Submitted to *Operations Research*
manuscript (Please, provide the manuscript number!)

EFX: A Simpler Approach and an (Almost) Optimal Guarantee via Rainbow Cycle Number

Hannaneh Akrami
Max Planck Institute for Informatics and Graduiertenschule Informatik, Universität des Saarlandes, hakrami@mpi-inf.mpg.de

Noga Alon
Princeton University, nalon@math.princeton.edu

Bhaskar Ray Chaudhury
University of Illinois at Urbana-Champaign, braycha@illinois.edu

Jugal Garg
University of Illinois at Urbana-Champaign, jugal@illinois.edu

Kurt Mehlhorn
Max Planck Institute for Informatics and Universität des Saarlandes, mehlhorn@mpi-inf.mpg.de

Ruta Mehta
University of Illinois at Urbana-Champaign, rutameht@illinois.edu

The existence of EFX allocations is a fundamental open problem in discrete fair division. Given a set of agents and indivisible goods, the goal is to determine the existence of an allocation where no agent envies another, following the removal of any single good from the other agent's bundle. Since the general problem has been elusive, progress is made on two fronts: (i) proving existence when the number of agents is small, and (ii) proving the existence of relaxations of EFX. In this paper, we improve and simplify the state-of-the-art results on both fronts with new techniques.

For the case of three agents, the existence of EFX was first shown with additive valuations (Chaudhury et al. 2020) and then extended to nice-cancelable valuations (Berger et al. 2022). Both results are obtained through an algorithm that moves in the space of partial EFX allocations, improving a certain potential as long as there are unallocated goods. However, the update rules to move from one partial EFX allocation to another are very involved, cumbersome, and fail if any one agent has a general monotone valuation function. As our first main result, we simplify and improve this result by showing the existence of EFX allocations when two of the agents have general monotone valuations and one has additive or, more generally, MMS-feasible valuation (a strict generalization of nice-cancelable valuation functions that subsumes additive, budget-additive and unit demand valuation functions). In contrast to the approaches in (Chaudhury et al. 2020, Berger et al. 2022), our new algorithm moves in the space of complete allocations, improving a potential as long as the allocation is not EFX. This approach is significantly simpler, as it also avoids using the standard concepts of envy-graph, champion-graph, and half-bundles and may find use in other fair-division problems.

Secondly, we consider relaxations of EFX allocations, namely, approximate EFX allocations and EFX allocations with few unallocated goods (charity). Through a promising new method using a problem in extremal combinatorics called Rainbow Cycle Number (RCN), Chaudhury et al. (2021a) managed to show the existence of $(1-\epsilon)$-EFX allocation with sub-linear charity, namely $\mathcal{O}((n/\epsilon)^{\frac{4}{5}})$ charity, where n is the number of agents. This is done by upper bounding the RCN by $\mathcal{O}(d^4)$ in d-dimension. They conjectured this number to be $\mathcal{O}(d)$ and gave a matching lower bound. We almost settle this conjecture by improving the upper bound to $\mathcal{O}(d \log d)$ and thereby get (almost) optimal charity of $\tilde{\mathcal{O}}((n/\epsilon)^{\frac{1}{2}})$ that is possible through this method. Our technique is simpler than the ones used for upper-bounding RCN in (Chaudhury et al. 2021a, Berendsohn et al. 2022, Jahan et al. 2023), and is based on the *probabilistic method*. We also derandomize the approach to construct such an allocation in deterministic polynomial time. Finally, we note that some of our techniques can be used to prove improved upper bounds on a problem in zero-sum combinatorics (Alon and Krivelevich 2021, Mészáros and Steiner 2021).

Key words: Discrete Fair Division, EFX Allocations, Rainbow Cycle Number

1. Introduction

Fair division of scarce resources is a fundamental problem in many disciplines, including computer science, economics, operations research, and social choice theory. In a classical fair division problem, the goal is to "fairly" allocate a set of goods among a set of agents (Steinhaus 1948). Such problems find very early historical mentions, for instance, in ancient Greek mythology and the Bible. Even more so today, many real-life scenarios are paradigmatic of the problems in this domain, e.g., division of family inheritance (Pratt and Zeckhauser 1990), divorce settlements (Brams and Taylor 1996), spectrum allocation (Etkin et al. 2005), air traffic management (Vossen 2002), course allocation (Budish and Cantillon 2010) and many more.*

In this paper, we focus on an important open problem in discrete fair division, where a set M of m indivisible goods needs to be allocated to a set $[n]$ of n agents. Each agent i is equipped with a valuation function $v_i \colon 2^M \to \mathbb{R}_{\geq 0}$ which captures the utility i derives from any bundle that can be allocated to her. One of the most well studied classes of valuations are *additive valuations*, i.e., $v_i(S) = \sum_{g \in S} v_i(\{g\})$ for all $S \subseteq M$. The goal is to determine a partition $X = \langle X_1, X_2, \ldots, X_n \rangle$ of M such that X_i is allocated to agent i, which is *fair*. Depending on the notion of fairness used, this setting has several different problems.

Envy-freeness up to any good (EFX) The quintessential notion of fairness is that of envy-freeness. An allocation $X = \langle X_1, X_2, \ldots, X_n \rangle$ is envy-free if every agent prefers her bundle as much as she

*See `www.spliddit.org` and `www.fairoutcomes.com` for a more detailed explanation of fair division protocols used in day-to-day problems.

prefers the bundle of any other agent, i.e., $v_i(X_i) \geq v_i(X_{i'})$ for all $i, i' \in [n]$. However, an envy-free allocation does not always exist, e.g., consider dividing a single valuable good among two agents. In any feasible allocation, the agent with no good will envy the agent that has been allocated the good. This necessitates the study of relaxed notions of envy-freeness. In this paper, we consider the relaxation known as *envy-freeness up to any good* (EFX). An allocation $X = \langle X_1, X_2, \dots, X_n \rangle$ is EFX if and only if for all pairs of agents i and i', we have $v_i(X_i) \geq v_i(X_{i'} \setminus \{g\})$ for all $g \in X_{i'}$, i.e., the envy should disappear following the removal of any single good from i''s bundle. EFX is, in fact, considered to be the "closest analogue of envy-freeness" in discrete fair division (Caragiannis et al. 2019). Unfortunately, the existence of EFX allocations is still unresolved despite significant efforts by several researchers (Moulin 2019, Caragiannis et al. 2016) and is considered one of the most important open problems in fair division (Procaccia 2020). There have been studies on

- the existence of EFX allocations in restricted settings. In particular, EFX existence has been studied when there are a small number of agents (Plaut and Roughgarden 2020, Chaudhury et al. 2020) and when agents have specific valuation functions (Halpern et al. 2020).
- The existence of relaxations of EFX allocations has also been investigated, e.g., approximate EFX allocations (Plaut and Roughgarden 2020, Amanatidis et al. 2020), EFX with bounded charity (Chaudhury et al. 2021b, Berger et al. 2022), approximate EFX with bounded charity (Chaudhury et al. 2021a).

Improving the understanding in both settings is a systematic direction toward the big problem. We first mention the existing results in the above two settings and some of their pitfalls. Thereafter, we highlight the main results of this paper and show how they address the said pitfalls. In particular, we focus on the existence of EFX allocations with a small number of agents and the existence of approximate EFX allocations with bounded charity.

Existence of EFX Allocations with Small Number of Agents. Plaut and Roughgarden (2020) first showed the existence of EFX allocations when there are two agents, using the *cut and choose protocol.* The existence of EFX allocations gets notoriously difficult with three or more agents. The existence of EFX allocations for three agents with *additive valuations* was shown by Chaudhury et al. (2020). Thereafter, Berger et al. (2022) showed the existence of EFX allocations with three agents when agents have *nice-cancelable valuation functions* – a class that subsumes additive, budget-additive, unit demand, and multiplicative valuation functions. However, this technique does not extend, as soon as one of the agents has a general monotone valuation function. Despite its fundamentality and ongoing efforts, the existence of EFX allocations with three agents under general valuation functions remains elusive. (Plaut and Roughgarden 2020) remarks, "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". In this paper, we make progress on this problem. As our first main result, we show the existence of EFX allocations, when *two* agents have general monotone valuation functions.

Theorem 1. *EFX allocations exist with three agents as long as at least one agent has an additive valuation function (the other two agents have general monotone valuation functions).*

In fact, our proof gives a stronger version of Theorem 1: we can show the existence of EFX allocations when two agents have general monotone valuation functions and one of the agents has an *MMS-feasible valuation function* – a valuation class that strictly generalizes nice-cancelable valuation functions– definitions and properties are described in Section 2. Thus, we strictly generalize the result in (Berger et al. 2022).

Theorem 2. *EFX allocations exist with three agents as long as at least one agent has an MMS-feasible valuation function.*

We briefly remark on our technique to prove Theorem 2 and how it crucially differs from the existing techniques in (Chaudhury et al. 2020, Berger et al. 2022, Mahara 2021). The algorithms in (Chaudhury et al. 2020, Berger et al. 2022, Mahara 2021) move in the space of partial EFX allocations (where not all goods are allocated) iteratively improving the vector $\langle v_1(X_1), v_2(X_2), v_3(X_3)\rangle$ lexicographically, where $v_i(\cdot)$ is the valuation function of agent i. However, (Chaudhury et al. 2021a) exhibit an instance with four agents and a partial EFX allocation X, such that in any complete EFX allocation X', $v_1(X'_1) < v_1(X_1)$, i.e., agent 1 (which is the highest priority agent) is better off in X than in any complete EFX allocation. This necessitates the study of a different approach for the existence of EFX allocations. *Our algorithm moves in the space of complete allocations (instead of partial allocations), iteratively improving a certain potential as long as the current allocation is not EFX. Furthermore, this proof turns out to be simpler and significantly shorter than the ones in (Chaudhury et al. 2020, Berger et al. 2022), as it does not use the notions of champions, champion-graphs, half-bundles, and even the envy-graph.*

Existence of Approximate EFX with Bounded Charity. Caragiannis et al. (2019) introduced the notion of EFX with charity. The goal here is to find "good" partial EFX allocations, i.e., partial EFX allocations where the set of unallocated goods is not very valuable. In particular, they show that there always exists a partial EFX allocation X such that for each agent i, we have $v_i(X_i) \geq 1/2 \cdot v_i(X^*_i)$, where $X^* = \langle X^*_1, X^*_2, \ldots, X^*_n\rangle$ is the allocation with maximum *Nash welfare*. The Nash welfare of an allocation Y is the geometric mean of agents' valuations, $\left(\prod_{i\in[n]} v_i(Y_i)\right)^{1/n}$. It is often considered a direct measure of the fairness and efficiency of an allocation. Following the same line of work, Chaudhury et al. (2021b) showed the existence of a partial EFX allocation X such that no agent envies the set of unallocated goods and the total number of unallocated goods is at most $n-1$. Quite recently, Chaudhury et al. (2021a) showed the existence of a $(1-\varepsilon)$-EFX allocation with $\mathcal{O}((n/\varepsilon)^{4/5})$ charity for any $\varepsilon > 0$, where an allocation X is said to be $(1-\varepsilon)$-EFX if and only if $v_i(X_i) \geq (1-\varepsilon) \cdot v_i(X_{i'} \setminus \{g\})$, $\forall i, i'$ and $\forall g \in X_{i'}$. While the last result is not a strict improvement of the result in (Chaudhury et al. 2021b) (since it ensures $(1-\varepsilon)$-EFX instead of exact EFX), it is the best relaxation of EFX that we can compute in polynomial time, as the algorithm in (Chaudhury et al. 2021b) can only be modified to give $(1-\varepsilon)$-EFX with $n-1$ charity in polynomial time. Another key aspect of the technique in (Chaudhury et al. 2021a) is the reduction of the problem of improving the bounds on charity to a purely graph-theoretic problem. In particular, (Chaudhury et al. 2021a) defines the notion of a *rainbow cycle number*: Given an integer $d > 0$, the rainbow cycle number $R(d)$ is the largest k such that there exists a k-partite graph $G = (V_1 \cup V_2 \cup \cdots \cup V_k, E)$ such that

- each part has at most d vertices, i.e., $|V_i| \leq d$, and
- every vertex in G has exactly one incoming edge from every part in G except the part containing it, and
- there exists no cycle C in G that visits each part at most once.

Let $h^{-1}(d)$ denote the smallest integer ℓ such that $h(\ell) = \ell \cdot R(\ell) \geq d$. Then there always exist an $(1-\varepsilon)$-EFX allocation with $\mathcal{O}(\frac{n}{\varepsilon \cdot h^{-1}(n/\varepsilon)})$. So, the smaller the upper bound on $h(\ell)$, the lower the number of unallocated goods. (Chaudhury et al. 2021a) shows that $R(d) \in \mathcal{O}(d^4)$ and thus establish the existence of $(1-\varepsilon)$-EFX allocation with $\mathcal{O}((n/\varepsilon)^{4/5})$ charity. An upper bound of $\mathcal{O}(d^2 2^{(\log\log d)^2})$ was obtained by Berendsohn et al. (2022), thereby showing the existence of EFX allocations with $\mathcal{O}((n/\varepsilon)^{0.67})$ charity. In this paper, we close this line of improvements by proving an almost tight upper bound on d (matching the lower bound up to a log factor). We note that our technique (and analysis) is simpler than the ones used for upper-bounding rainbow cycle number in (Chaudhury et al. 2021a, Berendsohn et al. 2022, Jahan et al. 2023). Although the core argument is based on the *probabilistic method*, we also derandomize the approach to construct such an allocation in deterministic polynomial time.

Theorem 3. *Given any integer $d > 0$, the rainbow cycle number $R(d) \in \mathcal{O}(d \log d)$.*

For any allocation X, let us denote the Nash welfare of X by $NW(X)$. As a consequence of the improved bound in Theorem 3, we obtain:

Theorem 4. *There exists a deterministic polynomial time algorithm that determines a partial $(1-\varepsilon)$-EFX allocation X such that no agent envies the set of unallocated goods and the total number of unallocated goods is $\tilde{\mathcal{O}}((n/\varepsilon)^{1/2})^{\dagger}$. Furthermore, $NW(X) \geq 1/(2e^{1/e}) \cdot NW(X^*)$ where X^* is the allocation with maximum Nash welfare.*

$^{\dagger}\tilde{\mathcal{O}}$ ignores logarithmic factors.

Rainbow Cycle and Zero-sum Combinatorics. We believe that investigating tighter bounds on $R(d)$ is interesting in its own right. Recently, Berendsohn et al. (2022) showed intriguing connections between rainbow cycle number and zero-sum problems in extremal combinatorics. Zero-sum problems in graphs ask questions of the following flavor: Given an edge/vertex weighted graph, whether there exists a certain substructure (for example, cliques, cycles, paths, etc.) with a zero-sum (modulo some integer). In particular, (Berendsohn et al. 2022) shows that the rainbow cycle number is a natural generalization of the zero-sum problems studied by Alon and Krivelevich (2021) and Mészáros and Steiner (2021). Both papers (Alon and Krivelevich 2021, Mészáros and Steiner 2021) aim to upper bound the maximum number of vertices of a complete bidirected graph with integer edge labels avoiding a zero-sum cycle (modulo d). (Berendsohn et al. 2022) shows through a simple argument that this is upper bounded by the *permutation rainbow cycle number* $R_p(d)$, which is defined by introducing an additional constraint in the definition of $R(d)$ that for all i, j, each vertex in V_i has exactly one *outgoing* edge to some vertex in V_j (in addition to exactly one incoming edge from some vertex in V_j). In Section 5.2, we show through a simple argument that $R_p(d) \leq 2d-2$, thereby also improving the upper bounds of $\mathcal{O}(d\log(d))$ in (Alon and Krivelevich 2021) and $8d-1$ in (Mészáros and Steiner 2021).

LEMMA 1. *For $d > 1$, we have $R_p(d) \leq 2d-2$. Therefore, by the Observation made by Berendsohn et al. (2022), the maximum number of vertices of a complete bidirected graph with integer edge labels avoiding a zero-sum cycle (modulo d) is at most $2d-2$.*

1.1. Further Related Work

Fair division has received significant attention since the seminal work of Steinhaus (1948). Other than envy-freeness, another fundamental fairness notion is that of *proportionality*. Recall that, in an envy-free allocation 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 she values $1/n$ times her valuation on the entire set of goods. Since envy-freeness and proportionality cannot always be guaranteed while dividing indivisible goods, various relaxations 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. While the existence of EFX allocations is open, EF1 allocations are known to exist for any number of agents, even when agents have general monotone valuation functions (Lipton et al. 2004). Also, another relaxation of envy-freeness called EF2X was studied in (Akrami et al. 2022), which shows the existence of EF2X allocations under restricted additive valuations.

While EF1 and EFX are fairness notions that relax envy-freeness, the most popular notions of fairness that relax proportionality for indivisible goods are *maximin share* (MMS), proportionality up to one good (PROP1), proportionality up to any good (PROPx), and proportionality up to the maximin good (PROPm). The MMS was introduced by Budish (2011). While MMS allocations do not always exist (Kurokawa et al. 2018), there has been extensive work to come up with approximate MMS allocations (Budish 2011, Bouveret and Lemaître 2016, Amanatidis et al. 2017, Barman and Krishnamurthy 2017, Kurokawa et al. 2018, Ghodsi et al. 2018, Garg et al. 2019, Garg and Taki 2020, Akrami et al. 2023). On the other hand, PROPx is stronger than PROPm, which is stronger than PROP1. While PROPx allocations do not always exist (Moulin 2019), PROPm allocations are guaranteed to exist (Baklanov et al. 2021). Some works assume ordinal ranking over the goods, as opposed to cardinal values, e.g., (Aziz et al. 2015, Brams et al. 2017).

Alongside fairness, the efficiency of an allocation is also a desirable property. Two standard measures of efficiency are Pareto-optimality and Nash welfare. Caragiannis et al. (2016) showed that any allocation that has the maximum Nash welfare is guaranteed to be Pareto-optimal (efficient) and EF1 (fair). Barman et al. (2018) gave a pseudo-polynomial algorithm to find an allocation that is both EF1 and Pareto-optimal, which was recently improved by Murhekar and Garg (2021). Other works explore relaxations of EFX with high Nash welfare (Caragiannis et al. 2019, Chaudhury et al. 2021b).

Independent Work. Independently and concurrently to our work, Berendsohn et al. (2022) also investigated upper bounds on rainbow cycle number. They obtained the same upper bound of $2d-2$ for $R_p(d)$. Jahan et al. (2023) also independently investigated upper bounds on the rainbow cycle number, and they showed $R(d) \in \mathcal{O}(d\log(d))$.

2. Preliminaries

For any non-negative integer n, let $[n] = \{1, 2, \ldots, n\}$. An instance of discrete fair division is given by the tuple $\langle [n], M, \mathcal{V} \rangle$, where $[n]$ is the set of agents, M is the set of indivisible goods, and $\mathcal{V} = (v_1(\cdot), v_2(\cdot), \ldots, v_n(\cdot))$ where each $v_i : 2^M \to \mathbb{R}_{\geq 0}$ denotes the valuation of agent i. Typically, the valuation functions are assumed to be *monotone*, i.e., for each agent i, $v_i(S \cup \{g\}) \geq v_i(S)$ for all $S \subseteq M$ and $g \notin S$, and *normalized*, i.e., for each agent i, $v_i(\emptyset) = 0$. A valuation $v_i(\cdot)$ is said to be *additive* if $v_i(S) = \sum_{g \in S} v_i(\{g\})$ for all $S \subseteq M$. For ease of notation, we use g instead of $\{g\}$. We also use $S \oplus_i T$ for $v_i(S) \oplus v_i(T)$ with $\oplus \in \{\leq, \geq, <, >\}$.

Given an allocation $X = \langle X_1, X_2, \ldots, X_n \rangle$, we say that an agent i *strongly envies* an agent i' if and only if $X_i <_i X_{i'} \setminus \{g\}$ for some $g \in X_{i'}$. Thus, an allocation is an EFX allocation if no strong envy exists between any pair of agents. We now introduce certain definitions and recall certain concepts that will be useful in the upcoming sections.

Definition 1 (EFX feasibility). Given a partition $X = (X_1, X_2, \ldots, X_n)$ of M, a bundle X_k is EFX-feasible to agent i if and only if $X_k \geq_i max_{j \in [n]} max_{g \in X_j} X_j \setminus g$. Therefore, an allocation $X = \langle X_1, X_2, \ldots, X_n \rangle$ is EFX if for each agent i, X_i is EFX-feasible .

Chaudhury et al. (2020) introduced the notion of non-degenerate instances where no agent values two distinct bundles the same. They showed that to prove the existence of EFX allocations in the additive setting, it suffices to show the existence of EFX allocations for all non-degenerate instances. We adapt their approach and show that the same claim holds, even when agents have general monotone valuations.

Non-Degenerate Instances (Chaudhury et al. 2020) We call an instance $\mathcal{I} = \langle [n], M, \mathcal{V}\rangle$ non-degenerate if and only if no agent values two different sets equally, i.e., $\forall i \in [n]$ we have $v_i(S) \neq v_i(T)$ for all $S \neq T$. We extend the technique in (Chaudhury et al. 2020) and show that it suffices to deal with non-degenerate instances when there are n agents with general valuation functions, i.e., if there exists an EFX allocation in all non-degenerate instances, then there exists an EFX allocation in all instances. We defer the reader to the appendix for detailed proof.

Henceforth, we assume that the given instance is non-degenerate, implying that all goods are positively valued by all agents.

MMS-feasible Valuations. In this paper, we introduce a new class of valuation functions called MMS-feasible valuations which are natural extensions of additive valuations in a fair division setting.

DEFINITION 2. A valuation function $v: 2^M \to \mathbb{R}_{\geq 0}$ is MMS-feasible if for every subset of goods $S \subseteq M$ and every partitions $A = (A_1, A_2)$ and $B = (B_1, B_2)$ of S, we have

$$\max(v(B_1), v(B_2)) \geq \min(v(A_1), v(A_2)).$$

Informally, these are the valuations under which an agent always has a bundle in any 2-partition of any subset of the goods that she values at least as much as her MMS value, i.e., given an agent i with an MMS-feasible valuation $v(\cdot)$, in any 2-partition of $S \subseteq M$, say $B = (B_1, B_2)$, we have $max(v(B_1), v(B_2)) \geq MMS_i^2(S)$, where $MMS_i^2(S)$ is the MMS value of agent i on the set S when there are 2 agents. Also, note that if there are two agents and one of the agents has an MMS-feasible valuation function, then irrespective of the valuation function of the other agent, MMS allocations always exist: Consider an instance where agent 1 has an MMS-feasible valuation function and agent 2 has a general monotone valuation function. Consider agent 2's optimal MMS partition of the good set $A = (A_1, A_2)$. Let agent 1 pick her favorite bundle from A. Then, agent 1 has a bundle that she values at least as much as her MMS value (as she has an MMS-feasible valuation function), and agent 2 has a bundle that she values at least as much as her MMS value as A is an MMS optimal partition according to agent 2.

MMS-feasible valuations strictly generalize the *nice-cancelable valuation functions* introduced in (Berger et al. 2022). A valuation function $v: 2^M \to \mathbb{R}_{\geq 0}$ is nice-cancelable if for every $S, T \subset M$

S	$\{g_1\}$	$\{g_2\}$	$\{g_3\}$	$\{g_1, g_2\}$	$\{g_1, g_3\}$	$\{g_2, g_3\}$	$\{g_1, g_2, g_3\}$
v	1	2	3	10	4	5	13

Table 1 valuation function v is MMS-feasible but not nice-cancelable.

and $g \in M \setminus (S \cup T)$, we have $v(S \cup \{g\}) > v(T \cup \{g\}) \Rightarrow v(S) > v(T)$. Nice-cancelable valuations include *budget-additive* ($v(S) = min(\sum_{s \in S} v(s), c)$), *unit demand* ($v(S) = max_{j \in S} v(s)$), and *multiplicative* ($v(S) = \prod_{s \in S} v(s)$) valuations (Berger et al. 2022).

LEMMA 2. *Every nice-cancelable function is MMS-feasible.*

We first make an observation about a nice-cancelable valuation function.

OBSERVATION 5. If v is a nice-cancelable valuation, then for every $S, T \subset M$ and $Z \subseteq M \setminus (S \cup T)$, we have $v(S \cup Z) > v(T \cup Z) \Rightarrow v(S) > v(T)$.

Let v be a nice-cancelable function. For a subset of goods $S \subseteq M$, consider any two partitions $A = (A_1, A_2)$ and $B = (B_1, B_2)$ of S. Without loss of generality assume $v(A_1 \cap B_1) < v(A_2 \cap B_2)$. Since $(A_1 \cap B_2)$ is disjoint from $(A_1 \cap B_1) \cup (A_2 \cap B_2)$, by the contrapositive of Observation 5 applied to nice-cancelable valuation v, we have,

$$v((A_1 \cap B_1) \cup (A_1 \cap B_2)) < v((A_2 \cap B_2) \cup (A_1 \cap B_2)). \tag{1}$$

Therefore,

$$\begin{aligned}
\min(v(A_1), v(A_2)) &\leq v(A_1) \\
&= v((A_1 \cap B_1) \cup (A_1 \cap B_2)) && A_1 = (A_1 \cap B_1) \cup (A_1 \cap B_2) \\
&< v((A_2 \cap B_2) \cup (A_1 \cap B_2)) && \text{Inequality (1)} \\
&= v(B_2) && B_2 = (A_2 \cap B_2) \cup (A_1 \cap B_2) \\
&\leq \max(v(B_1), v(B_2)).
\end{aligned}$$

To prove that MMS-feasible functions strictly generalize nice-cancelable functions, we present an example of a valuation function that is MMS-feasible but not nice-cancelable.

EXAMPLE 1. Let $M = \{g_1, g_2, g_3\}$. The value of $v(S)$ is given in Table 1 for all $S \subseteq M$. First note that $v(g_1 \cup g_2) > v(g_3 \cup g_2)$ but $v(g_1) < v(g_3)$. Therefore, v is not nice-cancelable. Now we prove that v is MMS-feasible. Let $S \subseteq M$ and $A = (A_1, A_2)$, $B = (B_1, B_2)$ be two partitions of S. Without loss of generality, assume $|A_1| \leq |A_2|$. If $A_1 = \emptyset$, $\min(v(A_1), v(A_2)) = 0 \leq \max(v(B_1), v(B_2))$. Hence, we assume $|A_1| \geq 1$ and therefore, we have $|S| \geq 2$. Moreover, if $A = B$, then $\max(v(B_1), v(B_2)) = \max(v(A_1), v(A_2)) \geq \min(v(A_1), v(A_2))$. Thus, we also assume $A \neq B$. If $S = \{g, g'\}$, the only two different possible partitioning of S is $A = (\{g\}, \{g'\})$ and $B = (\emptyset, \{g, g'\})$. For all $g, g' \in M$,

$v(\{g, g'\}) > \max(v(g), v(g'))$. Therefore, $\max(v(B_1), v(B_2)) \geq \min(v(A_1), v(A_2))$. If $S = \{g_1, g_2, g_3\}$, then $|A_1| = 1$ and therefore, $\min(v(A_1), v(A_2)) \leq v(A_1) \leq \max_{g\in M}(v(g)) = 3$. Without loss of generality, let $g_3 \in B_1$. For all $T \subseteq M$ such that $g_3 \in T$, we have $v(T) \geq 3$. Thus, $\max(v(B_1), v(B_2)) \geq v(B_1) \geq 3 \geq \min(v(A_1), v(A_2))$.

The next lemma follows from Lemma 2 and Example 1.

LEMMA 3. *The class of MMS-feasible valuation functions is a strict superclass of nice-cancelable valuation functions.*

2.1. Rainbow Cycle Number

Chaudhury et al. (2021a) reduced the problem of finding approximate EFX allocations with sublinear charity to a problem in extremal graph theory. In particular, they introduced the notion of a rainbow cycle number.

DEFINITION 3. Given an integer $d > 0$, the rainbow cycle number $R(d)$ is the largest k such that there exists a k-partite directed graph $G = (V_1 \cup V_2 \cup \cdots \cup V_k, E)$ such that

- each part has at most d vertices, i.e., $|V_i| \leq d$, and
- every vertex has exactly one incoming edge from every part other than the one containing it[‡], and
- there exists no cycle C in G that visits each part at most once.

We also refer to cycles that visit each part at most once as "rainbow" cycles.

They show that any finite upper bound on $R(d)$ implies the existence of approximate EFX allocations with sublinear charity. Better upper bounds on $R(d)$ gives better bounds on the charity. In particular, they prove the following theorem.

THEOREM 6. *(Chaudhury et al. 2021a) Let $G = (V_1 \cup V_2 \cup \ldots V_k, E)$ be a k-partite digraph such that (i) each part has at most d vertices and (ii) each vertex in G has an incoming edge from every part other than the one containing it. Furthermore, let $k > T(d) \geq R(d)$. If there exists a polynomial time algorithm that can find a cycle visiting each part at most once in G, then there exists a polynomial time algorithm that determines a partial EFX allocation X such that*

- *the total number of unallocated goods is in $\mathcal{O}(n/\varepsilon \cdot h^{-1}(n/\varepsilon))$ where $h^{-1}(d)$ is the smallest integer ℓ such that $h(\ell) = \ell \cdot T(\ell) \geq d$.*
- *$NW(X) \geq 1/(2e^{1/e}) \cdot NW(X^*)$, where X^* is the allocation with maximum Nash welfare.*

[‡] In the original definition of the rainbow cycle number $R(d)$ in (Chaudhury et al. 2021a), every vertex can have more than one incoming edge from a part. However, by reducing the number of edges, we can only increase the upper bound on $R(d)$.

3. Technical Overview

In this section, we briefly highlight the main technical ideas used to show our results.

3.1. EFX Existence beyond Additivity

We present an algorithmic proof for the existence of EFX allocations when agents have valuations more general than additive valuations. The main takeaway of our algorithm is that it does not require the sophisticated concepts introduced by the techniques in (Chaudhury et al. 2021b, 2020) that rely on improving a potential function while moving in the space of partial EFX allocations. In fact, our algorithm does not even require the concept of envy-graph which is a very fundamental concept used by the algorithms in (Chaudhury et al. 2021b, 2020) and also in (Plaut and Roughgarden 2020, Lipton et al. 2004) to prove the existence of weaker relaxations of envy-freeness (like EF1 and 1/2-EFX).

The crucial idea in our technique is to move in the space of partitions (of the good set), improving a certain potential as long as we cannot find an EFX allocation from the current partition, i.e., we cannot find a *complete* allocation of the bundles in the partition such that the EFX property is satisfied. In particular, we always maintain a partition $X = (X_1, X_2, X_3)$ such that (i) agent 1 finds X_1 and X_2 EFX-feasible and (ii) at least one of agent 2 and agent 3 finds X_3 EFX-feasible. Note that such allocations always exist: Agent 1 can determine a partition such that all bundles are EFX-feasible for her (such a partition is possible as agent 1 can run the algorithm in (Plaut and Roughgarden 2020) to find an EFX allocation assuming all three agents have agent 1's valuation function, thereby making all bundles EFX-feasible for her) and we call agent 2's favorite bundle in the partition X_3 (which is obviously EFX-feasible for her) and the remaining bundles X_1 and X_2 arbitrarily. Then, we have a partition that satisfies the invariant.

Note that if any one of agent 2 or 3 finds one of X_1 or X_2 EFX-feasible, then we easily get an EFX allocation. Indeed, assume without loss of generality that agent 2 finds X_3 EFX-feasible. Now, if

- agent 3 finds X_2 EFX-feasible, then we have an EFX allocation: agent 1 $\leftarrow X_1$, agent 2 $\leftarrow X_3$, and agent 3 $\leftarrow X_2$. We can give a symmetric argument when agent 3 finds X_1 EFX-feasible.
- Similarly, if agent 2 finds X_2 EFX-feasible, then we can let agent 3 pick her favorite bundle in the partition (which is obviously EFX-feasible for her) and still give agents 1 and 2 an EFX-feasible bundle. We can give a symmetric argument when agent 2 finds X_1 EFX-feasible.

Therefore, we only need to consider the scenario where only X_3 is EFX-feasible for agents 2 and 3. Essentially, in this scenario, X_3 is "too valuable" to agents 2 and 3, as they do not find any of the remaining bundles EFX-feasible. *A natural attempt would be to remove some good(s) from*

X_3 and allocate it to X_1 or X_2, i.e., we can increase the valuation of agent 1 for her EFX-feasible bundle(s) by removing the excess goods allocated to the only EFX-feasible bundle of agents 2 and 3. This brings us to our potential function: $\phi(X) = min(v_1(X_1), v_1(X_2))$. Now, the non-triviality lies in determining the set of goods to be removed from X_3 and then allocating them to X_1 and X_2 such that we maintain our invariants. Although non-trivial, this turns out to be significantly simpler than the procedure used in (Chaudhury et al. 2020) and also holds when agents 1 and 2 have general monotone valuation functions and agent 3 has an MMS-feasible valuation function. The entire procedure is elaborated in Section 4.

3.2. Improved Bounds on Rainbow Cycle Number

Our technique to achieve the improved bound involves the probabilistic method. It is significantly simpler and yields better guarantees. We briefly sketch our algorithmic proof. Let there be k parts in $G = (V_1 \cup V_2 \cup \dots V_k, E)$. Note that each part has at most d vertices and each vertex has at least one incoming edge from every part. We pick one vertex v_i from each part V_i uniformly and independently at random. Now, it suffices to show that with non-zero probability, the induced graph on the vertices $v_1, v_2, \dots, v_k$ is cyclic for some $k \in \mathcal{O}(d \log d)$. Note that if every vertex in $G[v_1, \dots, v_k]$ has an incoming edge, then $G[v_1 \dots v_k]$ is cyclic. So we need to show a non-zero lower bound on the probability of each vertex having at least one incoming edge or equivalently show an upper bound on the probability that each vertex has no incoming edge in $G[v_1 \dots v_k]$. To this end, let E_{v_i} denote the event that vertex v_i has no incoming edge in $G[v_1 \dots v_k]$. Note that $\mathbf{P}[E_{v_i}] \le (1 - 1/d)^{k-1}$: v_i has at least one incoming edge from each part, and therefore, the probability that there is no incoming edge from v_j to v_i is at most $(1-1/d)$ for all j. Since all v_j's are independently chosen, the probability that v_i has no incoming edge from any part is at most $(1 - 1/d)^{(k-1)}$. Then, by union bound, $\mathbf{P}[\cup_{i\in[n]} E_{v_i}] \le \sum_{i\in[n]} \mathbf{P}[E_{v_i}] \le k(1-1/d)^{(k-1)}$. Therefore, the probability that $G[v_1 \dots v_k]$ is cyclic is at least $1 - k(1-1/d)^{(k-1)}$ which is strictly positive for $k \in \mathcal{O}(d \log d)$.

4. EFX Existence beyond Additivity

Before we give the new algorithm, we first give the reader a quick recap of the Plaut and Roughgarden (PR) algorithm (2020) that determines an EFX allocation when all agents have the same valuation function, $v(\cdot)$ (the only assumption on $v(\cdot)$ is that it is monotone). The algorithm starts with any arbitrary allocation X (which may not be EFX) and makes minor reallocations to improve the valuation of the agent who has the lowest value, i.e., it modifies X to X' such that $min_{i\in[n]} v(X'_i) > min_{i\in[n]} v(X_i)$. We now elaborate on the reallocation procedure: Let ℓ be the agent with the lowest valuation in X. If X is not EFX, then there exists agents i and j such that $v(X_i) < v(X_j \setminus \{g\})$ for some $g \in X_j$. Since $v(X_\ell) < v(X_i)$, we also have $v(X_\ell) < v(X_j \setminus \{g\})$. The algorithm removes

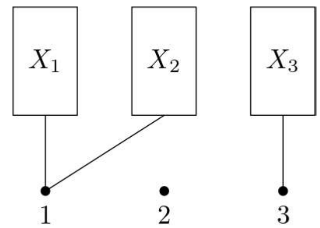

Figure 1 The nodes correspond to agents and an edge from agent i to a bundle X_j means that X_j is EFX-feasible for i. In this example, X_1 and X_2 are EFX-feasible for agent 1, and X_3 is EFX-feasible for agent 3. Therefore, the invariants hold.

the good g from j's bundle and allocates it to ℓ. Observe that $v(X_k) > v(X_\ell)$ for all $k \neq \ell$ as we assume non-degeneracy. Also, we have $v(X_\ell \cup \{g\})$ and $v(X_j \setminus \{g\})$ greater than $v(X_\ell)$. Therefore, the valuation of every new bundle is strictly larger than the valuation of X_ℓ. Thus, the valuation of the agent with the lowest valuation improves. This implies that the reallocation procedure will never revisit a particular allocation, and as a result, this process will eventually converge to an EFX allocation Y with $v(Y_i) > v(X_\ell)$ for all $i \in [n]$. Formally,

Lemma 4 **(Plaut and Roughgarden (2020))**. *Let $X = (X_1, X_2, X_3)$ be an arbitrary 3-partition. Running the PR algorithm with any monotone valuation v results in an EFX-partition $X' = (X'_1, X'_2, X'_3)$ with*

$$\min(v(X_1), v(X_2), v(X_3)) \leq \min(v(X'_1), v(X'_2), v(X'_3)).$$

We have equality only if the input is already EFX for v.

In contrast to the algorithms in (Chaudhury et al. 2020, 2021b, Berger et al. 2022, Plaut and Roughgarden 2020), our algorithm moves in the space of complete EFX allocations, iteratively maintaining some invariants. As long as our allocation is not EFX, we make some reallocations to the existing allocation and improve a certain potential. We give the proof here assuming only monotonicity for the valuation functions of agents 1 and 2 and assuming MMS-feasibility for the valuation of agent 3, i.e., $v_1(\cdot)$ and $v_2(\cdot)$ are general monotone valuation functions and $v_3(\cdot)$ is MMS-feasible. We now elaborate on our algorithm. We maintain a partition (X_1, X_2, X_3) of the good set such that

- X_1 and X_2 are EFX-feasible for agent 1.
- X_3 is EFX-feasible for at least one of agents 2 and 3.

See Figure 1 for better intuition.

One can show the existence of allocations satisfying the above invariants by running the PR algorithm and initializing: Agent 1 runs the PR algorithm with $v = v_1$ to determine a partition (X_1, X_2, X_3) such that all the three bundles are EFX-feasible for her. Then, agent 3 picks her favorite bundle out of the three, say X_3. Clearly, X_3 is EFX-feasible for agent 3, and X_1 and X_2 are EFX-feasible for agent 1. Thus X satisfies the invariants.

We define our potential function as $\phi(X) = \min(v_1(X_1), v_1(X_2))$. We now elaborate on how to modify X and improve the potential when we cannot determine an EFX allocation from the partition X, i.e., we cannot determine an allocation of the bundles in X to the agents that satisfy the EFX property.

4.1. Reallocation When We Cannot Get an EFX Allocation from X

Let $X = (X_1, X_2, X_3)$ be a partition satisfying the invariants. Without loss of generality, let us assume that agent 2 finds X_3 EFX-feasible. Observe that if any one of agents 2 or 3 finds bundles X_1 or X_2 EFX-feasible, then we are done: If agent 3 finds one of X_1 or X_2 EFX-feasible, then we can allocate agent 3's EFX-feasible bundle to her, X_3 to agent 2 and the remaining bundle of X_1 and X_2 to agent 1 and get an EFX allocation. Similarly, if agent 2 finds X_1 or X_2 EFX-feasible, we ask agent 3 to pick her favorite bundle out of X_1, X_2, and X_3. Now, note that no matter which bundle agent 3 picks, there is always a way to allocate agents 1 and 2 their EFX-feasible bundles as agent 1 finds X_1 and X_2 EFX-feasible and agent 2 finds X_3 and at least one of X_1 or X_2 EFX-feasible. If agent 3 picks X_1, allocate X_2 to agent 1 and X_3 to agent 2. If agent 3 picks X_2, allocate X_1 to agent 1 and X_3 to agent 2. Finally, if she picks X_3, allocate the bundle among X_1 and X_2, which is EFX-feasible for agent 2 to agent 2 and the remaining bundle to agent 1. Therefore, from here on we assume that neither agent 2 nor agent 3 finds X_1 or X_2 EFX-feasible. Let g_i be the good in X_3 such that $X_3 \setminus g_i \geq_i X_3 \setminus h$ for all $h \in X_3$, i.e., $X_3 \setminus g_i$ is the most valued proper subset of X_3 for agent i.

Observation 7. For $i \in \{2, 3\}$, we have $X_3 \setminus g_i >_i max_i(X_1, X_2)$.

We prove for $i = 2$. The proof for $i = 3$ is identical. Let us assume otherwise and say without loss of generality $X_1 >_2 X_3 \setminus g_2$. Then, the only reason why X_1 is not EFX-feasible for agent 2 is if $X_1 <_2 X_2 \setminus g$ for some $g \in X_2$. But, in that case, we have $X_2 >_2 X_1 >_2 X_3 \setminus g_2$. Therefore, we have $X_2 >_2 max_{\ell \in [3]} max_{h \in X_\ell} X_\ell \setminus h$, implying that X_2 is EFX-feasible, which is a contradiction.

With loss of generality assume that $X_1 <_1 X_2$, implying that $\phi(X) = v_1(X_1)$. We now distinguish two cases depending on how valuable the bundle $X_1 \cup g_i$ is to agent i for $i \in \{2, 3\}$ and give the appropriate reallocations in both cases. In particular, we first look into the case where $X_3 \setminus g_i$ is still more valuable to agent i than $X_1 \cup g_i$ for at least one $i \in \{2, 3\}$.

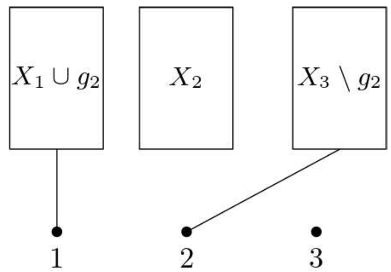

Figure 2 Assuming that $X_3 \setminus g_2 >_2 X_1 \cup g_2$, the edge between agent 2 and $X_3 \setminus g_2$ exists. Also, the edge between agent 1 and $X_1 \cup g_2$ exists.

Case: $\boldsymbol{X_3 \setminus g_2 >_2 X_1 \cup g_2}$ *or* $\boldsymbol{X_3 \setminus g_3 >_3 X_1 \cup g_3}$*, i.e.,* $\boldsymbol{X_3 \setminus g_i}$ *is the favorite bundle for agent* $\boldsymbol{i}$ *in the partition* $\boldsymbol{X_1 \cup g_i}$*,* $\boldsymbol{X_2}$ *and* $\boldsymbol{X_3 \setminus g_i}$ *for at least one* $\boldsymbol{i \in \{2,3\}}$. We provide the reallocation rules assuming that $X_3 \setminus g_2 >_2 X_1 \cup g_2$. The rules for the case $X_3 \setminus g_3 >_3 X_1 \cup g_3$ is symmetric. Now, consider the partition $(X_1 \cup g_2, X_2, X_3 \setminus g_2)$. See Figure 3.

By Observation 7, $X_3 \setminus g_2 >_2 X_2$ and by our current case $X_3 \setminus g_2 >_2 X_1 \cup g_2$, implying that $X_3 \setminus g_2$ is an EFX-feasible bundle for agent 2. Let X_1' be a minimal subset of $X_1 \cup g_2$ with respect to set inclusion that agent 1 values more than X_1, i.e., agent 1 values X_1 more than any proper subset of X_1' and $X_1' >_1 X_1$. Let $X_2' = X_2$ and $X_3' = (X_3 \setminus g_2) \cup ((X_1 \cup g_2) \setminus X_1')$. We define the partition $X' = (X_1', X_2', X_3')$. Observe that $\phi(X') > \phi(X)$ as $X_2' = X_2 >_1 X_1$ (by assumption) and $X_1' >_1 X_1$ (by definition). Also, note that X_3' is EFX-feasible for agent 2 as it is the most valuable bundle in X' for agent 2. Now, if X_1' and X_2' are EFX-feasible for agent 1, all invariants are maintained, and we are done. So now we look into the case when at least one of X_1' and X_2' is not EFX-feasible for agent 1 in X'.

We first make an observation on agent 1's valuation on the bundles X_1' and X_2'.

Observation 8. We have $X_1' >_1 X_2' \setminus g$ for all $g \in X_2'$ and $X_2' >_1 X_1' \setminus h$ for all $h \in X_1'$.

Note that $X_1' >_1 X_1$ by definition of X_1' and $X_1 >_1 X_2 \setminus g$ for all $g \in X_2$ as X_1 was EFX-feasible for agent 1 in X. Since $X_2' = X_2$, we have $X_1' >_1 X_2' \setminus g$ for all $g \in X_2'$.

Similarly, $X_2 >_1 X_1$ by assumption. Furthermore $X_1 >_1 X_1' \setminus h$ for all $h \in X_1'$ by the definition of X_1'. Since $X_2' = X_2$, we have $X_2' >_1 X_1' \setminus h$ for all $h \in X_1'$.

By Observation 8, if X_1' and X_2' are not EFX-feasible for agent 1 in X', then $X_3' \setminus g >_1 min_1(X_1', X_2')$ for some $g \in X_3'$. However, in that case, we run the PR algorithm on the partition X' with agent 1's valuation. Let $Y = (Y_1, Y_2, Y_3)$ be the final partition at the end of the PR algorithm. We have,

$$
\begin{aligned}
min(v_1(Y_1), v_1(Y_2), v_1(Y_3)) &> min(v_1(X_1'), v_1(X_2'), v_1(X_3')) && \text{(by Lemma 4)} \\
&= min(v_1(X_1'), v_1(X_2')) && \text{(as } v_1(X_3') > min(v_1(X_1'), v_1(X_2')))
\end{aligned}
$$

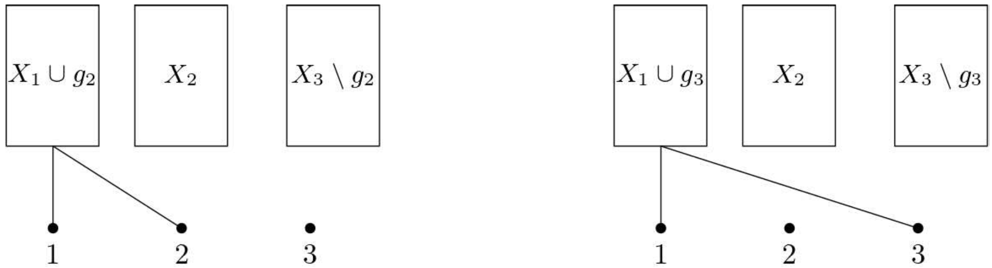

Figure 3 Assuming that $X_3 \setminus g_2 <_2 X_1 \cup g_2$ and $X_3 \setminus g_3 <_3 X_1 \cup g_3$, the edge between agent i and $X_1 \cup g_i$ exists in addition to the edge between agent 1 and $X_1 \cup g_i$.

$$
\begin{aligned}
&= \phi(X') \\
&> \phi(X).
\end{aligned}
$$

We then let agent 2 pick her favorite bundle out of Y_1, Y_2, and Y_3. Let us assume without loss of generality that she chooses Y_3. Then, allocation Y satisfies the invariants and we have $\phi(Y) = min(v_1(Y_1), v_1(Y_2)) \geq min(v_1(Y_1), v_1(Y_2), v_1(Y_3)) > \phi(X)$. Thus, we are done.

Remark: Note that we have not used the MMS-feasibility of $v_3(\cdot)$ yet. All the arguments, in this case, hold when all three valuation functions are general monotone. We use MMS-feasibility crucially in the upcoming case.

Case: $\boldsymbol{X_3 \setminus g_2 <_2 X_1 \cup g_2}$ *and* $\boldsymbol{X_3 \setminus g_3 <_3 X_1 \cup g_3}$*, i.e.,* $\boldsymbol{X_1 \cup g_i}$ *is the favourite bundle in the partition* $\boldsymbol{X_1 \cup g_i}$*,* $\boldsymbol{X_2}$ *and* $\boldsymbol{X_3 \setminus g_i}$ *for all* $\boldsymbol{i \in \{2,3\}}$*.* From Observation 7, we have $X_3 \setminus g_i >_i X_2$ for $i \in \{2,3\}$. Therefore, we have,

$$X_2 <_2 X_3 \setminus g_2 <_2 X_1 \cup g_2 \quad \text{and} \quad X_2 <_3 X_3 \setminus g_3 <_3 X_1 \cup g_3.$$

By MMS-feasibility of valuation function $v_3(\cdot)$, we conclude that $X_2 <_3 max_3(Z, Z')$ where (Z, Z') is any valid 2-partition of the good set $X_1 \cup X_3$, as MMS-feasibility implies that $max_3(Z, Z') \geq min_3(X_1 \cup g_3, X_3 \setminus g_3) >_3 X_2$. We run the PR algorithm on the 2-partition $(X_1 \cup g_2, X_3 \setminus g_2)$ with agent 2's valuation $(v_2(\cdot))$. Note that this time we run the PR algorithm with $n = 2$ instead of the usual $n = 3$ in the prior cases. Let (Y_2, Y_3) be the output of the PR algorithm. We let agent 3 choose her favorite among Y_2 and Y_3. Assume without loss of generality she chooses Y_3. Now, consider the allocation X':

$$\text{agent } 1: \; X_2 \quad \text{agent } 2: \; Y_2 \quad \text{agent } 3: \; Y_3.$$

We now analyze the strong envy in the allocation. To this end, we first observe that agents 2 and 3 do not strongly envy anyone.

Observation 9. Y_2 is EFX-feasible for agent 2 and Y_3 is EFX-feasible for agent 3 in X'.

Since (Y_2, Y_3) is the output of the PR algorithm run on $(X_1 \cup g_2, X_3 \setminus g_2)$ with agent 2's valuation function, (i) $Y_2 >_2 Y_3 \setminus h$ for all $h \in Y_3$, and (ii) $Y_2 \geq min_2(X_1 \cup g_2, X_3 \setminus g_2) >_2 X_2$, where the first inequality follows from Lemma 4 and the second inequality follows from the fact that $X_1 \cup g_2 >_2 X_3 \setminus g_2 >_2 X_2$. Therefore Y_2 is EFX-feasible for agent 2.

Now, we look into agent 3. Note that $Y_3 = max_3(Y_2, Y_3)$ as agent 3 picks her favourite among Y_2 and Y_3. Therefore, $Y_3 >_3 Y_2$ where the strict inequality follows due to non-degeneracy. Furthermore, due to the MMS-feasibility of $v_3(\cdot)$ and the fact that (Y_2, Y_3) is a valid 2-partition of the good set $X_1 \cup X_3$, we have $Y_3 = max_3(Y_2, Y_3) >_3 X_2$. Therefore, $Y_3 >_3 max_3(Y_2, X_2)$ and thus Y_3 is an EFX-feasible bundle for agent 3.

Therefore, the only possible strong envy is from agent 1. We now enlist the possible strong envy that may arise from agent 1 and show corresponding reallocations.

- Agent 1 does not strongly envy Y_2 and Y_3: Then we are done as X' is an EFX allocation.
- Agent 1 strongly envies both Y_2 and Y_3: In this case, we have $Y_2 >_1 X_2$ and $Y_3 >_1 X_2$. We run the PR algorithm on the partition (X_2, Y_2, Y_3) with agent 1's valuation function $(v_1(\cdot))$ and let agent 2 pick her favorite bundle from the final partition X'' returned by the PR algorithm. Then, we have a partition that satisfies the invariants and $\phi(X'') > \phi(X)$ as $min_1(X''_1, X''_2, X''_3) >_1 min_1(X_2, Y_2, Y_3) = X_2 >_1 X_1 = \phi(X)$, where the first inequality follows from Lemma 4.
- Agent 1 strongly envies one of Y_2 and Y_3: Let us assume without loss of generality that agent 1 strongly envies Y_2. Let $\overline{Y}_2$ be the minimal subset of Y_2 with respect to set inclusion that agent 1 values more than X_2. Then, consider the partition $X'' = (X''_1, X''_2, X''_3)$ where $X''_1 = X_2$, $X''_2 = \overline{Y}_2$ and $X''_3 = Y_3 \cup (Y_2 \setminus \overline{Y}_2)$. First note that X''_3 is EFX-feasible for agent 3 as $X'_3 = Y_3$ was EFX-feasible in allocation X' and now the bundle X''_1 remains the same, the bundle X'_2 has been compressed further in X'', and $X'_3 \subset X''_3$. Also note that $\phi(X'') = min(v_1(X''_1), v_1(X''_2)) = min(v_1(X_2), v_1(\overline{Y}_2)) = v_1(X_2) > v_1(X_1) = \phi(X)$. If X''_1 and X''_2 are EFX-feasible for agent 1, then partition X'' satisfies the invariants and $\phi(X'') > \phi(X)$ and we are done. So now consider the case when at least one of X''_1 and X''_2 is not EFX-feasible for agent 1. Note that $X''_1 >_1 X''_2 \setminus h$ for all $h \in X''_2$ and $X''_2 >_1 X''_1$ by the fact that $X''_1 = X_2$ and by the definition of $X''_2 = \overline{Y}_2$. Thus, if one of X''_1 or X''_2 is not EFX-feasible for agent 1, then we must have $X''_3 \setminus h' >_1 min_1(X''_1, X''_2)$ for some $h' \in X''_3$. In this case, we run the PR algorithm on the partition (X''_1, X''_2, X''_3) with agent 1's valuation function $v_1(\cdot)$ and let agent 2 pick her favorite bundle from the final partition Z returned by the PR algorithm. Then Z satisfies the invariants and

$$\begin{aligned}\phi(Z) &\geq min(v_1(Z_1), v_1(Z_2), v_1(Z_3)) \\ &\geq min(v_1(X''_1), v_1(X''_2), v_1(X''_3)) \\ &= v_1(X_2) \\ &> v_1(X_1) = \phi(X).\end{aligned}$$

So we are done.

This brings us to the main result of this section.

THEOREM 2. *EFX allocations exist with three agents as long as at least one agent has an MMS-feasible valuation function.*

5. Bounds on Rainbow Cycle Number

In this section, we improve the upper bounds on the rainbow cycle number introduced in (Chaudhury et al. 2021a), thereby implying the existence of approximate EFX allocations with $\tilde{\mathcal{O}}((n/\varepsilon)^{1/2})$ charity. Chaudhury et al. (2021a) gave an upper bound of $R(d) \in \mathcal{O}(d^4)$ and they showed it results in the existence of a $(1-\varepsilon)$-EFX allocation with $\mathcal{O}((n/\varepsilon)^{4/5})$ charity. In the same paper, (Chaudhury et al. 2021a) shows a lower bound of d on $R(d)$. In this section, we show improved bounds on $R(d)$. In particular, we first show in Section 5.1 that $R(d) \in \mathcal{O}(d \log d)$ (making the upper bound almost tight), thereby implying the existence of $(1-\varepsilon)$-EFX allocations with $\tilde{\mathcal{O}}((n/\varepsilon)^{1/2})$ charity. Secondly, in Section 5.2, we show an upper bound of $2d-2$ assuming that every vertex $v \in V_i$ has exactly one incoming edge from any other part $V_j \neq V_i$ and exactly one outgoing edge to some vertex in V_j. We call this number $R_p(d)$. We remark that the lower bound of d in (Chaudhury et al. 2021a) also holds for $R_p(d)$. The upper bound of $2d-2$ immediately improves the upper bound on the zero-sum extremal problem studied in (Alon and Krivelevich 2021, Mészáros and Steiner 2021).

5.1. Almost Tight Upper Bound on $R(d)$

Recall that $R(d)$ is the largest k such that there exists a k-partite digraph G with k classes of vertices V_i so that each part V_i has at most d vertices, for all distinct i, j each vertex in V_i has an incoming edge from some vertex in V_j, and there exists no (directed) rainbow cycle, namely, a cycle in G that contains at most one vertex of each V_i. In this section, we prove the following improved bound which is tight up to the logarithmic factor.

THEOREM 10. *If*

$$k(1-1/d)^{k-1} < 1 \tag{2}$$

then $R(d) < k$. Therefore, $R(d) \leq (1+o(1))d \log d$.

Suppose $k(1-1/d)^{k-1} < 1$. Let S be a random set of k vertices of G obtained by picking a single vertex v_i in each V_i, randomly and uniformly among all vertices of V_i, where all choices are independent. For each vertex v, let E_v be the event that S contains v and contains no other vertex u so that uv is a directed edge. We claim that if $v \in V_i$ then the probability of E_v is at most

$$\frac{1}{|V_i|}(1-1/d)^{k-1}.$$

Indeed, the probability that $v \in S$ is $1/|V_i|$. Conditioning on that, since for every $j \neq i$ there is some $u_j \in V_j$ so that $u_j v$ is a directed edge, and the probability that u_j is in S is $1/|V_j| \geq 1/d$, the probability that v has no in-neighbor in V_j is at most $1 - 1/d$. As the choices are independent, the claim follows. By the union bound, the probability that there is a vertex v so that the event E_v occurs is at most

$$\sum_{i=1}^{k} |V_i| \frac{1}{|V_i|} (1 - {}^{1}\!/_{d})^{k-1} = k(1 - 1/d)^{k-1} < 1.$$

Therefore, with positive probability, every vertex in the induced subgraph of G on S has an in-neighbor. Hence there is such an S and in this induced subgraph, there is a cycle, which contains at most one vertex from each V_i. Thus $R(d) < k$, completing the proof.

Theorems 6 and Theorem 10 then imply Theorem 4.

Remark. The proof above can be derandomized using the method of conditional expectations (cf., e.g., Alon and Spencer (1992), Chapter 16), giving the following.

Proposition 1. *Let G be a k-partite digraph with classes of vertices V_i, each having at most d vertices. Suppose that for all distinct i, j each vertex in V_i has an incoming edge from some vertex in V_j and vice versa, and suppose that (2) holds. Then a rainbow cycle in G can be found by a deterministic polynomial time algorithm.*

We apply the method of conditional expectations to produce a set $S = \{s_1, s_2, \ldots, s_k\}$ of vertices of G, where $s_i \in V_i$, so that every indegree in the induced subgraph of G on S is positive. This is done by choosing the vertices s_i one by one in order, maintaining a potential function ϕ whose value is the conditional expectation of the number of events E_v that hold, given the choices of the vertices s_i made so far.

In the beginning, there are no choices made, and the value of ϕ is the sum

$$\sum_{i=1}^{k} |V_i| \frac{1}{|V_i|} (1 - 1/d)^{k-1} = k(1 - 1/d)^{k-1} < 1.$$

Assuming $s_1, s_2, \ldots, s_{i-1}$ have already been chosen, and the above conditional expectation is still smaller than 1, choose $s_i \in V_i$ as the vertex that minimizes the updated value of the conditional expectation. As the expectation is the average over all possible choices of s_i, this minimum stays below 1. The computation of the required conditional expectations for each of the possible $|V_i| \leq d$ choices of $s_i \in V_i$ can be done efficiently. At the end of the process, the value of the potential function is exactly the number of events E_v that hold, and since this number is below 1, none of them holds. This supplies the required set S. Starting in any vertex of S and moving repeatedly to an in-neighbor of it in S until we reach a vertex we have already visited supplies the desired rainbow cycle.

5.2. A Linear Upper Bound on $R_p(d)$

In this section, we assume graph G satisfies all the properties in Definition 3 and also for all different parts V_i and V_j, each vertex in V_i has exactly one outgoing edge to a vertex in V_j. We call these graphs permutation graphs since the set of edges from any part to any other part defines a permutation.

Definition 4. Given an integer $d > 0$, the permutation rainbow cycle number $R_p(d)$ is the largest k such that there exists a k-partite directed graph $G = (V_1 \cup V_2 \cup \cdots \cup V_k, E)$ such that

- each part has exactly d vertices, i.e., $|V_i| = d$, and
- every vertex has exactly one incoming edge from every part other than the one containing it, and
- every vertex has exactly one outgoing edge to every part other than the one containing it, and
- there exists no cycle C in G that visits each part at most once.

Theorem 11. *For all integers $d > 1$, $R_p(d) < 2d - 1$.*

In the rest of this section, we prove Theorem 11. The proof is by induction.

Basis: For the base case, consider $d = 2$. For the sake of contradiction assume $R(2) \geq 3$ and let $V_1 = \{a_1, a_2\}$, $V_2 = \{b_1, b_2\}$ and $V_3 = \{c_1, c_2\}$ be three different parts. Without loss of generality, we can assume there is an edge from a_1 to b_1 and one from b_1 to c_1. Assuming there is no cycle in this graph, a directed edge from c_1 to a_1 cannot exist, and therefore, a directed edge from c_2 to a_1 exists. Thus, no edge from a_1 to c_2 exists, implying an edge from a_1 to c_1. Also, since there is an edge from b_1 to c_1, there must be an edge from c_1 to b_2 (since there can be none to b_1). Now if there is an edge from b_1 to a_1, the cycle (a_1, b_1) exists, and if there is an edge from b_2 to a_1, the cycle (a_1, c_1, b_2) exists which is a contradiction. Therefore, $R(2) < 3$.

Moreover, we prove that $R_p(1) = 1$. Towards a contradiction assume $R_p(1) \geq 2$ and there are two different parts $V_1 = \{a\}$ and $V_2 = \{b\}$. Then, there exists an edge from a to b and one from b to a, forming a cycle.

Induction step: For $d > 2$, we assume for all $1 < d' < d$, $R_p(d') < 2d' - 1$ and prove $R_p(d) < 2d - 1$. In particular, since $R_p(1) = 1$, for all $d' < d$ we have

$$R_p(d') \leq 2d' - 1. \tag{3}$$

First, we define i-restricted paths, which are the paths that use each part at most once, and except for the last vertex, all vertices are in the first i parts.

Definition 5. We call path $P = v_1 \to v_2 \to \cdots \to v_t$ an i-restricted path if

- $v_1, \ldots, v_{t-1} \in V_1 \cup V_2 \cup \cdots \cup V_i$, and

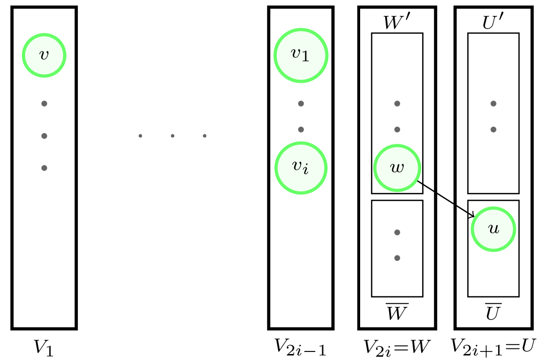

Figure 4 W' has an outgoing edge to $\overline{U}$

- P visits each part at most once.

Note that for all $j > i$, every i-restricted path is also a j-restricted path. Now we prove the following claim.

Claim 1. *If $k \geq 2d-1$, for every vertex v, there is a way of reindexing the parts such that $v \in V_1$ and for all $i \in [d]$, there are i nodes in V_{2i-1} which are reachable from v via $(2i-2)$-restricted paths.*

The proof of the claim is also by induction. For the base case, let $i = 1$. If $v \in U$, set $V_1 = U$, and the claim follows. For the induction step, we assume $V_1, V_2, \ldots, V_{2i-1}$ are already defined and there is a $(2i-2)$-restricted path from v to $v_1, v_2, \ldots, v_i \in V_{2i-1}$. Consider any part $U \notin \{V_1, V_2, \ldots, V_{2i-1}\}$. For all $j \in [i]$, let $v_j \to u_j$ be the outgoing edge from v_j to U. Since each node in V_{2i-1} has exactly one outgoing edge to U, and each node in U has exactly one incoming edge from V_{2i-1}, $u_1, u_2, \ldots, u_i$ are distinct. Therefore, at least i nodes in U are reachable from v via $(2i-1)$-restricted paths. Let $U' \subseteq U$ be the vertices that are reachable from v via $(2i-1)$-restricted paths and let $\overline{U} = U \setminus U'$. If $|U'| \geq i+1$, we set $V_{2i} = W$ for some $W \notin \{V_1, V_2, \ldots, V_{2i-1}, U\}$ and set $V_{2i+1} = U$ and the claim follows. Otherwise, for all $U \notin \{V_1, V_2, \ldots, V_{2i-1}\}$, we have $|U'| = i$ and $|\overline{U}| = d-i$. If there exist $U, W \notin \{V_1, V_2, \ldots, V_{2i-1}\}$ such that $w \in W'$ has an outgoing edge to $u \in \overline{U}$, then we set $V_{2i} = W$ and $V_{2i+1} = U$. Note that all nodes in U' are reachable from v using $(2i-1)$-restricted paths, and u is reachable via a $(2i)$-restricted path. Therefore, in total $i+1$ vertices in $U = V_{2i+1}$ are reachable from v via $(2i)$-restricted paths. See Figure 4 for an illustration.

Let $V(G) = V_1 \cup V_2 \cup \cdots \cup V_{2i-1} \cup U_1 \cup U_2 \cup \cdots \cup U_{k-2i+1}$. The only remaining case is that for all $j \in [k-2i+1]$, $|\overline{U}_j| = d-i$ and for all $j, \ell \in [k-2i+1]$, there is no edge from U'_j to $\overline{U}_\ell$. This means that all the $d-i$ incoming edges of $\overline{U}_\ell$ come from $\overline{U}_j$. Hence all the $d-i$ outgoing edges of $\overline{U}_j$ go to $\overline{U}_\ell$. Therefore, the induced subgraph on $\overline{U}_1 \cup \overline{U}_2 \cup \cdots \cup \overline{U}_{k-2i+1}$, forms a permutation graph. See Figure 5. By Inequality (3), we know $R_p(d-i) \leq 2d-2i-1$ and hence, $k-2i+1 \leq 2d-2i-1$. This contradicts the assumption of the claim, which requires $k \geq 2d-1$. Therefore, this case cannot occur.

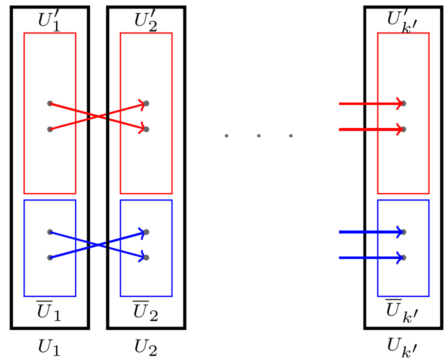

Figure 5 $k' \geq k - 2i - 1$ and for all $j, \ell \in [k']$, there exists no edge between U'_j and $\overline{U}_\ell$.

Back to the assumption step, we want to prove $R_p(d) < 2d - 1$. Towards a contradiction, assume $R_p(d) \geq 2d - 1$ and consider a graph G with $|R_p(d)|$ parts satisfying properties of Definition 4. Now pick an arbitrary vertex v. By setting $i = d$ in Claim 1, there exists a reindexing of the parts such that all d nodes in part V_{2d-1} are reachable from v using $(2d-2)$-restricted paths. Let $u \in V_{2d-1}$ be the vertex with an outgoing edge to v. Then a $(2d-2)$-restricted path from v to u followed by the edge $u \to v$ forms a rainbow cycle. Hence, $R_p(d) < 2d - 1$.

Appendix A:

Non-Degenerate Instances (Chaudhury et al. 2020). We call an instance $\mathcal{I} = \langle [n], M, \mathcal{V} \rangle$ non-degenerate if and only if no agent values two different sets equally, i.e., $\forall i \in [n]$ we have $v_i(S) \neq v_i(T)$ for all $S \neq T$. We extend the technique in (Chaudhury et al. 2020) and show that it suffices to deal with non-degenerate instances when there are n agents with general valuation functions, i.e., if there exists an EFX allocation in all non-degenerate instances, then there exists an EFX allocation in all instances.

Let $M = \{g_1, g_2, \ldots, g_m\}$. We perturb any instance $\mathcal{I}$ to $\mathcal{I}(\varepsilon) = \langle [n], M, \mathcal{V}(\varepsilon) \rangle$, where for every $v_i \in \mathcal{V}$ we define $v'_i \in \mathcal{V}(\varepsilon)$, as

$$v'_i(S) = v_i(S) + \varepsilon \cdot \sum_{g_j \in S} 2^j \qquad \forall S \subseteq M$$

LEMMA 5. *Let $\delta = \min_{i \in [n]} \min_{S,T:\, v_i(S) \neq v_i(T)} |v_i(S) - v_i(T)|$ and let $\varepsilon > 0$ be such that $\varepsilon \cdot 2^{m+1} < \delta$. Then*

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

For the first statement of the lemma, observe that

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

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

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

Acknowledgments

Noga Alon was supported by NSF Grant DMS-2154082 and BSF Grant 2018267. Bhaskar Ray Chaudhury was partially supported by NSF Grants CCF-1942321 and CCF-1750436. Jugal Garg was supported by NSF Grant CCF-1942321. Ruta Mehta was supported by NSF Grant CCF-1750436.

References

Akrami H, Garg J, Sharma E, Taki S (2023) Simplification and improvement of MMS approximation. *Proc. 32nd Intl. Joint Conf. Artif. Intell. (IJCAI).*

Akrami H, Rezvan R, Seddighin M (2022) An EF2X allocation protocol for restricted additive valuations. *Proc. 31st Intl. Joint Conf. Artif. Intell. (IJCAI)*, 17–23.

Alon N, Krivelevich M (2021) Divisible subdivisions. *J. Graph Theory* 98(4):623–629.

Alon N, Spencer J (1992) *The Probabilistic Method* (John Wiley).

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

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

Aziz H, Gaspers S, Mackenzie S, Walsh T (2015) Fair assignment of indivisible objects under ordinal preferences. *Artif. Intell.* 227:71–92.

Baklanov A, Garimidi P, Gkatzelis V, Schoepflin D (2021) Achieving proportionality up to the maximin item with indivisible goods. *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI*, 5143–5150.

Barman S, Krishnamurthy SK (2017) Approximation algorithms for maximin fair division. *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 647–664.

Barman S, Krishnamurthy SK, Vaish R (2018) Finding fair and efficient allocations. *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 557–574.

Berendsohn BA, Boyadzhiyska S, Kozma L (2022) Fixed-point cycles and EFX allocations. *CoRR* 2201.08753.

Berger B, Cohen A, Feldman M, Fiat A (2022) (Almost full) EFX exists for four agents (and beyond). *Proc. 36th Conf. Artif. Intell. (AAAI).*

Bouveret S, Lemaître M (2016) Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems (AAMAS) 30, 2*, 259–290.

Brams SJ, Kilgour DM, Klamler C (2017) Maximin envy-free division of indivisible items. *Group Decision and Negotiation* 26(1):115–131.

Brams SJ, Taylor AD (1996) *Fair division - from cake-cutting to dispute resolution* (Cambridge University Press).

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

Budish E, Cantillon E (2010) The multi-unit assignment problem: Theory and evidence from course allocation at Harvard. *American Economic Review* 102.

Caragiannis I, Gravin N, Huang X (2019) Envy-freeness up to any item with high Nash welfare: The virtue of donating items. *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, 527–545.

Caragiannis I, Kurokawa D, Moulin H, Procaccia AD, Shah N, Wang J (2016) The unreasonable fairness of maximum Nash welfare. *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, 305–322.

Chaudhury BR, Garg J, Mehlhorn K (2020) EFX exists for three agents. *Proc. 21st Conf. Economics and Computation (EC)*, 1–19 (ACM).

Chaudhury BR, Garg J, Mehlhorn K, Mehta R, Misra P (2021a) Improving EFX guarantees through rainbow cycle number. *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 310–311 (ACM).

Chaudhury BR, Kavitha T, Mehlhorn K, Sgouritsa A (2021b) A little charity guarantees almost envy-freeness. *SIAM J. Comput.* 50(4):1336–1358.

Etkin R, Parekh A, Tse D (2005) Spectrum sharing for unlicensed bands. *In Proceedings of the first IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks.*

Garg J, McGlaughlin P, Taki S (2019) Approximating maximin share allocations. *Proceedings of the 2nd Symposium on Simplicity in Algorithms (SOSA)*, volume 69, 20:1–20:11.

Garg J, Taki S (2020) An improved approximation algorithm for maximin shares. *EC*, 379–380 (ACM).

Ghodsi M, Hajiaghayi MT, Seddighin M, Seddighin S, Yami H (2018) Fair allocation of indivisible goods: Improvements and generalizations. *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 539–556.

Halpern D, Procaccia AD, Psomas A, Shah N (2020) Fair division with binary valuations: One rule to rule them all. *WINE*, volume 12495 of *Lecture Notes in Computer Science*, 370–383 (Springer).

Jahan SC, Seddighin M, Seyed-Javadi SM, Sharifi M (2023) Rainbow cycle number and EFX allocations: (almost) closing the gap. *Proc. 32nd Intl. Joint Conf. Artif. Intell. (IJCAI).*

Kurokawa D, Procaccia AD, Wang J (2018) Fair enough: Guaranteeing approximate maximin shares. *Journal of ACM* 65(2):8:1–27.

Lipton RJ, Markakis E, Mossel E, Saberi A (2004) On approximately fair allocations of indivisible goods. *Proc. 5th Conf. Economics and Computation (EC)*, 125–131.

Mahara R (2021) Extension of additive valuations to general valuations on the existence of EFX. *ESA*, volume 204 of *LIPIcs*, 66:1–66:15 (Schloss Dagstuhl - Leibniz-Zentrum für Informatik).

Mészáros T, Steiner R (2021) Zero sum cycles in complete digraphs. *Eur. J. Comb.* 98:103399.

Moulin H (2019) Fair division in the internet age. *Annual Review of Economics* 11(1):407–441.

Murhekar A, Garg J (2021) On fair and efficient allocations of indivisible goods. *Proc. 35th Conf. Artif. Intell. (AAAI)*, 5595–5602.

Plaut B, Roughgarden T (2020) Almost envy-freeness with general valuations. *SIAM J. Discret. Math.* 34(2):1039–1068.

Pratt JW, Zeckhauser RJ (1990) The fair and efficient division of the Winsor family silver. *Management Science* 36(11):1293–1301.

Procaccia AD (2020) Technical perspective: An answer to fair division's most enigmatic question. *Commun. ACM* 63(4):118, ISSN 0001-0782, URL `http://dx.doi.org/10.1145/3382131`.

Steinhaus H (1948) The problem of fair division. *Econometrica* 16(1):101–104.

Vossen TW (2002) *Fair allocation concepts in air traffic management.* Ph.D. thesis, University of Maryland, College Park.