation with sublinear many unallocated goods is also a 2.88 approximation of the Nash welfare as the valuations of the agents in the final allocation are at least their valuations in X. Therefore, we have the following theorem.

Theorem 8. *In polynomial time, we can determine a* $(1-\varepsilon)$*-EFX allocation X with* $\mathcal{O}((n/\varepsilon)^{4/5})$ *goods unallocated such that* $NW(X) \geq 1/(2.88) \cdot NW(X^*)$, *where* X^* *is the Nash welfare-maximizing allocation. Furthermore, using the existence of partial EFX allocations with* $1/2$ *approximation to Nash welfare* (*Caragiannis et al.* [17]), *there exists a* $(1-\varepsilon)$*-EFX allocation X with* $\mathcal{O}((n/\varepsilon)^{4/5})$ *goods unallocated such that* $NW(X) \geq 1/2 \cdot NW(X^*)$.

Acknowledgments

The authors thank the anonymous reviewers for their detailed reading and great suggestions, which improved the statements, the exposition, and the presentation of the main results presented in this paper. In fact, Masoud Seddighin (one of the reviewers) observed that the upper bound on $R(d)$ can be improved to $\mathcal{O}(d^3)$ through a subtle modification to our existing algorithm: In particular, the algorithm needs to construct the "bypass parts" ($B^{i,j}$'s) between the parts (V_i, V_{i+1}) instead of all (V_i, V_j) with $(i,j) \in [d] \times [d]$ as the final path uses parts only from $B^{i,i+1}$, except when $q = j$, where one can use the intermediate part (i.e., part $V_{\tilde{\ell}}$). We do not discuss this improvement in detail in our paper, as there are results by Akrami et al. [1] and Berendsohn et al. [12] that improve the bound in our paper using this observation and additional clever techniques.

Appendix. Limitations of the Approach in Chaudhury et al. [19]

In Chaudhury et al. [19], an algorithmic proof to the existence of an EFX allocation is shown for three agents with additive valuations. We briefly sketch the proof technique in Chaudhury et al. [19] and then, highlight why it does not work for determining $(1-\varepsilon)$-EFX allocations with just four agents for arbitrarily small ε. It would be interesting to investigate whether we can rule out proving existence of $(1-\varepsilon)$-EFX for larger values of ε (e.g., $\varepsilon = 1/3$) using the technique in Chaudhury et al. [19].

Let the three agents be a, b, and c, and for any allocation X, let $\phi(X)$ be the vector $\langle v_a(X_a), v_b(X_b), v_c(X_c)\rangle$. The algorithm starts with an empty allocation, which is trivially EFX, and as long as there is an unallocated good, the algorithm determines another EFX allocation X' such that $\phi(X')$ is lexicographically larger than $\phi(X)$ (i.e., either $v_a(X'_a) > v_a(X_a)$ or $v_a(X'_a) = v_a(X_a)$ and $v_b(X'_b) > v_b(X_b)$ or $v_a(X'_a) = v_a(X_a)$, $v_b(X'_b) = v_b(X_b)$, and $v_c(X'_c) > v_c(X_c)$). In this paper, we show that such a technique cannot be used to show the existence of $(1-\varepsilon)$-EFX allocations for four agents.

We remark that our instance builds on the instance in Chaudhury et al. [19] that is used to show the existence of a partial EFX allocation, which is not Pareto dominated by any complete EFX allocation. We now construct an instance I with four agents, say $\{a, b, c, d\}$ with additive valuations and nine goods $\{g_i \mid i \in [9]\}$. Let $\phi(X) = \langle v_a(X_a), v_b(X_b), v_c(X_c), v_d(X_d)\rangle$. We show a $(1-\varepsilon)$-EFX allocation X of eight goods among four agents. Then, we show in any complete $(1-\varepsilon)$-EFX allocation, that the valuation of agent a will be strictly less than (almost half of) her valuation in X. This shows that for any complete $(1-\varepsilon)$-EFX allocation Y, we have that $\phi(X)$ is lexicographically larger than $\phi(Y)$.

The full description of our instance is captured by Table A.1. We choose our $\varepsilon \ll 1$. The subinstance defined by the agents b, c, and d and the goods $\cup_{i\in[6]} g_i \cup g_9$ is the instance in Chaudhury et al. [19] used to show the existence of a partial EFX allocation,

Table A.1. An instance where showing that the technique in Chaudhury et al. [19] cannot be used to determine $(1-\varepsilon)$-EFX allocations with four agents.

	g_1	g_2	g_3	g_4	g_5	g_6	g_7	g_8	g_9
a	0	0	0	0	0	0	6	4	0
b	16	4	24	4	0	34	31	0	2
c	10	0	18	8	20	0	29	0	6
d	0	0	0	0	18	20	19	0	4

Notes. In particular, given a $(1-\varepsilon)$-EFX allocation X and the unallocated good g_9, there is no complete $(1-\varepsilon)$-EFX allocation where the valuation of agent a does not strictly decrease (i.e., in any complete $(1-\varepsilon)$-EFX allocations Y, we have $v_a(Y_a) < v_a(X_a)$).

which is not Pareto dominated by any complete EFX allocation. We now specify the allocation X:

$$X_a = \{g_7, g_8\} \qquad X_b = \{g_2, g_3, g_4\}$$
$$X_c = \{g_1, g_5\} \qquad X_d = \{g_6\}.$$

The good g_9 is unallocated. We will show that in any complete $(1-\varepsilon)$-EFX allocation, agent a cannot have both g_7 and g_8. This would imply that agent a's valuation in any final $(1-\varepsilon)$-EFX allocation is strictly less than her valuation in X (as agent a's valuation for all goods other than g_7 and g_8 is zero). We prove this claim by contradiction. So, assume that Y is a complete $(1-\varepsilon)$-EFX allocation and $\{g_7, g_8\} \subseteq Y_a$. Note that $v_b(g_7) = 31$, $v_c(g_7) = 29$, and $v_d(g_7) = 19$. Because Y_a contains at least one other good (namely, g_8), each of the agents b, c, and d needs to be allocated bundles that they value at least 31, 29, and 19, respectively.

First, consider the case that $g_6 \in Y_b$. Then, we have $v_b(Y_b) \ge 34$. Now, to ensure $v_d(Y_d) \ge 19$, we need to allocate g_5 and g_9 to d, as d values all the other goods zero. We are left with goods g_1, g_2, g_3, and g_4. In order to ensure $v_c(Y_c) \ge 29$, we definitely need to allocate g_1, g_3, and g_4 to c. Now, even if we allocate the remaining good g_2 to b, we have $v_b(Y_b) = v_b(\{g_2, g_6\}) = 38 < (1-\varepsilon)\cdot 40 = (1-\varepsilon)\cdot v_b(\{g_1, g_3\}) \le (1-\varepsilon)\cdot v_b(Y_c \setminus g_4)$. Therefore, b will strongly envy c. Thus, $g_6 \notin Y_b$.

If $g_6 \notin Y_b$ and $v_b(Y_b) \ge 31$, Y_b must contain g_3 (the total valuation for b of all the goods other than g_3, g_6, g_7, and g_8 is less than 31). Now, we consider some more subcases.

Let us first assume that $g_1 \in Y_b$. Because Y_b already contains g_1 and g_3, the goods that can be allocated to c and d are g_2, g_4, g_5, g_6, and g_9. In order to ensure $v_c(Y_c) \ge 29$, we need to allocate g_4, g_5, and g_9 to c. Now, even if we allocate all the remaining goods (g_2 and g_6) to d, we have $v_d(Y_d) = v_d(\{g_3, g_6\}) = 20 < (1-\varepsilon)\cdot 22 = (1-\varepsilon)\cdot v_d(\{g_5, g_7\}) \le (1-\varepsilon)\cdot v_d(Y_c \setminus g_4)$. Therefore, d will strongly envy c.

Thus, $g_1 \notin Y_b$. Because neither g_1 nor g_6 belong to Y_b, the only way to ensure that $v_b(Y_b) \ge 31$ is to at least allocate g_2, g_3, and g_4 to b (we can allocate more). Similarly, given that the goods not allocated yet are g_1, g_5, g_6, and g_9, the only way to ensure that $v_c(Y_c) \ge 29$ is to allocate at least g_1 and g_5 to c. Similarly, the only way to ensure $v_d(Y_d) \ge 19$ now is to allocate at least g_6 to d. Now, we only have to allocate g_9. We show that adding g_9 to any one of the existing bundles will cause a violation of the $(1-\varepsilon)$-EFX property.

- Adding g_9 to Y_a. b, c, and d strongly envy a as $v_b(Y_b) = 32 < (1-\varepsilon)\cdot 33 = (1-\varepsilon)\cdot v_b(\{g_7, g_9\}) \le (1-\varepsilon)\cdot v_b(Y_a \setminus g_8)$. Similarly, we have $v_c(Y_c) = 30 < (1-\varepsilon)\cdot 35 = (1-\varepsilon)\cdot v_c(\{g_7, g_9\}) \le (1-\varepsilon)\cdot v_c(Y_a \setminus g_8)$ and $v_d(Y_d) = 20 < (1-\varepsilon)\cdot 23 = (1-\varepsilon)\cdot v_d(\{g_7, g_9\}) \le (1-\varepsilon)\cdot v_d(Y_a \setminus g_8)$.
- Adding g_9 to Y_b. c strongly envies b as $v_c(Y_c) = 30 < (1-\varepsilon)\cdot 32 = (1-\varepsilon)\cdot v_c(\{g_3, g_4, g_7\}) = (1-\varepsilon)\cdot v_c(Y_b \setminus g_2)$.
- Adding g_9 to Y_c. d strongly envies c as $v_d(Y_d) = 20 < (1-\varepsilon)\cdot 22 = (1-\varepsilon)\cdot v_d(\{g_5, g_9\}) = (1-\varepsilon)\cdot v_d(Y_c \setminus g_1)$.
- Adding g_9 to Y_d. b strongly envies d as $v_b(Y_a) = 32 < (1-\varepsilon)\cdot 34 = (1-\varepsilon)\cdot v_b(g_6) = (1-\varepsilon)\cdot v_b(Y_d \setminus g_9)$.

This shows that $\{g_7, g_8\} \nsubseteq Y_a$ for any complete $(1-\varepsilon)$-EFX allocation Y. This implies that agent a's valuation in Y is strictly less than her valuation in X, implying that $\phi(X)$ is lexicographically larger than $\phi(Y)$. This shows that the approach from Chaudhury et al. [19] cannot be generalized to guarantee $(1-\varepsilon)$-EFX allocation when there are four or more agents.

Endnotes

[1] It implies other notions of efficiency, like *Pareto optimality*. An allocation $X = \langle X_1, \ldots, X_n \rangle$ is Pareto optimal if there is no allocation $Y = \langle Y_1, \ldots, Y_n \rangle$, where $v_i(Y_i) \ge v_i(X_i)$ for all $i \in [n]$ and $v_j(Y_j) > v_j(X_j)$ for some j.

[2] See $(1-\varepsilon)$-EFX allocation in Chaudhury et al. [19] and Chaudhury et al. [20] and 1/2-EFX allocation in Plaut and Roughgarden [39].

[3] They are called the "most envious agents" in Chaudhury et al. [20].

[4] Because we are dealing with $(1-\varepsilon)$-EFX allocations and not EFX allocations, we have changed the definition of champions and champion graphs appropriately. Chaudhury et al. [19] and Chaudhury et al. [20]) also use this definition in their algorithms as the polynomial-time algorithms also deal with $(1-\varepsilon)$-EFX allocations. Furthermore, for notational convenience, we will use a slightly different definition of champions in the future sections (mentioned in Section 3).

[5] Let C be an envy cycle. For each edge (i, j) of the cycle, one assigns in X' the bundle X_j to i. One continues in this way as long as there is a cycle in the envy graph.

References

[1] Akrami H, Alon N, Chaudhury BR, Garg J, Mehlhorn K, Mehta R (2023) EFX: A simpler approach and an (almost) optimal guarantee via rainbow cycle number. Leyton-Brown K, Hartline JD, Samuelson L, eds. *Proc. 24th ACM Conf. Econom. Comput., EC 23* (ACM, New York), 61.

[2] Alon N, Yuster R, Zwick U (1995) Color-coding. *J. ACM* 42(4):844–856.

[3] Amanatidis G, Markakis E, Ntokos A (2020) Multiple birds with one stone: Beating 1/2 for EFX and GMMS via envy cycle elimination. *Theoret. Comput. Sci.* 841:94–109.

[4] Amanatidis G, Markakis E, Nikzad A, Saberi A (2017) Approximation algorithms for computing maximim share allocations. *ACM Trans. Algorithms* 13(4):52:1–52:28.

[5] Amanatidis G, Birmpas G, Filos-Ratsikas A, Hollender A, Voudouris AA (2021) Maximum Nash welfare and other stories about EFX. *Theor. Comput. Sci.* 863:69–85.

[6] Amanatidis G, Aziz H, Birmpas G, Filos-Ratsikas A, Li B, Moulin H, Voudouris AA, Wu X (2023) Fair division of indivisible goods: Recent progress and open questions. *Artificial Intelligence* 322:103965.

[7] Anari N, Gharan SO, Saberi A, Singh M (2017) Nash social welfare, matrix permanent, and stable polynomials. Papadimitriou CH, ed. *8th Innovations Theoret. Comput. Sci. Conf. (ITCS)* (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Wadern, Germany), 1–12.
[8] Aziz H, Gaspers S, Mackenzie S, Walsh T (2015) Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227:71–92.
[9] Barman S, Krishnamurthy SK (2020) Approximation algorithms for maximin fair division. *ACM Trans. Economics and Comput.* 8(1):5:1–5:28.
[10] Barman S, Krishnamurthy SK, Vaish R (2018) Finding fair and efficient allocations. Tardos É, Elkind E, Vohra R, eds. *Proc. 19th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 557–574.
[11] Barman S, Krishnamurthy SK, Vaish R (2018) Greedy algorithms for maximizing Nash social welfare. André E, Koenig S, Dastani M, Sukthankar G, eds. *Proc. 17th Internat. Conf. Autonomous Agents MultiAgent Systems (AAMAS)* (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC/ACM, New York), 7–13.
[12] Berendsohn BA, Boyadzhiyska S, Kozma L (2022) Fixed-point cycles and approximate EFX allocations. Szeider S, Ganian R, Silva A, eds. *47th Internat. Sympos. Mathematical Foundations of Computer Science, MFCS 2022*, LIPIcs, vol. 241 (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Wadern, Germany), 17:1–17:13.
[13] Berger B, Cohen A, Feldman M, Fiat A (2022) Almost full EFX exists for four agents. *Thirty-sixth AAAI Conf. Artificial Intelligence, AAAI 2022, Thirty-fourth Conf. Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Sympos. Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event* (AAAI Press, Palo Alto, CA), 4826–4833.
[14] Bouveret S, Lemaître M (2016) Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents Multi-Agent Systems* 30(2):259–290.
[15] Brams SJ, Kilgour DM, Klamler C (2017) Maximin envy-free division of indivisible items. *Group Decision Negotiation* 26(1):115–131.
[16] Budish E (2011) The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. Political Econom.* 119(6):1061–1103.
[17] Caragiannis I, Gravin N, Huang X (2019) Envy-freeness up to any item with high Nash welfare: The virtue of donating items. Karlin A, Immorlica N, Johari R, eds. *Proc. 20th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 527–545.
[18] Caragiannis I, Kurokawa D, Moulin H, Procaccia AD, Shah N, Wang J (2016) The unreasonable fairness of maximum Nash welfare. Conitzer V, Bergemann D, Chen Y, eds. *Proc. 17th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 305–322.
[19] Chaudhury BR, Garg J, Mehlhorn K (2020) EFX exists for three agents. Biró P, Hartline JD, Ostrovsky M, Procaccia AD, eds. *Proc. 21st ACM Conf. Econom. Comput. (EC)* (ACM, New York), 1–19.
[20] Chaudhury BR, Kavitha T, Mehlhorn K, Sgouritsa A (2020) A little charity guarantees almost envy-freeness. Chawla S, ed. *Proc. 31st Sympos. Discrete Algorithms (SODA)* (SIAM, Philadelphia), 2658–2672.
[21] Chaudhury BR, Garg J, Mehlhorn K, Mehta R, Misra P (2021) Improving EFX guarantees through rainbow cycle number. Biró P, Chawla S, Echenique F, eds. *Proc. 22nd ACM Conf. Econom. Comput. (EC)* (ACM, New York), 310–311.
[22] Cole R, Gkatzelis V (2018) Approximating the Nash social welfare with indivisible items. *SIAM J. Comput.* 47(3):1211–1236.
[23] Conlon D, Ferber A (2021) Lower bounds for multicolor Ramsey numbers. *Adv. Math.* 378:107528.
[24] Cygan M, Fomin FV, Kowalik L, Lokshtanov D, Marx D, Pilipczuk M, Pilipczuk M, Saurabh S (2015) *Parameterized Algorithms*, vol. 5 (Springer, Cham, Switzerland).
[25] Darmann A, Schauer J (2015) Maximizing Nash product social welfare in allocating indivisible goods. *Eur. J. Oper. Res.* 247(2):548–559.
[26] Diestel R (2012) *Graph Theory*, 4th ed., Graduate Texts in Mathematics, vol. 5 (Springer, Cham, Switzerland).
[27] Erdős P, Szekeres G (1935) A combinatorial problem in geometry. *Compositio Mathdmatica* 2:463–470.
[28] Garg J, Murhekar A (2021) Computing fair and efficient allocations with few utility values. Caragiannis I, Hansen KF, eds. *Proc. 14th Internat. Sympos. Algorithmic Game Theory (SAGT)*, vol. 12885 (Springer, New York), 345–359.
[29] Garg J, Taki S (2021) An improved approximation algorithm for maximin shares. *Artificial Intelligence* 300:103547.
[30] Garg J, McGlaughlin P, Taki S (2019) Approximating maximin share allocations. Fineman JT, Mitzenmacher M, eds. *Proc. 2nd Sympos. Simplicity Algorithms (SOSA)*, vol. 69 (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Wadern, Germany), 1–20.
[31] Ghodsi M, Hajiaghayi MT, Seddighin M, Seddighin S, Yami H (2018) Fair allocation of indivisible goods: Improvements and generalizations. Tardos É, Elkind E, Vohra R, eds. *Proc. 19th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 539–556.
[32] Kurokawa D, Procaccia AD, Wang J (2018) Fair enough: Guaranteeing approximate maximin shares. *J. ACM* 65(2):8:1–8:27.
[33] Lefmann H (1987) A note on Ramsey numbers. *Studia Scientiarum Mathematicarum Hungarica.* 22(1–4):445–446.
[34] Lipton RJ, Markakis E, Mossel E, Saberi A (2004) On approximately fair allocations of indivisible goods. Breese JS, Feigenbaum J, Seltzer MI, eds. *Proc. 5th ACM Conf. Electronic Commerce (EC)* (ACM, New York), 125–131.
[35] Mahara R (2021) Extension of additive valuations to general valuations on the existence of EFX. *ESA LIPIcs*, vol. 204 (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Wadern, Germany), 66:1–66:15.
[36] Manurangsi P, Suksompong W (2021) Closing gaps in asymptotic fair division. *SIAM J. Discret. Math.* 35(2):668–706.
[37] Moulin H (2019) Fair division in the internet age. *Annual Rev. Econom.* 11(1):407–441.
[38] Naor M, Schulman LJ, Srinivasan A (1995) Splitters and near-optimal derandomization. *Proc. IEEE 36th Annual Foundations Comput. Sci.* (IEEE, Piscataway, NJ), 182–191.
[39] Plaut B, Roughgarden T (2020) Almost envy-freeness with general valuations. *SIAM J. Discret. Math.* 34(2):1039–1068.
[40] Procaccia AD (2020) Technical perspective: An answer to fair division's most enigmatic question. *Comm. ACM* 63(4):118.
[41] Steinhaus H (1949) Sur la division pragmatique. *Econometrica* 17:315–319.
[42] Walsh T (2020) Fair division: The computer scientist's perspective. Bessiere C, ed. *Proc. Twenty-Ninth Internat. Joint Conf. Artificial Intelligence (IJCAI)* (IJCAI, CA), 4966–4972.