

EF2X Exists for Four Agents

Arash Ashuri¹, Vasilis Gkatzelis², Alkmini Sgouritsa³,

¹Sharif University of Technology, Iran

²Drexel University, USA

³Athens University of Economics and Business, and Archimedes/Athena RC, Greece
arash.ashoori0330@sharif.edu, gkatz@drexel.edu, alkmini@aueb.gr

Abstract

We study the fair allocation of indivisible goods among a group of agents, aiming to limit the envy between any two agents. The central open problem in this literature, which has proven to be extremely challenging, is regarding the existence of an EFX allocation, i.e., an allocation such that any envy from some agent i toward another agent j would vanish if we were to remove any single good from the bundle allocated to j . Prior work has shown that when the agents' valuations are additive, which has been the main focus of prior works, an EFX allocation is guaranteed to exist for all instances involving up to three agents. Subsequent work extended this guarantee to more general valuations, like nice-cancelable and MMS-feasible. However, the existence of EFX allocations for instances involving four agents remains open, even for additive valuations.

We contribute to this literature by focusing on EF2X, a relaxation of EFX which requires that any envy toward some agent would vanish if any *two* of the goods allocated to that agent were to be removed. Our main result shows that EF2X allocations exist for any instance with four agents, even for the class of cancelable valuations, which is more general than additive. Our proof is constructive, proposing an algorithm that computes such an allocation in pseudo-polynomial time. Furthermore, for instances involving three agents we provide an algorithm that computes an EF2X allocation in polynomial time, in contrast to EFX for which the fastest known algorithm for three agents is only pseudo-polynomial.

1 Introduction

During the last decade, the main focus of the fair division literature has been on allocating a set M of indivisible goods, and one of the main fairness goals within this line of work is to limit the amount of envy between any two agents. Since it is well-known that some amount of envy is inevitable when allocating indivisible goods (e.g., one of the goods could be more valuable than all other goods combined, so any agent who receives it is bound to be envied), a vast amount of work has focused on relaxations of “envy-freeness” and the extent to which they can be achieved or approximated.

A central problem within this literature is regarding the existence of allocations satisfying the EFX property, introduced by Caragiannis et al. (2019). An allocation is EFX

if removing any good from the bundle allocated to an agent would ensure that no agent envies the remaining bundle. Despite a plethora of attempts to prove the existence (or non-existence) of EFX allocations in general, this problem remains open. Plaut and Roughgarden (2020) proved the existence of EFX allocations for two agents with general monotone valuations. For instances with three agents, Chaudhury, Garg, and Mehlhorn (2024) were able to prove the existence of EFX allocations for additive valuations, and this was subsequently extended to nice-cancelable valuations (a special case of cancelable valuations) by Berger et al. (2022) and even more general (e.g., MMS-feasible) valuations by (Akrami et al. 2023). The existence of EFX allocations for instances involving four agents remains a very challenging open problem, even in the case of additive valuations.

Our work focuses on an interesting relaxation of EFX, namely EF2X. An allocation is EF2X if removing any *two* goods from the bundle allocated to an agent would ensure that no agent envies the remaining bundle. This notion was introduced by Akrami, Rezvan, and Seddighin (2022), who proved that an EF2X allocation exists for instances involving any number of agents, but with restricted additive valuations: an agent's value for a bundle equals the sum of their value for each good, but their value for each good g is restricted to be either 0 or $v(g)$, where the latter value is the same for all agents. In a very recent work, Kaviani, Seddighin, and Shahrezaei (2024) proved the same result for a different class of valuations ($(\infty, 1)$ -bounded valuations). This class defines a good to be relevant for an agent if its marginal value is not always zero, and the crucial restriction is that any two agents share at most one relevant good.

Our Results

Rather than significantly restricting the valuations that the agents may have, we focus on valuations that are much more expressive, and we study the existence of EF2X allocations for instances involving four agents. Specifically, we consider *cancelable valuations*, introduced by Berger et al. (2022), which strictly generalize the well-studied class of additive valuations. A valuation v is cancelable if for any two bundles $S, T \subset M$ and any good $g \in M \setminus (S \cup T)$, we have that $v(S \cup \{g\}) > v(T \cup \{g\})$ implies $v(S) > v(T)$.

Our main result uses a constructive argument to prove that EF2X allocations always exist for any instance involving

four agents with cancelable valuations.

Theorem: For every instance involving four agents with cancelable valuation functions and any number of goods there exists an EF2X allocation.

In fact, we only require three of these agents to have cancelable valuations; our result holds even if one agent has arbitrary monotone valuations. Furthermore, we show that we can compute such an allocation in pseudopolynomial time.

Proving this result is quite demanding, as it inherits a lot of the complexity of the EFX problem. Note that proofs of EFX existence for three agents tend to be quite long, requiring non-trivial techniques and careful case analysis. As a result, extending those arguments to instances with four agents becomes rather intractable. We overcome this obstacle by introducing some new notions and techniques that are likely to be useful more broadly, e.g., for solving the EFX problem; in fact, the allocation that we compute is often EFX, not just EF2X. These contributions allow us to make the problem more tractable, but our algorithm and its analysis are still quite long and technically demanding.

As a secondary result, we apply the same tools to instances involving three agents, leading to an algorithm that computes an EF2X allocation in polynomial time.

Theorem: For every instance involving three agents with cancelable valuations and any number of goods we can compute an EF2X allocation in polynomial time.

Due to space limitations, we provide only an overview of our results and defer most of the technical content to the full version of the paper (Ashuri, Gkatzelis, and Sgouritsa 2024).

Further Related Work

EFX allocations have been shown to exist for any number of agents when these agents have identical valuations (Plaut and Roughgarden 2020), lexicographic preferences (Hosseini et al. 2021), additive valuations with at most two types of goods (Gorantla, Marwaha, and Velusamy 2023), as well as binary valuations (Halpern et al. 2020), subsequently extended to bi-valued valuations (Amanatidis et al. 2021). Christodoulou et al. (2023) proved EFX allocations exist when the agent’s valuations are captured by a graph.

Another line of research has aimed to achieve multiplicative approximations of EFX. Plaut and Roughgarden (2020) showed the existence of $1/2$ -EFX allocations for subadditive valuations (subsequently derived in poly-time by Chan et al. (2019)), and Amanatidis, Markakis, and Ntokos (2020) proved the existence of $1/\phi \approx 0.618$ -EFX allocations for additive valuations. For more restrictive valuation functions, Markakis and Santorinaios (2023) and Amanatidis, Filos-Ratsikas, and Sgouritsa (2024) proved the existence of $2/3$ -EFX allocations. Barman, Kar, and Pathak (2024) achieved improved EFX approximations for restricted settings.

For more general settings, a lot of work has focused on relaxations of EFX. One such relaxation has focused on “partial allocations,” donating some of the goods to charity and achieving EFX with the rest. This was first studied by Caragiannis, Gravin, and Huang (2019) who showed the exist-

tence of a partial allocation that satisfies EFX and its Nash social welfare is half of the maximum possible. Chaudhury et al. (2021) proved that EFX allocations exist for n agents if up to $n - 1$ goods can be donated; moreover, no agent envies the donated bundle. Berger et al. (2022) improved this number to $n - 2$ goods with nice cancelable valuations, and Mahara (2024) extended this to monotone valuations. Berger et al. (2022) further showed that an EFX allocation exists for 4 agents with at most one donated good. The number of donated goods was subsequently improved at the expense of achieving $(1 - \varepsilon)$ -EFX, instead of exact EFX (Chaudhury et al. 2023; Akrami et al. 2023; Berendsohn, Boyadzhyska, and Kozma 2022; Jahan et al. 2023).

The existence of (approximate) EFX has also been studied along with Nash social welfare (NSW) guarantees. Amanatidis et al. (2021) showed that for additive bi-valued preferences, maximizing the NSW is always EFX. The tradeoff between EFX and NSW was recently studied by Feldman, Mauras, and Ponitka (2024).

Another well-studied relaxation of EFX is *envy freeness up to one good* (EF1), introduced by Budish (2010). An allocation is EF1 if for every pair of agents i and j , there exists some good in j ’s bundle whose removal would ensure that i does not envy j . Lipton et al. (2004) showed that EF1 allocations always exist, even for general monotone valuations, and can be computed efficiently. There have been several other relaxations for EFX that are stronger than EF1, like EFL (Barman et al. 2018), EFR (Farhadi et al. 2021), EEFX and MXS (Caragiannis et al. 2023), and their combinations (e.g., MXS and EFL (Ashuri and Gkatzelis 2024)).

For the case of chore allocation, recently Christoforidis and Santorinaios (2024) showed that the existence of EFX is not guaranteed even for instances with just 3 agents and 6 chores. The non-existence of EFX was also known for mixed goods and chores (mixed manna), where the valuations are not monotone (Bérczi et al. 2024; Hosseini et al. 2023).

A broader overview of discrete fair division can be found in a recent survey by Amanatidis et al. (2023).

2 Preliminaries

An instance of discrete fair division is a tuple $\langle N, M, V \rangle$, where $N = [n] = \{1, 2, \dots, n\}$ is a set of agents, M is a set of m indivisible goods, and $V = (v_1, v_2, \dots, v_n)$ is a profile of valuation functions, where $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$ for each agent $i \in N$ determines i ’s value for each subset of goods. Whenever $v_i(S) < v_i(T)$, agent i strictly prefers set T to set S , and we denote this preference by $S \prec_i T$; similarly, we use $S \preceq_i T$ to denote weak preference, meaning that $v_i(S) \leq v_i(T)$. For notational simplicity, we sometimes use $v(g)$ to denote $v(\{g\})$, i.e., the value for a single good, $S \cup g$ to denote $S \cup \{g\}$, and $S \setminus g$ to denote $S \setminus \{g\}$. Moreover, when we perform a union over disjoint sets, we sometimes use \sqcup instead of \cup to emphasize their disjointness.

Types of valuation functions. We consider valuation functions that are monotone, i.e., for any $S \subseteq T \subseteq M$, $v(S) \leq v(T)$. A valuation function v is *non-degenerate* if $v(S) \neq v(T)$ for any two different bundles S, T . For convenience, throughout the paper, we assume that the valua-

tion function of agent 1 is non-degenerate (in the full version of the paper (Ashuri, Gkatzelis, and Sgouritsa 2024) we show that this is without loss of generality). One of the most well-studied classes of valuation functions is *additive*: a valuation function v is additive if the value for any bundle $S \subseteq M$ is equal to the sum of the values of its goods, i.e., $v(S) = \sum_{g \in S} v(g)$. In this paper, we focus on the more general class of *cancelable* valuations: a valuation function v is *cancelable* (Berger et al. 2022) if for any two bundles $S, T \subset M$ and any good $g \in M \setminus (S \cup T)$, if $v(S \cup g) > v(T \cup g)$, then $v(S) > v(T)$, i.e., removing the same good from two different bundles would not change the relative preference between the two. It is easy to verify that for any cancelable valuation v and bundles S, T, R such that $R \subseteq M \setminus (S \cup T)$ we have the following two properties, which we heavily use throughout the paper:

$$\begin{aligned} v(S \cup R) > v(T \cup R) &\Rightarrow v(S) > v(T) \quad \text{and} \\ v(S \cup R) \leq v(T \cup R) &\Leftarrow v(S) \leq v(T). \end{aligned}$$

A class of valuations that is even more general than cancelable valuations¹ is *MMS-feasible* valuations: a valuation function v is MMS-feasible if for every bundle $S \subseteq M$ and any two bipartitions of S , $\mathbf{Y} = (Y_1, Y_2)$ and $\mathbf{Y}' = (Y'_1, Y'_2)$, we have $\max(v(Y_1), v(Y_2)) \geq \min(v(Y'_1), v(Y'_2))$.

Sets of bundles and notation. Throughout the paper, we use $\mathbf{X} = (X_1, X_2, \dots, X_k)$ and variants such as \mathbf{X}' and $\tilde{\mathbf{X}}$ to denote a partition of all goods M into k (pairwise disjoint) bundles, i.e. $\bigcup_{i \in [k]} X_i = M$. If the union of the bundles does not need to be equal to the set of all goods, we instead use $\mathbf{Y} = (Y_1, Y_2, \dots, Y_k)$. We let $L_i(\mathbf{Y}) \in \arg \min_{j \in [k]} v_i(Y_j)$ be a least valuable bundle in \mathbf{Y} from agent i 's perspective, and $H_i(\mathbf{Y}) \in \arg \max_{j \in [k]} v_i(Y_j)$ to be a most valuable bundle. In the extreme case where $v_i(Y_j)$ is the same for all $j \in [k]$ (and $k \geq 2$), we assign different bundles to $L_i(\mathbf{Y})$ and $H_i(\mathbf{Y})$, so it always holds that $L_i(\mathbf{Y}) \neq H_i(\mathbf{Y})$. For simplicity we write $L_i(Y_1, Y_2, \dots, Y_k)$ instead of $L_i((Y_1, Y_2, \dots, Y_k))$, and similarly for H_i .

Given an ordering $g_1 \preceq_i g_2 \preceq_i \dots \preceq_i g_{|S|}$ of the goods of some bundle S induced by the preferences of some agent i , we use $\ell_i^k(S)$ for some natural number $k \leq |S|$ to denote the k -th least valued good in bundle S from agent i 's perspective, i.e., $\ell_i^k(S) = g_k$. We also denote by $h_i^k(S)$ the k -th highest valued good in S from agent i 's perspective, i.e., $h_i^k(S) = g_{|S|+1-k}$. In the extreme case that $v_i(g)$ is the same for all $g \in S$ (and $|S| \geq 2$), we assign different goods to $\ell_i^1(S)$ and $h_i^1(S)$, so we always have $\ell_i^1(S) \neq h_i^1(S)$.

{EF, EFX, EF2X}-envy. We say that an agent i envies some bundle T relative to bundle S , if $S \prec_i T$. We say that agent i EFX-envies T relative to S if there exists some $g \in T$ such that $S \prec_i T \setminus g$, and we denote this by $S \ll_i T$. We

further write $S \not\ll_i T$ if agent i does not EFX-envy T relative to S . For simplicity, we use $(Y_1, \dots, Y_p) \not\ll_i (Y'_1, \dots, Y'_q)$ if i does not EFX-envy any of the bundles Y'_1, \dots, Y'_q relative to any of the bundles Y_1, \dots, Y_p . We say that agent i EF2X-envies T relative to S if there exist two distinct goods $g_1, g_2 \in T$ such that $S \prec_i T \setminus \{g_1, g_2\}$.

{EFX, EF2X}-feasibility. Given a set of disjoint bundles of goods $\mathbf{Y} = (Y_1, \dots, Y_k)$, we say that bundle Y_j is EFX-feasible for agent i (or for valuation v_i) in \mathbf{Y} if agent i does not EFX-envy any other bundle in \mathbf{Y} relative to Y_j . Similarly, we say that Y_j is EF2X-feasible for agent i (or for valuation v_i) in \mathbf{Y} if agent i does not EF2X-envy any bundle in \mathbf{Y} relative to Y_j . We say that partition $\mathbf{X} = X_1, \dots, X_n$ is EFX-feasible (respectively EF2X-feasible) if there is a way to assign each bundle in \mathbf{X} to a distinct agent in N , such that each bundle is EFX-feasible (respectively EF2X-feasible) for the agent it is assigned to in \mathbf{X} .

{EFX, EF2X}-best sets. Given a bundle $S \subseteq M$ and some natural number $k \leq |S|$, we use $G_k(S) = \{T \subseteq S : |T| = k\}$ to denote the set of subsets of S of size k . Then, if $T^* \in \arg \max_{T \in G_k(S)} v_i(S \setminus T)$, we refer to $S \setminus T^*$, i.e., the best subset of S that one can get after removing k of its goods, as the best- k -remaining goods of S w.r.t. agent i 's valuation. We refer to $w_i^k(S) = v_i(S \setminus T^*)$ as the *best- k -value* of a set S for agent i ; if $k > |S|$, we define the best- k -value of a set S for any agent i to be 0. We say that bundle X_j is *EFX-best* and *EF2X-best* for agent i in \mathbf{X} , if it has the maximum best-1-value and best-2-value, respectively, among all bundles in \mathbf{X} . We write $\text{EFXF}_i(\mathbf{X})$, $\text{EF2XF}_i(\mathbf{X})$, and $\text{EFXBest}_i(\mathbf{X})$ to denote the set of EFX-feasible, EF2X-feasible, and EFX-best bundles for agent i in partition \mathbf{X} , respectively.

The Plaut-Roughgarden (PR) algorithm. This local search algorithm, dubbed the *PR algorithm* by Akrami et al. (2023), takes as input a set of bundles $\mathbf{Y} = (Y_1, Y_2, \dots, Y_k)$ and a valuation function v , and it returns a set of bundles $\mathbf{Y}' = (Y'_1, Y'_2, \dots, Y'_k)$ that is EFX-feasible for valuation v in \mathbf{Y}' (i.e., all of its bundles are EFX-feasible in \mathbf{Y}'). If Y_j is the least valuable bundle in \mathbf{Y} w.r.t. valuation v , the algorithm checks whether Y_j is EFX-feasible w.r.t. v in \mathbf{Y} . If it is, the algorithm terminates and returns \mathbf{Y} . On the other hand, if Y_j is not EFX-feasible in \mathbf{Y} , then there must exist another bundle Y_i such that $v(Y_j) < v(Y_i \setminus g)$ for some good $g \in Y_i$. The PR algorithm then removes g from Y_i , adds g to Y_j , and repeats the same sequence of steps until it reaches an EFX-feasible set of bundles.

3 Overview of Algorithm for Four Agents

Our main result (Theorem 1) proves the existence of an EF2X allocation for instances involving an arbitrary number of goods and four agents with cancelable valuations; in fact, one of them can have an arbitrary monotone valuation.

Theorem 1. *For every instance involving four agents with cancelable valuation functions and any number of goods, there exists an EF2X allocation and we can compute one in pseudo-polynomial time.*

¹The fact that every cancelable valuation is MMS-feasible is shown in Lemma 2 of Akrami et al. (2023). They claim that “nice” cancelable valuation are MMS-feasible, but their definition of nice cancelable coincides with our definition of a cancelable function, which is also the one originally used by Berger et al. (2022).

As the theorem statement suggests, instead of a purely existential argument, we provide a constructive proof using an algorithm that computes an EF2X allocation in pseudo-polynomial time. The algorithm starts with some arbitrary initial partition of the goods into four bundles and it gradually manipulates this until it reaches a partition that is EF2X-feasible. Given the complexity of the problem, the ways in which the algorithm gradually transforms a partition aiming to achieve EF2X-feasibility are quite non-trivial and technical. Therefore, to be able to clearly present this algorithm, we define a sequence of stages (formally defined later on), each of which determines a list of structural properties that a partition in that stage needs to satisfy. Given these stages, we can provide a high-level description of this complicated algorithm using the simple diagram of Figure 1.

Our algorithm follows a sequence of steps, each of which takes as input some partition $\mathbf{X} = (X_1, X_2, X_3, X_4)$ and returns a partition $\tilde{\mathbf{X}} = (\tilde{X}_1, \tilde{X}_2, \tilde{X}_3, \tilde{X}_4)$. The way in which each step transforms \mathbf{X} into $\tilde{\mathbf{X}}$ depends on the stage that \mathbf{X} is in, i.e., on the structural properties that \mathbf{X} satisfies. Using these properties, our algorithm carefully re-arranges the goods across bundles, and our proof shows that the resulting partition, $\tilde{\mathbf{X}}$, satisfies all of the corresponding properties of a new stage. The diagram of Figure 1 illustrates the high-level structure of the algorithm by providing a node for each stage that a partition \mathbf{X} may be in, and a directed edge between two stages for each step that transforms a partition \mathbf{X} to $\tilde{\mathbf{X}}$. E.g., the edge from stage B1 to stage B2 corresponds to a transformation that takes as input a partition \mathbf{X} in stage B1 and returns a partition $\tilde{\mathbf{X}}$ in stage B2 (see Theorem 22).

The algorithm starts with an arbitrary initial partition and it uses the PR algorithm to get a partition whose bundles are all EFX-feasible for agent 1. It then turns this into an in-stage partition (the first stage in Figure 1). From that point onward, at the end of every step (i.e., after using each edge of Figure 1), any bundle in the resulting partition $\tilde{\mathbf{X}}$ will have at least one agent for whom it is EFX-feasible in $\tilde{\mathbf{X}}$. If we can match the agents to bundles that are EF2X-feasible for them, then the algorithm terminates. If not, then we essentially identify an “under-demanded” bundle (one that is EFX-feasible for agent 1, but not EF2X-feasible for any other agent) and some “over-demanded” bundles. Each step then carefully re-distributes goods across these bundles to make the under-demanded bundle more appealing.

Potential function and running time. To ensure that the algorithm is “making progress” toward its goal of reaching an EF2X allocation, we use a potential function that weakly increases after each step. Specifically, we set aside one agent (Agent 1), who plays a special role throughout the algorithm, and we use that agent’s value for the first bundle of the partition as the potential of that partition.

Definition 2 (potential function). *The potential of a partition \mathbf{X} is $\phi(\mathbf{X}) = v_1(X_1)$.*

Note that the bold edges in Figure 1 (all edges going upward) correspond to steps of the algorithm that guarantee a strict increase of the potential. All other edges ensure that the potential does not drop. Using this observation, combined

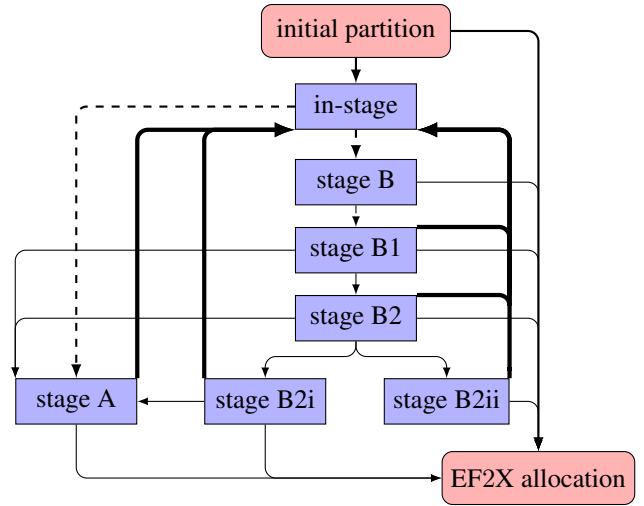


Figure 1: Dashed edges correspond to transitions where the partition remains the same. Non-dashed edges correspond to transitions where the partition may change. These edges are bold when the potential strictly increases during the transition and non-bold when the potential weakly increases.

with the fact that the number of potential increases can be at most pseudo-polynomial and that each step takes pseudo-polynomial time, we can conclude that the overall running time of this algorithm is pseudo-polynomial (details deferred to full version (Ashuri, Gkatzelis, and Sgouritsa 2024)).

Formal Definitions of the Algorithm’s Stages

We now provide the formal definitions of the stages.

Definition 3 (stage A). *We say partition \mathbf{X} is in stage A if $X_1, X_2, X_3 \in \text{EFXF}_1(\mathbf{X})$ with $X_1 = L_1(X_1, X_2, X_3)$, and at least one of the following holds:*

- i) X_4 is EFX-best for at least one agent $j \neq 1$ and $X_4 \setminus \ell_j^1(X_4) \succeq_j X_1 \cup \ell_j^1(X_4)$,
- ii) X_4 is EFX-best for at least two distinct agents $i, j \neq 1$.

Definition 4 (stage B). *We say partition \mathbf{X} is in stage B if $X_1, X_2 \in \text{EFXF}_1(\mathbf{X})$, $X_1 = L_1(X_1, X_2)$, $X_3 \in \text{EFXF}_i(\mathbf{X})$ for some agent $i \neq 1$, and $X_4 \in \text{EFXF}_j(\mathbf{X})$ for some agent $j \notin \{1, i\}$.*

Definition 5 (in-stage partition). *We say that a partition is in-stage if it is either in stage A or in stage B.*

Definition 6 (stage B1). *We say partition \mathbf{X} is in stage B1 if it is in stage B, $X_4 \in \text{EFXBest}_j(\mathbf{X})$, and also:*

$$X_4 \setminus \ell_j^1(X_4) \succ_j (X_1, X_2)$$

Definition 7 (swap-optimal sets). *We say that two sets (S, T) are swap-optimal w.r.t. the agents i and j if there exist no $g \in S, h \in T$ such that $g \succeq_i h$, $h \preceq_j g$ and at least one of these two preferences is strict.*

Definition 8 (stage B2). *We say partition \mathbf{X} is in stage B2 if it is in stage B1 and (X_3, X_4) are swap-optimal w.r.t. the agents i and j that satisfy the conditions of stage B.*

Definition 9 (stage B2i). We say partition \mathbf{X} is in stage B2i if it is in stage B2 and also the following conditions hold:

$$\begin{aligned} X_4 \setminus \{\ell_i^1(X_4), \ell_i^2(X_4)\} \succ_i X_1, X_2 \\ X_4 \setminus \{\ell_j^1(X_4), \ell_j^2(X_4)\} \succ_j X_1, X_2 \\ H_i(X_1, X_2) \succeq_i X_3 \setminus \ell_i^1(X_3) \\ X_3 \cup \ell_j^1(X_4) \gg_j X_4 \setminus \ell_j^1(X_4) \end{aligned}$$

Definition 10 (stage B2ii). We say partition \mathbf{X} is in stage B2ii, if it is in stage B2 and also:

$$X_3 \setminus \ell_i^1(X_3) \succ_i (X_1, X_2).$$

Outline of the Paper. Given the multiple subtleties that our algorithm needs to address, it is infeasible to fit it all within the page limit. In Section 4, we provide some useful lemmas. Then, as a warm-up, Section 5 provides the transformations applied to partitions that are either in stage A or in stage B. Section 6 discusses the swap-optimization process, and Section 7 provides a sample of some more technical arguments required for dealing with partitions in stage B1. The even more technical arguments used for partitions in stage B2, stage B2i, stage B2ii, as well as the process for computing EF2X allocation for three agents in polynomial time are deferred to (Ashuri, Gkatzelis, and Sgouritsa 2024). We conclude with Section 8, providing a high-level description of some of the broader technical contributions.

4 Some Useful Lemmas

We now provide some lemmas and observations that we use repeatedly in subsequent proofs. Due to space limitations we defer their proofs to (Ashuri, Gkatzelis, and Sgouritsa 2024), where we also include additional lemmas and observations.

Lemma 11. Given a bundle S and an agent i with a cancellable valuation function, if g_1, \dots, g_k are the k least valued goods in S w.r.t. agent i , then $w_i^k(S) = v_i(S \setminus \{g_1, \dots, g_k\})$.

The following lemma establishes an important property of EFX-best bundles: if some bundle X_j is not EFX-feasible for some agent i in some partition, then this agent EFX-envies the EFX-best bundle in the partition relative to X_j .

Lemma 12. Given partition \mathbf{X} , if $X_a \in \text{EFXBest}_i(\mathbf{X})$, then $X_j \in \text{EFXF}_i(\mathbf{X})$ if and only if $X_j \succeq_i X_a \setminus \ell_i^1(X_a)$. Similar statement holds for EF2X-feasible and EF2X-best bundles.

The following three lemmas are related to some processes for redistributing goods, which we will be using repeatedly.

Lemma 13. Given any partition \mathbf{X} , we can construct a new partition $\tilde{\mathbf{X}}$ that is either EF2X-feasible or in-stage with $\phi(\tilde{\mathbf{X}}) \geq \min_{r \in [4]} v_1(X_r)$. Also, if at least one bundle in \mathbf{X} is not EFX-feasible for agent 1 in \mathbf{X} , then $\phi(\tilde{\mathbf{X}}) > \min_{r \in [4]} v_1(X_r)$.

Lemma 14. Given any partition \mathbf{X} such that X_4 is EFX-feasible for some agent $j \neq 1$, we can construct a new partition $\tilde{\mathbf{X}}$ that is either EF2X-feasible or in-stage with $\phi(\tilde{\mathbf{X}}) \geq \min_{r \in [3]} v_1(X_r)$.

Lemma 15. Suppose \mathbf{X} is a partition such that $X_3 \sqcup X_4 = Y_3 \sqcup Y_4$, $(X_3, X_4) \not\prec_i (X_1, X_2)$ for some agent $i \neq 1$, and $(Y_3, Y_4) \not\prec_j (X_1, X_2)$ for some agent $j \notin \{1, i\}$. Then we can construct a new partition $\tilde{\mathbf{X}}$ that is either EF2X-feasible or in-stage with $\phi(\tilde{\mathbf{X}}) \geq \min_{r \in [2]} v_1(X_r)$.

The following is one of the most important lemmas of the paper. The algorithm starting with a partition \mathbf{X} in stage B will aim to either find an EF2X-feasible partition or satisfy the following lemma's conditions, so it can ensure progress.

Lemma 16. Given a partition \mathbf{X} in stage B, suppose that \mathbf{Y} is a set of disjoint bundles such that $Y_1 = X_1 \cup g$ with $g \in X_3 \cup X_4$, $Y_2 = X_2$, $Y_3 \sqcup Y_4 \subseteq (X_3 \sqcup X_4) \setminus g$. Suppose further, that there exists some agent $i \neq 1$, such that $Y_3 \in \text{EFXF}_i(\mathbf{Y})$, and some agent $j \notin \{1, i\}$, such that $Y_4 \in \text{EFXF}_j(\mathbf{Y})$. Then we can construct a new partition $\tilde{\mathbf{X}}$ that is either EF2X-feasible or in-stage with $\phi(\tilde{\mathbf{X}}) > \phi(\mathbf{X})$.

5 Dealing with Partitions in Stage A or B

We now prove that if we start with a partition \mathbf{X} in Stage A, then we either construct an EF2X-feasible partition, or we construct another in-stage partition with higher potential. We further prove (Theorem 18) that a partition \mathbf{X} that is in stage B is either EF2X-feasible or in stage B1.

Theorem 17. Given any partition \mathbf{X} in stage A, we can construct a new partition $\tilde{\mathbf{X}}$ that is either EF2X-feasible or in-stage with $\phi(\tilde{\mathbf{X}}) > \phi(\mathbf{X})$.

Proof. We first address the case where \mathbf{X} satisfies the condition i of stage A, and then the condition ii .

Stage Ai (X_4 is EFX-best for some agent $j \neq 1$, $X_4 \setminus \ell_j^1(X_4) \succeq_j X_1 \cup \ell_j^1(X_4)$): We define

$$\mathbf{X}' = (X_1 \cup \ell_j^1(X_4), X_2, X_3, X_4 \setminus \ell_j^1(X_4)).$$

Since $X_4 \in \text{EFXBest}_j(\mathbf{X})$, agent j does not EFX-envy bundles X_2, X_3 relative to bundle $X_4 \setminus \ell_j^1(X_4)$. Also, since $X_1 \cup \ell_j^1(X_4) \preceq_j X_4 \setminus \ell_j^1(X_4)$, we get that $X'_4 \in \text{EFXF}_j(\mathbf{X}')$. Additionally, we have that $\min_{r \in [3]} v_1(X'_r) = \min(v_1(X_1 \cup \ell_j^1(X_4)), v_1(X'_2), v_1(X'_3)) > v_1(X_1) = \phi(\mathbf{X})$, since agent 1's valuation is non-degenerate. Therefore, by Lemma 14, we can construct a new partition $\tilde{\mathbf{X}}$ that is either EF2X-feasible or in-stage with $\phi(\tilde{\mathbf{X}}) \geq \min_{r \in [3]} v_1(X'_r) > \phi(\mathbf{X})$, and the theorem follows for this case.

Stage Aii (X_4 is EFX-best for at least two distinct agents $i, j \neq 1$): We consider the case that \mathbf{X} is not in stage Ai, otherwise we would show the theorem as in the previous case. Therefore, for any $u \in \{i, j\}$, we have $X_4 \setminus \ell_u^1(X_4) \preceq_u X_1 \cup \ell_u^1(X_4)$. By rearranging the bundles we define

$$\mathbf{X}' = (X_2, X_3, X_1, X_4).$$

Then, since $X_4 \in \text{EFXBest}_u(\mathbf{X})$, it holds that $X_4 \setminus \ell_u^1(X_4) \not\prec_u (X_2, X_3)$, and so it also holds that $X_1 \cup \ell_u^1(X_4) \not\prec_u (X_2, X_3)$. Also, note that $X_1 \sqcup X_4 = (X_1 \cup \ell_i^1(X_4)) \sqcup (X_4 \setminus \ell_i^1(X_4)) = (X_1 \cup \ell_j^1(X_4)) \sqcup (X_4 \setminus \ell_j^1(X_4))$. Then, by Lemma 15, we can construct a new partition $\tilde{\mathbf{X}}$ that is either EF2X-feasible or in-stage with $\phi(\tilde{\mathbf{X}}) \geq \min(v_1(X_2), v_1(X_3)) > v_1(X_1) = \phi(\mathbf{X})$, since agent 1's valuation is non-degenerate, which concludes the proof. \square

Theorem 18. Any partition \mathbf{X} that is in stage B is either EF2X-feasible or in stage B1.

Proof. \mathbf{X} is in stage B, so $(X_1, X_2) \in \text{EFXF}_1(\mathbf{X})$, and for some distinct agents $i, j \neq 1$, we have $X_3 \in \text{EFXF}_i(\mathbf{X})$ and $X_4 \in \text{EFXF}_j(\mathbf{X})$. Let the remaining agent be $u \notin \{1, i, j\}$. If any of the bundles X_1, X_2 is EFX-feasible for agent u , then \mathbf{X} is EFX-feasible: we assign X_3 to i , X_4 to j , agent u is given their EFX-feasible bundle from $\{X_1, X_2\}$, and the remaining bundle is assigned to agent 1; therefore \mathbf{X} is also EF2X-feasible. If, on the other hand, neither X_1 nor X_2 are EFX-feasible for agent u in \mathbf{X} , then the EFX-best bundle for agent u is either X_3 or X_4 . Assume, w.l.o.g., that X_4 is EFX-best for agent u in \mathbf{X} . Since X_1, X_2 are not EFX-feasible for u , by Lemma 12, we get that $X_4 \setminus \ell_u^1(X_4) \succ_u (X_1, X_2)$, so \mathbf{X} is in stage B1 by setting $j = u$. \square

6 Swap-Optimization Procedure

This section focuses on the swap-optimization process, which plays a crucial role in the rest of the paper.

Definition 19. Given two disjoint bundles S, T and two distinct agents i, j , we call the following operation swap-optimization of (S, T) with respect to (i, j) : as long as there exists some $g \in S$ and $h \in T$ such that $g \preceq_i h$, $g \succeq_j h$, and at least one of these two preferences is strict, swap these two goods across the bundles.

The following lemmas regarding the swap-optimization process are used repeatedly for our subsequent results. We defer their proofs to (Ashuri, Gkatzelis, and Sgouritsa 2024).

Lemma 20. If we swap-optimize two bundles (S, T) w.r.t. two agents i, j , this weakly increases $w_i^k(S)$ and weakly decreases $w_i^k(T)$ for every $k \in \mathbb{N}$. Analogously, $w_j^k(S)$ weakly decreases and $w_j^k(T)$ weakly increases.

Lemma 21. Suppose \mathbf{X} is a partition and we swap-optimize two of its bundles, (X_i, X_j) , w.r.t. agents (i, j) to get bundles (X'_i, X'_j) , leading to partition \mathbf{X}' with $X'_r = X_r$ for all $r \notin \{i, j\}$. Then, for every $u \in \{i, j\}$ and $k \in \mathbb{N}$, if X_u is EFkX-best or EFX-feasible for agent u in \mathbf{X} , then X'_u is EFkX-best or EFX-feasible, respectively, for agent u in \mathbf{X}' .

7 Dealing with Partitions in Stage B1

To provide the reader with a sense of what each step of the algorithm looks like, this section describes how this algorithm handles partitions that are in stage B1.

Theorem 22. Given any partition \mathbf{X} in stage B1, we can construct a new partition $\tilde{\mathbf{X}}$ s.t. one of the following holds:

- $\tilde{\mathbf{X}}$ is EF2X-feasible.
- $\tilde{\mathbf{X}}$ is in-stage with $\phi(\tilde{\mathbf{X}}) > \phi(\mathbf{X})$.
- $\tilde{\mathbf{X}}$ is in stage B2 with $\phi(\tilde{\mathbf{X}}) \geq \phi(\mathbf{X})$.
- $\tilde{\mathbf{X}}$ is in stage A with $\phi(\tilde{\mathbf{X}}) \geq \phi(\mathbf{X})$.

To prove this theorem, we use the following lemma, whose proof we defer to the full version of the paper (Ashuri, Gkatzelis, and Sgouritsa 2024).

Lemma 23. Suppose \mathbf{X} is a partition s.t. $X_1 \prec_1 X_2$ and some agent $j \neq 1$ does not EF2X-envy bundle X_3 relative to bundle X_4 , does not EFX-envy bundle X_2 relative to bundle X_4 , and $X_4 \setminus \ell_j^1(X_4) \succeq_j X_1$. Then, if $X_1 \ll_1 X_3$, we can construct a partition $\tilde{\mathbf{X}}$ that is either EFX-feasible, or in-stage with $\phi(\tilde{\mathbf{X}}) > \phi(\mathbf{X})$, or in stage A with $\phi(\tilde{\mathbf{X}}) = \phi(\mathbf{X})$.

We can now prove the main theorem of this section.

Proof of Theorem 22. We show that there exists either a partition satisfying the theorem's conditions, or a partition $\tilde{\mathbf{X}}$ in stage B1, with $\phi(\tilde{\mathbf{X}}) \geq \phi(\mathbf{X})$ and $|\tilde{X}_3| > |X_3|$. In the second case, we set $\tilde{\mathbf{X}}$ as the initial partition and repeat; since the cardinality of the third bundle strictly increases at each iteration, the procedure will terminate after at most m steps to a partition satisfying the theorem's conditions.

Let i, j be the agents as defined in stage B1. After performing a Swap-Optimization on (X_3, X_4) w.r.t. agents (i, j) , let (X'_3, X'_4) be the two new bundles. Let

$$\mathbf{X}' = (X_1, X_2, X'_3, X'_4),$$

and note that in the swap-optimization, the cardinality of each bundle does not change, so $|X'_3| = |X_3|$. By Lemma 20 and Lemma 11:

$$X'_4 \setminus \ell_j^1(X'_4) \succeq_j X_4 \setminus \ell_j^1(X_4) \succ_j (X_1, X_2),$$

where the second preference is due to the definition of stage B1. Moreover, regarding \mathbf{X} we had $X_3 \in \text{EFXF}_i(\mathbf{X})$ and $X_4 \in \text{EFXBest}_j(\mathbf{X})$, so by Lemma 21, we still have $X'_3 \in \text{EFXF}_i(\mathbf{X}')$ and $X'_4 \in \text{EFXBest}_j(\mathbf{X}')$.

If $X'_1, X'_2 \in \text{EFXF}_1(\mathbf{X}')$, then \mathbf{X}' would be in stage B2 with $\phi(\mathbf{X}') = \phi(\mathbf{X})$, and so the theorem would follow. If we assume otherwise, it should be either $X'_1 \ll_1 X'_3$, or $X'_1 \ll_1 X'_4$, since $X'_1 = L_1(X'_1, X'_2)$. If the first was true, the conditions of Lemma 23 would be satisfied and the theorem would follow. So, we assume that $X'_1 \ll_1 X'_4$.

Let X''_4 be a minimal subset of X'_4 with $X''_4 \succ_1 X'_1$, and let $S = X'_4 \setminus X''_4$. Since $X'_1 \ll_1 X'_4$, it is $S \neq \emptyset$. We define

$$\mathbf{X}'' = (X_1, X_2, X'_3 \cup S, X'_4 \setminus S).$$

Since $S \neq \emptyset$, we get that $|X''_3| > |X'_3| = |X_3|$. Note that since $X'_3 \in \text{EFXF}_i(\mathbf{X}')$ and X''_3 has more goods than X'_3 , X''_3 can lose some good and still be EFX-feasible for agent i in the remaining partition; hence, $X''_3 \in \text{EFXBest}_i(\mathbf{X}'')$. If $X''_1 \ll_1 X''_3$, then by Lemma 13, there exists a partition $\tilde{\mathbf{X}}$ that is either EF2X-feasible or in-stage with $\phi(\tilde{\mathbf{X}}) > \min_{r \in [4]} v_1(X''_r) = v_1(X_1) = \phi(\mathbf{X})$, and the theorem follows. By assuming otherwise, since $L_i(X_1, X_2, X''_4) = X_1$, it holds that $X''_1, X''_2, X''_4 \in \text{EFXF}_1(\mathbf{X}'')$.

If for any of the agents $j \notin \{1, i\}$, it was $X''_3 \in \text{EFXBest}_j(\mathbf{X}'')$, \mathbf{X}'' (after swapping bundles X''_3 and X''_4) would be in the stage A (where the second condition is satisfied), with $\phi(\mathbf{X}'') = \phi(\mathbf{X})$. Considering now that those other two agents, which are not 1 or i , do not have X''_3 as their EFX-best bundle in \mathbf{X}'' , so if they had different EFX-best bundles in \mathbf{X}'' , then \mathbf{X}'' would be EFX-feasible (by the fact that X''_3 is EFX-feasible for i in \mathbf{X}'' and the fact that EFX-best bundles are EFX-feasible). So, consider the last case where both those agents have the same unique

EFX-best bundle in \mathbf{X}'' , which is different from X_3'' . Let this bundle be \tilde{X}_4 , and the two remaining bundles among X_1'', X_2'', X_4'' be \tilde{X}_1, \tilde{X}_2 such that $\tilde{X}_1 \prec_1 \tilde{X}_2$. We finally define the partition

$$\tilde{\mathbf{X}} = (\tilde{X}_1, \tilde{X}_2, X_3'', \tilde{X}_4).$$

Note that since $\tilde{X}_1, \tilde{X}_2 \in \text{EFXF}_1(\tilde{\mathbf{X}})$, $X_3'' \in \text{EFXF}_i(\tilde{\mathbf{X}})$ and \tilde{X}_4 is EFX-feasible for both agents other than 1 and i , if for any of the agents $j \notin \{1, i\}$, \tilde{X}_1 or \tilde{X}_2 is EFX-feasible in $\tilde{\mathbf{X}}$, this partition would be EFX-feasible, so assume otherwise. So for any agent $j \notin \{1, i\}$, $(\tilde{X}_1, \tilde{X}_2) \notin \text{EFXF}_j(\tilde{\mathbf{X}})$, and therefore, by Lemma 12, $\tilde{X}_4 \setminus \ell_j^1(\tilde{X}_4) \succ_j (\tilde{X}_1, \tilde{X}_2)$. Moreover, since $X_3'' \in \text{EFXBest}_i(\mathbf{X}'')$, it holds that $\tilde{X}_3 \in \text{EFXF}_i(\tilde{\mathbf{X}})$, since $\tilde{\mathbf{X}}$ is just a rearrangement of the bundles of \mathbf{X}'' . Overall, $\tilde{\mathbf{X}}$ is in stage B1, with $\phi(\tilde{\mathbf{X}}) \geq \min(v_1(X_1''), v_1(X_2''), v_1(X_4'')) = v_1(X_1) = \phi(\mathbf{X})$, and $|\tilde{X}_3| = |X_3''| > |X_3|$. By setting $\mathbf{X} = \tilde{\mathbf{X}}$ and repeatedly transforming the partition as above, we reach a partition that satisfies the conditions of the theorem since we strictly increase the cardinality of the third set in each round. \square

8 Broader Technical Contributions

We conclude with a high-level summary of some of the new notions and techniques that can be of independent interest.

EFkX-Best-Bundle. One concept that we introduce is the EFkX-best bundle in a set of bundles \mathbf{Y} for some agent i (see Section 2 for a definition). This satisfies the very useful property that any bundle that is EFkX-best (and, hence, also EFkX-feasible) for some agent i , remains EFkX-feasible even if we remove i 's k least valued goods from it (and keep those goods unallocated). This holds since the remaining bundle is, by definition, at least as valuable for i as any other bundle T , after the removal of any k goods from T . This is a property that we leverage when, during some reallocation process, we remove goods from some bundle, temporarily leading to a *partial* partition.

To extract useful conditions for *full* partitions, we combine the above idea with the following observation: Given a partition, let S be the EF2X-best bundle for some agent i and T be another bundle that is not EF2X-feasible for i . Then, if i received T , she would EF2X-envy S ; even if we removed i 's two least valuable goods from S , she would still prefer S to T . However, the trick here is to instead remove just the *one* least valuable good of S and add it to T . This gives a (full) partition where agent i still prefers the new bundle S to the new bundle T (this is due to cancelable valuations). Moreover, since S was EF2X-best for i before, it remains at least EF2X-feasible for i after this change. This way, some “underdemanded” bundle has gained a new good, and the one that was EF2X-best for i is still at least EF2X-feasible for her, even though it lost a good. Such re-allocations allow us to make progress toward the final allocation.

Although this procedure will eventually increase the number of bundles that are EF2X-feasible for an agent, we cannot use it by itself: it may compromise EF2X-feasibility for other agents, who may end up EF2X-envying the bundle that

gained a good. To avoid this, we require conditions for other agents (one example is achieved by the swap-optimization that we discuss next) and possibly some case analysis.

Swap-Optimization. We also introduce a procedure we call *swap-optimization*, in which two bundles, each associated with a different agent, swap items in such a way that those agents are more satisfied with their new bundles.

Swap-Optimization is useful to overcome some issues discussed at the end of the EFkX-Best-Bundle paragraph. Specifically, consider a partition of the goods such that three bundles, S, T , and Q , satisfy that S is EF2X-best for agent i , T is EF2X-best for agent j , and Q is not EF2X-feasible for either of them. Given such conditions, the EFkX-Best-Bundle paragraph suggested moving to Q either i 's least valued good from S or j 's least valued good from T . To ensure that this does not cause any envy, we can first perform swap-optimization between S, T w.r.t. i, j , which maintains these initial conditions (by Lemmas 21 and 12). If g is i 's least valuable good in S and h is j 's least valuable good in T after the swap-optimization, we move to Q i 's least valuable good in $\{g, h\}$. This guarantees that neither i nor j will envy the resulting bundle Q . To see this, suppose i weakly prefers h to g . Then, j also weakly prefers h to g , since S and T were swap-optimized. As we argued in the EFkX-Best-Bundle paragraph, i will not envy the resulting bundle Q after adding g to it. Similarly, agent j would not envy Q even if we were to move h from T to Q , so she will not envy Q if we add g to it, which she values less, and keep T unchanged. Similar arguments apply if the moved good is h .

Maintaining EFX-feasibility for 2 agents & more goods.

The following fact also plays a crucial role in our algorithm. Suppose i, j are two agents with MMS-feasible valuations, and $\mathbf{Y} = (Y_1, Y_2)$ are two disjoint bundles such that Y_1 and Y_2 are EFX-feasible for agents i, j , respectively, in \mathbf{Y} . Suppose now that we want to introduce a set S of additional goods and re-allocate, so that the above conditions are preserved. In the full version of the paper (Ashuri, Gkatzelis, and Sgouritsa 2024) we show that it is possible to rearrange all those goods and derive a set of two disjoint bundles $\mathbf{Y}' = (Y'_1, Y'_2)$ (where $Y'_1 \sqcup Y'_2 = Y_1 \sqcup Y_2 \sqcup S$) such that Y'_1 and Y'_2 are EFX-feasible for agents i, j , respectively, in \mathbf{Y}' , and moreover, both agents are weakly more satisfied than before, i.e., $Y'_1 \succeq_i Y_1$ and $Y'_2 \succeq_j Y_2$.

This tool is useful in cases where two agents are EFX-feasible with one distinct bundle each, and in our attempt to satisfy another agent (agent 1), we need to remove some goods from another bundle and then add them to the bundles associated with these two agents without breaking the condition that those bundles are EFX-feasible for them.

Acknowledgements

Vasilis Gkatzelis was partially supported by NSF CAREER award CCF-2047907. Regarding Alkmini Sgouritsa, the research project is implemented in the framework of H.F.R.I call “Basic research Financing (Horizontal support of all Sciences)” under the National Recovery and Resilience Plan “Greece 2.0” funded by the European Union–NextGenerationEU (H.F.R.I. Project Number: 15635).

References

- Akrami, H.; Alon, N.; Chaudhury, B. R.; Garg, J.; Mehlhorn, K.; and Mehta, R. 2023. EFX: A Simpler Approach and an (Almost) Optimal Guarantee via Rainbow Cycle Number. In *Proceedings of the 24th ACM Conference on Economics and Computation, EC 2023*, 61. ACM.
- Akrami, H.; Rezvan, R.; and Seddighin, M. 2022. An EF2X Allocation Protocol for Restricted Additive Valuations. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022*, 17–23. ijcai.org.
- Amanatidis, G.; Aziz, H.; Birmipas, G.; Filos-Ratsikas, A.; Li, B.; Moulin, H.; Voudouris, A. A.; and Wu, X. 2023. Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.*, 322: 103965.
- Amanatidis, G.; Birmipas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2021. Maximum Nash welfare and other stories about EFX. *Theor. Comput. Sci.*, 863: 69–85.
- Amanatidis, G.; Filos-Ratsikas, A.; and Sgouritsa, A. 2024. Pushing the Frontier on Approximate EFX Allocations. *CoRR*, abs/2406.12413.
- Amanatidis, G.; Markakis, E.; and Ntokos, A. 2020. Multiple Birds with One Stone: Beating $1/2$ for EFX and GMMS via Envy Cycle Elimination. 1790–1797. AAAI Press.
- Ashuri, A.; and Gkatzelis, V. 2024. Simultaneously Satisfying MXS and EFL. arXiv:2412.00358.
- Ashuri, A.; Gkatzelis, V.; and Sgouritsa, A. 2024. EF2X Exists For Four Agents. arXiv:2412.00254.
- Barman, S.; Biswas, A.; Murthy, S. K. K.; and Narahari, Y. 2018. Groupwise Maximin Fair Allocation of Indivisible Goods. 917–924. AAAI Press.
- Barman, S.; Kar, D.; and Pathak, S. 2024. Parameterized Guarantees for Almost Envy-Free Allocations. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2024, Auckland, New Zealand, May 6-10, 2024*, 151–159. International Foundation for Autonomous Agents and Multiagent Systems / ACM.
- Bérczi, K.; Bérczi-Kovács, E. R.; Boros, E.; Gedefa, F. T.; Kamiyama, N.; Kavitha, T.; Kobayashi, Y.; and Makino, K. 2024. Envy-free relaxations for goods, chores, and mixed items. *Theor. Comput. Sci.*, 1002: 114596.
- Berendsohn, B. A.; Boyadzhyska, S.; and Kozma, L. 2022. Fixed-Point Cycles and Approximate EFX Allocations. In *47th International Symposium on Mathematical Foundations of Computer Science, MFCS 2022*, volume 241 of *LIPICs*, 17:1–17:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Berger, B.; Cohen, A.; Feldman, M.; and Fiat, A. 2022. Almost Full EFX Exists for Four Agents. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, 4826–4833. AAAI Press.
- Budish, E. 2010. The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. 74:1. ACM.
- Caragiannis, I.; Garg, J.; Rathi, N.; Sharma, E.; and Varricchio, G. 2023. New Fairness Concepts for Allocating Indivisible Items. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China, 2554–2562*. ijcai.org.
- Caragiannis, I.; Gravin, N.; and Huang, X. 2019. Envy-Freeness Up to Any Item with High Nash Welfare: The Virtue of Donating Items. In Karlin, A. R.; Immorlica, N.; and Johari, R., eds., *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019*, 527–545. ACM.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Economics and Comput.*, 7(3): 12:1–12:32.
- Chan, H.; Chen, J.; Li, B.; and Wu, X. 2019. Maximin-Aware Allocations of Indivisible Goods. *CoRR*, abs/1905.09969.
- Chaudhury, B. R.; Garg, J.; and Mehlhorn, K. 2024. EFX Exists for Three Agents. *J. ACM*, 71(1): 4:1–4:27.
- Chaudhury, B. R.; Garg, J.; Mehlhorn, K.; Mehta, R.; and Misra, P. 2023. Improving Envy Freeness up to Any Good Guarantees Through Rainbow Cycle Number. *Math. Oper. Res.*, 0(0).
- Chaudhury, B. R.; Kavitha, T.; Mehlhorn, K.; and Sgouritsa, A. 2021. A Little Charity Guarantees Almost Envy-Freeness. *SIAM J. Comput.*, 50(4): 1336–1358.
- Christodoulou, G.; Fiat, A.; Koutsoupias, E.; and Sgouritsa, A. 2023. Fair allocation in graphs. In *Proceedings of the 24th ACM Conference on Economics and Computation, EC 2023, London, United Kingdom, July 9-12, 2023*, 473–488. ACM.
- Christoforidis, V.; and Santorinaios, C. 2024. On the Pursuit of EFX for Chores: Non-existence and Approximations. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, August 3-9, 2024*, 2713–2721. ijcai.org.
- Farhadi, A.; Hajiaghayi, M. T.; Latifian, M.; Seddighin, M.; and Yami, H. 2021. Almost Envy-freeness, Envy-rank, and Nash Social Welfare Matchings. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021*, 5355–5362. AAAI Press.
- Feldman, M.; Mauras, S.; and Ponitka, T. 2024. On Optimal Tradeoffs between EFX and Nash Welfare. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024*, 9688–9695. AAAI Press.
- Gorantla, P.; Marwaha, K.; and Velusamy, S. 2023. Fair allocation of a multiset of indivisible items. In Bansal, N.; and Nagarajan, V., eds., *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, 304–331. SIAM.

Halpern, D.; Procaccia, A. D.; Psomas, A.; and Shah, N. 2020. Fair Division with Binary Valuations: One Rule to Rule Them All. In *Web and Internet Economics - 16th International Conference, WINE 2020, Beijing, China, December 7-11, 2020, Proceedings*, volume 12495 of *Lecture Notes in Computer Science*, 370–383. Springer.

Hosseini, H.; Sikdar, S.; Vaish, R.; and Xia, L. 2021. Fair and Efficient Allocations under Lexicographic Preferences. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, 5472–5480. AAAI Press.

Hosseini, H.; Sikdar, S.; Vaish, R.; and Xia, L. 2023. Fairly Dividing Mixtures of Goods and Chores under Lexicographic Preferences. In Agmon, N.; An, B.; Ricci, A.; and Yeoh, W., eds., *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 - 2 June 2023*, 152–160. ACM.

Jahan, S. C.; Seddighin, M.; Javadi, S. M. S.; and Sharifi, M. 2023. Rainbow Cycle Number and EFX Allocations: (Almost) Closing the Gap. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023*, 2572–2580. ijcai.org.

Kaviani, A.; Seddighin, M.; and Shahrezaei, A. 2024. Almost Envy-free Allocation of Indivisible Goods: A Tale of Two Valuations. arXiv:2407.05139.

Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings 5th ACM Conference on Electronic Commerce (EC-2004)*, 125–131. ACM.

Mahara, R. 2024. Extension of Additive Valuations to General Valuations on the Existence of EFX. *Math. Oper. Res.*, 49(2): 1263–1277.

Markakis, E.; and Santorinaios, C. 2023. Improved EFX Approximation Guarantees under Ordinal-based Assumptions. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, 591–599. ACM.

Plaut, B.; and Roughgarden, T. 2020. Almost Envy-Freeness with General Valuations. *SIAM J. Discret. Math.*, 34(2): 1039–1068.