

Almost Full EFX Exists for Four Agents*

Ben Berger,¹ Avi Cohen,¹ Michal Feldman,¹ Amos Fiat¹

¹ Tel Aviv University

benberger1@tauex.tau.ac.il, avicohen2@mail.tau.ac.il, michal.feldman@cs.tau.ac.il, fiat@tau.ac.il

Abstract

The existence of EFX allocations of goods is a major open problem in fair division, even for additive valuations. The current state of the art is that no setting where EFX allocations are impossible is known, and yet, existence results are known only for very restricted settings, such as: (i) agents with identical valuations, (ii) 2 agents, and (iii) 3 agents with additive valuations. It is also known that EFX exists if one can leave $n - 1$ items unallocated, where n is the number of agents.

We develop new techniques that allow us to push the boundaries of the enigmatic EFX problem beyond these known results, and (arguably) to simplify proofs of earlier results. Our main result is that every setting with 4 additive agents admits an EFX allocation that leaves at most a single item unallocated. Beyond our main result, we introduce a new class of valuations, termed *nice cancelable*, which includes additive, unit-demand, budget-additive and multiplicative valuations, among others. Using our new techniques, we show that both our results and previous results for additive valuations extend to nice cancelable valuations.

1 Introduction

The question of justness, fairness and division of resources and commitments dates back to Aristotle (Chroust 1942). *Distributional justice*, the “just” allocation of limited resources, is fundamental in the work of (Rawls 1999). Some evidence of the great interest in Rawls’ work is that numerous editions of his book have been cited over 100,000 times.

The mathematical study of fair division is due to Hugo Steinhaus, Bronislaw Knaster and Stefan Banach (Steinhaus 1949) who considered proportional allocations, in which every one of the n agents gets at least a $1/n$ fraction of her total value for all the goods.

A stronger notion of fairness is that of an *envy free* (EF) allocation — introduced by (Gamow and Stern 1958) for cake cutting, and in the context of general resource allocation by (Foley 1967). Unfortunately, if goods are indivisible, envy free allocations need not exist. Consider the trivial case of one indivisible good — if some agent gets the good, others will be envious. Lipton et al. (2004) and Budish (2011) consider a relaxed notion of envy freeness, namely *envy freeness*

up to some item (EF1) — an allocation is EF1 if for every pair of agents Alice and Bob, there is an item that we can remove from Alice’s allocation such that Bob will not be interested in swapping his allocation with what remains of Alice’s allocation.

EF1 allocations always exist but their fairness guarantees are questionable. Consider for example a setting where Alice and Bob have identical valuations over 3 items a, b, c with respective values 1, 1, 2. Arguably, a fair allocation would assign a, b to one of the players, and c to the other one, giving each a value 2. However, the allocation that assigns a, c to Alice and b to Bob is also EF1.

The notion of envy freeness up to any item (EFX) was introduced by (Caragiannis et al. 2016, 2019). An allocation is EFX if for every pair of agents, Alice and Bob, Bob does not want to swap with what remains of Alice’s allocation when *any item is discarded*. *I.e.*, it suffices to consider removing the item with minimal marginal value (to Bob) from Alice’s allocation. Indeed, in the example above, the only EFX allocations are those that allocate a, b to one player and c to the other player.

A major open problem is “when do EFX allocations exist?”. The current state of our knowledge is somewhat embarrassing. We do not know how to rule out EFX allocations in any setting, and yet, they are known to exist only in several restricted cases. In particular, Plaut and Roughgarden (2020) prove that EFX valuations exist for 2 agents with arbitrary valuations, and for any number of agents with identical valuations. Even for the simple case of *additive* valuations (where the value of a bundle of items is simply the sum of values of individual goods), EFX is only known to exist in settings with 3 agents (Chaudhury, Garg, and Mehlhorn 2020), in settings with only one of two types of additive valuations (Mahara 2020), or when the value of every agent to every item can take one of two permissible values (Amanatidis et al. 2021).

Indeed, Procaccia (2020) recently wrote:

In my view, it (EFX existence) is the successor of envy-free cake cutting as fair division’s biggest problem.

Given that EFX valuations are known to exist in so few cases, the following question arises: Can one find a good *partial* EFX allocation? *I.e.*, an EFX allocation in which

*See (Berger et al. 2021) for the full version of this paper.
Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

only a small amount of items can be unallocated? The idea of partial EF and EFX allocations has appeared in multiple papers, e.g. (Brams, Kilgour, and Klamler 2013; Cole, Gkatzelis, and Goel 2013; Caragiannis, Gravin, and Huang 2019). Caragiannis, Gravin, and Huang (2019) show that discarding some items gives good EFX allocations for the rest (achieving 1/2 of the maximum Nash Welfare). Chaudhury et al. (2020) show that given n agents with arbitrary valuations, there always exists an EFX allocation with at most $n - 1$ unallocated items. Moreover, no agent prefers the set of unallocated items to her own allocation.

1.1 Our Results

In this paper we develop new techniques, based upon ideas that appear in (Chaudhury, Garg, and Mehlhorn 2020; Chaudhury et al. 2020). Chaudhury, Garg, and Mehlhorn (2020) introduced the notion of *champion edges* with respect to a single unallocated good, and used it to make progress with respect to the lexicographic potential function in order to eventually reach an EFX allocation. We extend the notion of champion edges beyond a single unallocated item, to *sets* of items, allocated or not, and derive useful structural properties that allow us to make more aggressive progress within a graph theoretic framework.

Our main result concerns EFX allocation for four agents. Extending EFX existence from three to four agents is highly non-trivial. Indeed, Chaudhury et al. (2021) discovered an instance with four additive agents in which there exists an EFX allocation with one unallocated item such that no progress can be made based on the lexicographic potential function. We show that one unallocated item is the only possible obstacle to EFX existence in any setting with four agents.

Theorem 1 (Main Result): Every setting with four additive agents admits an EFX allocation with at most a single unallocated item (which is not envied by any agent).

To prove Theorem 1, we show that for any EFX allocation with at least two unallocated items, it is possible to reshuffle bundles and reallocate them in such a way that advances the lexicographic potential function and preserves EFX. The proof requires solving a complex puzzle, and exemplifies the extensive use of our new techniques.

The immediate open problem is whether one can go the additional mile and allocate the one item that remains. A natural approach to solving this problem is by using a different potential function. Notably, our new techniques are orthogonal to the choice of the potential function, and may prove useful in analyzing other potential functions.

We believe that we have only scratched the surface of the power of our new techniques, and hope they will prove useful in making further progress on the EFX problem.

Further Results: Beyond our main result on 4 agents, in the full version we illustrate the applicability of our framework by extending previous results. In a subsequent paper (Mahara 2021) further improved these results while making use of our techniques and terminology introduced herein.

We show that for every setting with n additive agents there exists an EFX allocation with at most $n - 2$ unallocated items. Additionally, our new techniques greatly sim-

plify existing proofs of EFX existence for 3 agents (Chaudhury, Garg, and Mehlhorn 2020) and for the case of 2 types of additive valuations (Mahara 2020).

Extensions beyond additive valuations appear in the full version. Note that the full version deals with *nice-cancelable* valuations, a generalization of additive valuations.

1.2 Our Techniques

Our proof techniques lie within a graph theoretic framework. Given an EFX allocation \mathbf{X} , we describe a graph $M_{\mathbf{X}}$ (see Definition 3.2) where vertices are associated with agents and there are three types of edges: envy edges $i \rightarrow j$, champion edges $i \xrightarrow{g} j$, where g is an unallocated item, and generalized champion edges $i \xrightarrow{(H|S)} j$, where H is some subset of items (allocated or not) and S is a subset of j 's allocation in \mathbf{X} .

The use of such graphs, with envy and champion edges (but no generalized championship edges) has previously appeared in the literature and is a key component in the proof of an EFX allocation for 3 additive agents (Chaudhury, Garg, and Mehlhorn 2020). The new ingredient introduced in this paper is the notion of generalized champion edges. We show how to find such edges (Section 3.1), and use them to reach a new EFX allocation that advances the lexicographic potential function of (Chaudhury, Garg, and Mehlhorn 2020).

The key idea in all our results is to reshuffle the existing allocation to obtain a new allocation with higher potential, while preserving EFX. This follows the same proof template as in Chaudhury *et al.* — but we have more options to play with by using the generalized championship edges.

An envy edge $i \rightarrow j$ suggests a possible reshuffling where agent i gets j 's current allocation. A champion edge $i \xrightarrow{g} j$ suggests another reshuffling, where agent i gets a subset of agent j 's current allocation, along with the currently unallocated item g . A generalized champion edge $i \xrightarrow{(H|S)} j$ suggests giving agent i agent j 's allocation along with some arbitrary set of items H (that may be arbitrarily allocated among other agents, or be unallocated), while freeing up the set of items S .

Our proofs require solving a complex puzzle, where the goal is to find a cycle consisting of envy, champion, and generalized champion edges, such that the union of all sets S freed up along with the currently unallocated items suffice for the suggested allocation along the cycle.

Finding the appropriate generalized championship edges is a major technical component of our techniques (see Section 3.1). We show how to find such edges, based on existing edges. Then, these edges allow us to reshuffle the current allocation and advance the potential.

Remark 1.1. Due to space limitations some of the technical claims and lemmas in the body of the paper are presented without proofs. These can be found in the full version of the paper.

2 Preliminaries

We consider a setting with n agents, and a set M of m items. Each agent has a valuation $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$, which is normalized and monotone, i.e. $v(S) \leq v(T)$ whenever $S \subseteq T$

and $v(\emptyset) = 0$. We further assume that all valuations are additive, i.e., $v_i(S) = \sum_{s \in S} v_i(\{s\})$.

For two sets of items $S, T \subseteq M$, we write $S <_i T$ if $v_i(S) < v_i(T)$. Similarly we define $S >_i T$, $S \leq_i T$, $S \geq_i T$, $S =_i T$ if $v_i(S) > v_i(T)$, $v_i(S) \leq v_i(T)$, $v_i(S) \geq v_i(T)$, $v_i(S) = v_i(T)$, respectively.

We denote a valuation profile by $\mathbf{v} = (v_1, \dots, v_n)$. An allocation is a vector $\mathbf{X} = (X_1, \dots, X_n)$ of disjoint bundles, where X_i is the bundle allocated to agent i . Given an allocation \mathbf{X} , We say that agent i envies a set of items S if $X_i <_i S$. We say that agent i envies agent j , denoted $i \rightarrow j$, if i envies X_j . We say that agent i strongly envies a set of items S if there exists some $h \in S$ such that i envies $S \setminus \{h\}$. Likewise we say that agent i strongly envies agent j if i strongly envies X_j . \mathbf{X} is called *envy-free* (EF) if no agent envies another. \mathbf{X} is called *envy-free up to any good* (EFX) if no agent strongly envies another.

A valuation v is *non-degenerate* if $v(S) \neq v(T)$ for any two different bundles S, T . (Chaudhury, Garg, and Mehlhorn 2020) have shown that in order to prove the existence of an EFX allocation for a given valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ of additive valuations, it is without loss of generality to assume that all of the valuations are non-degenerate. Thus, for the remainder of this paper we assume that all valuations are non-degenerate.

Potential Functions and Progress Measures: All our EFX existence results follow the same paradigm: given an arbitrary EFX allocation \mathbf{X} with k unallocated goods, construct a new partial EFX allocation that advances some fixed potential function. Since there are finitely many allocations, there must exist an EFX allocation with at most $k - 1$ unallocated items.

A natural progress measure to consider is Pareto domination. Given two allocations \mathbf{X}, \mathbf{Y} , we say that \mathbf{Y} *Pareto dominates* \mathbf{X} if $Y_i \geq X_i$ for every $i \in [n]$, and there exists some i for which the inequality is strict. Chaudhury, Garg, and Mehlhorn (2020) have presented an instance with $n = 3$ agents that admits a partial EFX allocation which is not Pareto-dominated by any full EFX allocation (note that there can still exist Pareto-optimal full EFX allocations). To overcome this obstacle they introduced an alternative “lexicographic” progress measure which we shall also use:

Definition 2.1 (Chaudhury, Garg, and Mehlhorn (2020)). Fix some arbitrary ordering of the agents a_1, \dots, a_n . The allocation \mathbf{Y} *dominates* \mathbf{X} if for some $k \in [n]$, we have that $Y_{a_j} = X_{a_j}$ for all $1 \leq j < k$, and $Y_{a_k} > X_{a_k}$.

Note that Pareto-domination implies domination but not vice versa.

Lemma 2.2. If for every partial EFX allocation \mathbf{X} with k unallocated items, there exists a partial EFX allocation \mathbf{Y} that dominates \mathbf{X} , then there exists a partial EFX allocation with at most $k - 1$ unallocated items. Moreover, no agent envies the set of $k - 1$ unallocated items.

Hereinafter we fix a partial EFX allocation \mathbf{X} , and our goal is to find a dominating EFX allocation \mathbf{Y} . In fact, we almost always progress via Pareto-domination. In the few cases we do not, we find an allocation in which a_1 (the most

important agent in the ordering) is strictly better off. We denote this agent a_{vip} .

Most Envious Agents: Fix some unallocated good g . We denote by U the set of goods that are unallocated in \mathbf{X} (thus $g \in U$). The following are variants of definitions from (Chaudhury, Garg, and Mehlhorn 2020; Chaudhury et al. 2020).

We say that i is *most envious* of a set of items S , if there exists a subset $T \subseteq S$, such that i envies T and no agent strongly envies T . When more than one such T exists, we choose one of them arbitrarily unless stated otherwise. The set $S \setminus T$ is referred to as the corresponding *discard set*.

Definition 2.3 (Chaudhury, Garg, and Mehlhorn (2020)). We say that i *champions* j with respect to g , denoted $i \xrightarrow{g} j$, if i is most envious of $X_j \cup \{g\}$. The corresponding discard set is denoted $D_{i,j}^g$. Note that i envies the set $(X_j \cup \{g\}) \setminus D_{i,j}^g$, but no agent strongly envies it.

An important case considered frequently in the paper is where $g \notin D_{i,j}^g$. In this case $X_j = (X_j \setminus D_{i,j}^g) \cup D_{i,j}^g$. Following (Chaudhury, Garg, and Mehlhorn 2020), if $i \xrightarrow{g} j$ and $g \notin D_{i,j}^g$, then we say that i *g-decomposes* j into *top* and *bottom half-bundles* $(X_j \setminus D_{i,j}^g)$ and $D_{i,j}^g$, respectively (in short, i *g-decomposes* j). If there is no concern of ambiguity, then we denote the top and bottom half-bundles by T_j and B_j , respectively (note that different g -decomposers of j may induce different top and bottom half-bundles). Under this notation, we have $(X_j \cup \{g\}) \setminus D_{i,j}^g = T_j \cup \{g\}$.

In the following observations from (Chaudhury, Garg, and Mehlhorn 2020), $i \xrightarrow{g} j$ and $i \not\xrightarrow{g} j$ are the respective negations of $i \xrightarrow{g} j$, $i \rightarrow j$.

Observation 2.4. For every agent i , there exists an agent j who champions i with respect to g .

Observation 2.5. If $i \xrightarrow{g} j$ and $i \not\xrightarrow{g} j$, then $g \notin D_{i,j}^g$, i.e., i *g-decomposes* j .

3 Generalized Championship

A crucial component in our techniques is the extension of Definition 2.3 to an arbitrary set of items H . It will be useful to have a notation that contains some information regarding the discarded items.

Definition 3.1. i *champions* j with respect to $(H \mid S)$, denoted $i \xrightarrow{(H \mid S)} j$, where $H \subseteq M \setminus X_j$ and $S \subseteq X_j$, if i is most envious of $(X_j \setminus S) \cup H$. The corresponding discard set is denoted $D_{i,j}^{H \mid S}$.

As opposed to basic championship, not every agent j has an $(H \mid S)$ -champion (consider an extreme example where $H = \emptyset, S = X_j$). If $i \xrightarrow{(H \mid S)} j$, then giving i the desired bundle implied by the championship releases S to be reallocated to other agents. For example, if we also know that $k \xrightarrow{(S \mid S')} \ell$, then these two champion relations can be “used” simultaneously in a transition to a new EFX allocation.

We say that a set of items T is *released* by $i \xrightarrow{(H \mid S)} j$ if $T \subseteq S \cup D_{i,j}^{H \mid S}$. We denote the negation of $i \xrightarrow{(H \mid S)} j$ by $i \not\xrightarrow{(H \mid S)} j$.

Definition 3.2. The *champion graph* $M_{\mathbf{X}} = ([n], E)$ with respect to \mathbf{X} is a labeled directed multi-graph. The vertices correspond to the agents, and E consists of the following 3 types of edges:

1. Envy edges: $i \rightarrow j$ iff i envies j .
2. Champion edges: $i \xrightarrow{-g} j$ iff i champions j with respect to g , where g is an unallocated good.
3. Generalized champion edges: $i \xrightarrow{-(H|S)} j$ iff i champions j with respect to $H | S$.

We refer to envy and champion edges as *basic edges*. Hereinafter, the edge notations above will sometimes refer to the edges of the champion graph and sometimes refer to the statements they convey. For example, we will sometimes refer to “ $i \xrightarrow{-g} j$ ” as an edge in $M_{\mathbf{X}}$ and sometimes as shorthand that i is a g -champion of j , and the meaning will be clear from the context. Furthermore, it is not hard to verify that $i \xrightarrow{-g} j$ iff $i \xrightarrow{-\{g\}|\emptyset} j$ and that $i \rightarrow j$ iff $i \xrightarrow{-\emptyset|\emptyset} j$. Thus, we can treat every basic edge in $M_{\mathbf{X}}$ as a generalized champion edge.

Given a cycle $C = a_1 \rightarrow a_2 \rightarrow \dots \rightarrow a_k \rightarrow a_1$ of edges and an agent a_i in the cycle, $\text{succ}(a_i)$ and $\text{pred}(a_i)$ denote, respectively, the successor and predecessor of a_i along the cycle.

Definition 3.3. A cycle $C = a_1 \xrightarrow{-(H_1|S_1)} a_2 \xrightarrow{-(H_2|S_2)} \dots \xrightarrow{-(H_{k-1}|S_{k-1})} a_k \xrightarrow{-(H_k|S_k)} a_1$ in $M_{\mathbf{X}}$ is called *Pareto improvable* (PI) if for every $i, j \in [k]$ we have $H_i \cap H_j = \emptyset$, and either $H_i \subseteq U$ or H_i is released by some edge in the cycle¹.

A PI cycle which is composed entirely of basic edges is called a *basic PI cycle*.

By definition, every agent a_i along a PI cycle envies some subset $A_i \subseteq (X_{\text{succ}(a_i)} \setminus S_i) \cup H_i$ that no agent strongly envies. The following simple but useful lemma asserts that reallocating A_i to agent a_i for every a_i along the cycle produces a Pareto-dominating EFX allocation. Thus, finding a PI cycle in $M_{\mathbf{X}}$ suffices to advance our potential function.

Lemma 3.4. If $M_{\mathbf{X}}$ contains a Pareto improvable cycle, then there exists a (partial) EFX allocation \mathbf{Y} that Pareto dominates \mathbf{X} . Furthermore, every agent i along the cycle satisfies $Y_i >_i X_i$.

Corollary 3.5 (Following (Chaudhury et al. 2020)). If $M_{\mathbf{X}}$ contains an envy-cycle, a self- g -champion (an agent i satisfying $i \xrightarrow{-g} i$) or a cycle composed of envy edges and basic champion edges where for every $h \in U$ there is at most one h -champion edge in the cycle, then there exists a (partial) EFX allocation \mathbf{Y} that Pareto dominates \mathbf{X} . Note that these are exactly the basic PI cycles².

¹We could have defined a PI cycle more generally, e.g., to allow the set H_i to be a combination of unallocated goods and items released from several edges. The proposed definition is hopefully easier to digest and suffices for our purposes.

²Envy cycles, the simplest form of basic PI cycles, were considered in (Lipton et al. 2004). Basic PI-cycles were considered in (Chaudhury et al. 2020) using different terminology — championship was only defined in (Chaudhury, Garg, and Mehlhorn 2020). Our definition of a PI-cycle captures and generalizes these notions.

Remark 3.6. Lemma 3.4 can be generalized to handle disjoint cycles. The fact that C is a cycle is used in the proof of the lemma only to ensure that every agent whose bundle is reallocated, is also given an alternative bundle in the new allocation. The same is true if C is a set of vertex-disjoint cycles rather than a single cycle. We may then define C as an edge set $\{a_i \xrightarrow{-(H_i|S_i)} \text{succ}(a_i)\}_{i \in [k]}$, and if the conditions in the definition of a Pareto-improvable cycle are satisfied, then Lemma 3.4 still applies. In this case we refer to C as a *Pareto-improvable edge set*.

3.1 New Edge Discovery

In this section we describe a way to discover new generalized champion edges in $M_{\mathbf{X}}$. These will almost always be of the form $k \xrightarrow{-(S|B_j)} j$ where $B_j \subseteq X_j$ is some bottom half-bundle induced by a g -decomposer of j (see discussion below Definition 2.3). Therefore, to facilitate readability we use the following convention:

Convention 3.7. For any two agents k, j and any set S disjoint from X_j , we write $k \xrightarrow{-(S|\circ)} j$ as shorthand for $k \xrightarrow{-(S|B_j)} j$, where the half-bundle B_j will be clear from the context.

The following structure within the champion-graph is especially convenient for edge discovery.

Definition 3.8. A cycle $C = a_1 \xrightarrow{-g} a_2 \xrightarrow{-g} \dots \xrightarrow{-g} a_k \xrightarrow{-g} a_1$ with at least two g -champion edges in $M_{\mathbf{X}}$ is called a *good g -cycle* if:

1. All agents along the cycle are different.
2. There are no parallel envy edges, i.e., $a_i \not\rightarrow \text{succ}(a_i)$ for all i .
3. There are no internal g -champion edges, i.e., for every $i, j \in [k]$, $a_i \xrightarrow{-g} a_j$ iff $a_j = \text{succ}(a_i)$.

Observation 3.9. Agents j on a good g -cycle are g -decomposed by $\text{pred}(j)$ into $X_j = T_j \cup B_j$.

We next show how to discover new generalized champion edges in the presence of a good g -cycle. The following two observations are useful:

Observation 3.10. If $i \not\rightarrow j$ then $i \xrightarrow{-(B_j|\circ)} j$.

Observation 3.11. For any two agents i, j along a good g -cycle, $\text{pred}(i) \xrightarrow{-(B_i|\circ)} j$.

Lemma 3.12. Let C be a good g -cycle. For any agent i along the cycle, there exists an agent a such that $a \xrightarrow{-(B_i|\circ)} \text{succ}(i)$.

Lemma 3.13. Let C be a good g -cycle. Let i, j, k be agents along the cycle. If $k \xrightarrow{-(B_i|\circ)} j$, then there exists an agent a (not necessarily in the cycle) such that $a \xrightarrow{-(B_i|\circ)} \text{succ}(k)$.


For every bottom half-bundle B_i along a good g -cycle C , applying Lemma 3.12 provides an initial B_i -champion edge. If this edge is internal to the cycle, i.e., the source of the edge is in the cycle, then we can apply Lemma 3.13 to discover a new B_i -champion edge. Once again, if the new edge is internal to the cycle, then we can reapply Lemma 3.13. We can repeat this process to discover more and more B_i -champion edges, until either the new edge has already been previously

discovered, or it is external (i.e., its source is outside the cycle).

There are two particular types of internal B_i -champions edges that are useful to us.

Definition 3.14. Let C be a good g -cycle. Let i, j, k be three agents along C . If $i \xrightarrow{B_k | \circ} j$ and k is on the path from j to i in C , then we say that the edge $i \xrightarrow{B_k | \circ} j$ is a *good edge* (or good B_k -edge). If $\ell \xrightarrow{B_k | \circ} j$ for some agent ℓ outside the cycle C , then we say that the edge $\ell \xrightarrow{B_k | \circ} j$ is an *external edge* (or external B_k -edge).

The figure below illustrates Definition 3.14.



The red edges form a good g -cycle C among 4 agents, $C = 1 \xrightarrow{g} 2 \xrightarrow{g} 3 \xrightarrow{g} 4 \xrightarrow{g} 1$. The edge $2 \xrightarrow{B_1 | B_4} 4$ is a *good edge*, since 1 is on the path from 4 to 2 in C . On the other hand, $3 \xrightarrow{B_4 | B_2} 2$ is not a good edge (we call it a *bad edge* in the figure), since 4 is not on the path from 2 to 3 along C . $5 \xrightarrow{B_2 | B_3} 3$ is an external edge.

Theorem 3.15. Let C be a good g -cycle. For every agent j along the cycle, there exists either a good B_j -edge in C , or an external B_j -edge in C .


Proof. Assume without loss of generality that $C = 1 \xrightarrow{g} 2 \xrightarrow{g} \dots \xrightarrow{g} k \xrightarrow{g} 1$ and $j = 1$, i.e., we try to find B_1 -champion edges. By Lemma 3.12 there exists an edge $\ell_1 \xrightarrow{B_1 | \circ} 2$ for some agent ℓ_1 . If this is an external B_1 -edge we are done. Otherwise, ℓ_1 is an agent along C , and thus by Lemma 3.13 there exists an edge $\ell_2 \xrightarrow{B_1 | \circ} \text{succ}(\ell_1)$, for some agent ℓ_2 which can be equal to ℓ_1 . As long as the result of Lemma 3.13 is not an external edge we may apply the lemma repeatedly. Hence, if no such iteration results in an external edge, we obtain a sequence of agents $(\ell_i)_{i=1}^{\infty}$ such that $\ell_{i+1} \xrightarrow{B_1 | \circ} \text{succ}(\ell_i)$ for every $i \geq 1$.

If for some $i \geq 1$, we have $\ell_{i+1} \leq \ell_i$ then the edge $\ell_{i+1} \xrightarrow{B_1 | \circ} \text{succ}(\ell_i)$ is a good edge (since the path from $\text{succ}(\ell_i)$ to ℓ_{i+1} includes 1). Hence, if none of these edges are good, then $\ell_i < \ell_{i+1}$ for every $i \geq 1$, in contradiction to C being of finite length. Thus, one of these edges must be good, hence we are done. \square

The following observation and its corollary allow us to narrow down the possible configurations of B_j -edges obtained from Theorem 3.15.

Observation 3.16. If $i \xrightarrow{B_j | \circ} k$ and $i \not\rightarrow k$ then $B_k < B_j$.

Corollary 3.17. Let C be a good g -cycle. Consider the set of B_j -edges guaranteed by Theorem 3.15 for every agent j along the cycle. If all these edges are external, then they cannot all share the same source agent, unless that agent envies some agent along the cycle. (the figure below demonstrates an impossible configuration).



4 Proof of the Main Result: EFX for 4 additive agents with 1 unallocated good

In this section we prove our main result, namely that every setting with 4 additive agents admits an EFX allocation with at most one unallocated good. By Lemma 2.2 it suffices to prove:

Theorem 4.1. Let \mathbf{X} be an EFX allocation on 4 agents with additive valuations, with at least two unallocated items. Then, there exists an EFX allocation \mathbf{Y} that dominates \mathbf{X} .

The proof involves a rigorous case analysis, which exemplifies the extensive use of our new techniques. We have attempted to make the proofs as accessible as possible through the use of extensive aids such as figures and colors.

By assumption, there exist two unallocated goods which we denote g, h . The proof distinguishes between two main cases, namely whether \mathbf{X} is envy-free or not. When \mathbf{X} is envy-free, we show that a Pareto improvable (PI) cycle always exists. This is shown via a case analysis that depends on the lengths of the good g - and h -cycles which (we show) must exist in the champion graph $M_{\mathbf{X}}$.

When \mathbf{X} has envy, we show that one can restrict attention to cases where the basic edges in $M_{\mathbf{X}}$ follow some specific structure, modulo permuting the agent identities. Then, we show that there is an EFX allocation in which agent a_{vip} (per the lexicographic potential) is better off relative to \mathbf{X} . Since a_{vip} could be any one of the agents (due to the identity permutation), the proof deals with all cases irrespective of the identity of a_{vip} . Our approach here is heavily inspired by and follows a similar high-level structure to that of (Chaudhury, Garg, and Mehlhorn 2020) in their analysis of the envy case in their 3 agent result. Due to space constraints the complete proof of Theorem 4.1 is deferred to the full version. In what follows we present some of the cases where \mathbf{X} is envy-free.

4.1 \mathbf{X} is Envy-Free

We show that if \mathbf{X} is envy-free, then we can always find a PI cycle or edge set in $M_{\mathbf{X}}$ (see Remark 3.6), implying (by Lemma 3.4) the existence of a Pareto-dominating EFX allocation \mathbf{Y} .

Recall that every agent i has an incoming g -champion edge and an incoming h -champion edge (Observation 2.4), and thus $M_{\mathbf{X}}$ contains a g -cycle and an h -cycle. If there is a self g or h champion we are done by Corollary 3.5. Thus, these cycles are of size at least 2, and contain no envy edges, and are therefore good cycles. Denote them by C_g, C_h .


By Observation 3.9, C_g (resp. C_h) induces a g (resp. h)-decomposition of X_j for any agent j in the cycle. In the following we denote the g (resp. h)-decomposition by $X_j = T_j^g \cup B_j^g$ (resp. $X_j = T_j^h \cup B_j^h$). We shall make repeated use of Observations 3.10 and 3.11. For concise presentation, we write here the implications that will be repeatedly used in this section: For every two agents i, j we have:

(a) $i \xrightarrow{B_j^g} j$; (b) if i, j reside on the same good g -cycle, then $\text{pred}(i) \xrightarrow{B_j^g} j$. Analogous claims hold for h . We remind the reader that $i \xrightarrow{B_k^g} j, i \xrightarrow{B_k^h} j$ are shorthand for $i \xrightarrow{B_k^g | B_j^g} j, i \xrightarrow{B_k^h | B_j^h} j$, respectively.


Since there are 4 agents, C_g and C_h can be of size 2, 3, or 4. Assume w.l.o.g. that $|C_g| \leq |C_h|$. Thus, there are six cases to consider. The cases $|C_g| = 2, |C_h| = 4; |C_g| = |C_h| = 3; |C_g| = |C_h| = 4$ appear in the full version. The remaining cases are treated below.

Case 1: $|C_g| = |C_h| = 2$. Assume w.l.o.g. that $C_g = 1 \xrightarrow{g} 2 \xrightarrow{g} 1$. By Theorem 3.15, there exists either a good or external B_1^g -edge going into agent 2, and there exists either a good or external B_2^g -edge going into agent 1³. If one of these is good we are done: for example, the only possible good B_1^g -edge is $1 \xrightarrow{B_1^g} 2$ which closes the PI cycle $1 \xrightarrow{B_1^g} 2 \xrightarrow{g} 1$ (recall that the edge $2 \xrightarrow{g} 1$ releases B_1^g).

Thus both edges have to be external, i.e., their sources are agents 3 or 4. Assume w.l.o.g. that $3 \xrightarrow{B_1^g} 2$. We cannot also have $3 \xrightarrow{B_2^g} 1$ by Corollary 3.17. We conclude that $4 \xrightarrow{B_2^g} 1$, and we have the following structure:



Consider C_h . If $C_h = 1 \xrightarrow{h} 2 \xrightarrow{h} 1$ then we are done since we get the PI cycle $1 \xrightarrow{g} 2 \xrightarrow{h} 1$ (see Corollary 3.5). If $C_h = 3 \xrightarrow{h} 4 \xrightarrow{h} 3$, then following the analogous reasoning for C_g we can assume that we have external B_3^h and B_4^h edges, in which case we have one of the following two structures (highlighted edges are part of PI cycles or PI edge sets):




In the left graph, the two cycles $1 \xrightarrow{B_3^h} 4 \xrightarrow{B_2^g} 1$ and $2 \xrightarrow{B_4^h} 3 \xrightarrow{B_1^g} 2$ form a PI edge set, and in the right graph we have the PI-cycle $1 \xrightarrow{B_4^h} 3 \xrightarrow{B_1^g} 2 \xrightarrow{B_3^h} 4 \xrightarrow{B_2^g} 1$ (note that in both cases every one of the four used bottom half-bundles is released by a corresponding incoming edge to each one of the agents), and thus we are done. We remark that if M_X contains the good g -cycle $3 \xrightarrow{g} 4 \xrightarrow{g} 3$, then similar reasoning also shows that we have a PI cycle (or PI edge set). In other words, if M_X contains two disjoint g -cycles of size 2 or two disjoint h -cycles of size 2, then we are done.


It remains to consider the case where C_h intersects C_g at exactly one agent. If $C_h = 1 \xrightarrow{h} 3 \xrightarrow{h} 1$, then we get the PI cycle $1 \xrightarrow{h} 3 \xrightarrow{B_1^g} 2 \xrightarrow{g} 1$. The case $C_h = 2 \xrightarrow{h} 4 \xrightarrow{h} 2$ is symmetric. We are left with the cases $C_h = 2 \xrightarrow{h} 3 \xrightarrow{h} 2$ or $C_h = 1 \xrightarrow{h} 4 \xrightarrow{h} 1$, which

³Note that by Observation 3.10, there cannot be a B_1^g edge going into agent 1 or a B_2^g edge going into agent 2.

are also symmetric. Thus, w.l.o.g. we assume $C_h = 2 \xrightarrow{h} 3 \xrightarrow{h} 2$ and we get the structure:




As before, there must be two external B_2^h and B_3^h edges (coming out of agents 1 and 4). If $1 \xrightarrow{B_2^h} 3$ and $4 \xrightarrow{B_3^h} 2$, then we get the PI cycle $1 \xrightarrow{B_2^h} 3 \xrightarrow{h} 2 \xrightarrow{g} 1$. Thus $1 \xrightarrow{B_3^h} 2$ and $4 \xrightarrow{B_2^h} 3$, and we get the structure:




We now ask who is an h -champion of agent 4 (such exists by Observation 2.4). If $3 \xrightarrow{h} 4$, we are done via the PI edge set $1 \xrightarrow{B_3^h} 2 \xrightarrow{g} 1, 3 \xrightarrow{h} 4 \xrightarrow{B_2^h} 3$. If $2 \xrightarrow{h} 4$, we are done: $1 \xrightarrow{g} 2 \xrightarrow{h} 4 \xrightarrow{B_2^g} 1$. Thus, assume $1 \xrightarrow{h} 4$ (recall our assumption that there are no self-champions, i.e., $4 \not\xrightarrow{h} 4$).

Now we ask who is an h -champion of agent 1. If $4 \xrightarrow{h} 1$ we are done since we have two disjoint size 2 h -cycles, a situation we have already dealt with at the start of Case 1. If $2 \xrightarrow{h} 1$ we are done: $1 \xrightarrow{g} 2 \xrightarrow{h} 1$. Therefore, $3 \xrightarrow{h} 1$ and we have the structure:




Now we ask who is a g -champion of 3. If $1 \xrightarrow{g} 3$ or $2 \xrightarrow{g} 3$, we have a size 2 cycle with g and h edges and we are done. Thus, $4 \xrightarrow{g} 3$.

Finally, we ask who is a g -champion of 4. If $3 \xrightarrow{g} 4$ then we have two disjoint g -cycles of size 2 and we are done. If $2 \xrightarrow{g} 4$, then we are done: $2 \xrightarrow{g} 4 \xrightarrow{B_2^h} 3 \xrightarrow{h} 2$. Thus assume $1 \xrightarrow{g} 4$, and we have the structure:




In this case we are done via the PI edge set $1 \xrightarrow{g} 4 \xrightarrow{B_2^g} 1, 2 \xrightarrow{h} 3 \xrightarrow{B_1^g} 2$.

Case 2: $|C_g| = 2, |C_h| = 3$. Assume w.l.o.g. that $C_g = 1 \xrightarrow{g} 2 \xrightarrow{g} 1$. If C_h passes through both agents 1 and 2 then we are done since we are guaranteed to have a size 2 cycle with g and h edges. Thus C_h passes through exactly one of them, and we can assume w.l.o.g. that $C_h = 1 \xrightarrow{h} 4 \xrightarrow{h} 3 \xrightarrow{h} 1$ (note that the reverse direction of the cycle is symmetric by switching the roles of agents 3 and 4). We get the following structure:




As in the previous case we may assume that the B_1^g and B_2^g edges guaranteed by Theorem 3.15 are external. If we have $3 \xrightarrow{B_2^g} 1$ and $4 \xrightarrow{B_1^g} 2$, we are done: $1 \xrightarrow{h} 4 \xrightarrow{B_1^g} 2 \xrightarrow{g} 1$. Thus assume we have the edges $3 \xrightarrow{B_1^g} 2$ and $4 \xrightarrow{B_2^g} 1$.

We now ask who is a g -champion of 3. If $1 \xrightarrow{g} 3$, we are done: $1 \xrightarrow{g} 3 \xrightarrow{h} 1$. If $2 \xrightarrow{g} 3$, we are done: $1 \xrightarrow{h} 4 \xrightarrow{B_2^g} 1$, $2 \xrightarrow{g} 3 \xrightarrow{B_1^g} 2$. Thus $4 \xrightarrow{g} 3$, and we have the structure:




Consider C_h . We ask which agent i satisfies $i \xrightarrow{B_4^h} 3$ (such exists by Lemma 3.12, since $3 = \text{succ}(1)$ in C_h). We cannot have $1 \xrightarrow{B_4^h} 3$ since $1 = \text{pred}(4)$ in C_h (see Observation 3.11). If $4 \xrightarrow{B_4^h} 3$, we are done: $1 \xrightarrow{h} 4 \xrightarrow{B_4^h} 3 \xrightarrow{B_1^g} 2 \xrightarrow{g} 1$. If $2 \xrightarrow{B_4^h} 3$ we are done: $1 \xrightarrow{h} 4 \xrightarrow{B_2^g} 1$, $2 \xrightarrow{B_4^h} 3 \xrightarrow{B_1^g} 2$. Thus, we must have $3 \xrightarrow{B_4^h} 1$.⁴

We now ask which agent i satisfies $i \xrightarrow{B_4^h} 1$ (such exists by Lemma 3.13 since $3 \xrightarrow{B_4^h} 1$ and $1 = \text{succ}(4)$ in C_h). We cannot have $1 \xrightarrow{B_4^h} 1$, since $1 = \text{pred}(4)$ in C_h . If $3 \xrightarrow{B_4^h} 1$, we are done: $1 \xrightarrow{h} 4 \xrightarrow{g} 3 \xrightarrow{B_4^h} 1$. If $4 \xrightarrow{B_4^h} 1$, we are done: $1 \xrightarrow{h} 4 \xrightarrow{B_4^h} 1$. Thus $2 \xrightarrow{B_4^h} 1$, and we have the structure:



Finally, we ask which agent i satisfies $i \xrightarrow{B_1^h} 4$ (such exists by Lemma 3.12). We cannot have $2 \xrightarrow{B_1^h} 4$, as otherwise, together with $2 \xrightarrow{B_4^h} 1$ we have by Observation 3.16 that $B_1^h <_2 B_4^h <_2 B_1^h$, contradiction. We cannot have $3 \xrightarrow{B_1^h} 4$, as $3 = \text{pred}(1)$ in C_h . Thus, we must have $4 \xrightarrow{B_1^h} 4$, and we are done: $1 \xrightarrow{g} 2 \xrightarrow{B_4^h} 1$, $4 \xrightarrow{B_1^h} 4$.

Case 3: $|C_g| = 3$, $|C_h| = 4$. (Note that this is Case 5 in the full version.) Assume w.l.o.g. that $C_h = 1 \xrightarrow{h} 2 \xrightarrow{h} 3 \xrightarrow{h} 4 \xrightarrow{h} 1$ and that $C_g = 1 \xrightarrow{g} 3 \xrightarrow{g} 4 \xrightarrow{g} 1$ (note that if C_g is in the opposite direction we immediately get a PI cycle). We have the structure:




⁴Note that as opposed to a self g -loop, $3 \xrightarrow{B_4^h} 3$ is not a PI-cycle since B_4^h is not released within the cycle.

Consider C_g . In what follows we reason about possible B_1^g, B_3^g and B_4^g edges, starting with B_3^g . By Theorem 3.15, M_X contains a good or external B_3^g -edge. If it is a good B_3^g -edge, then it is one of $3 \xrightarrow{B_3^g} 4$, $3 \xrightarrow{B_3^g} 1$, $4 \xrightarrow{B_3^g} 1$, and in all cases we get a PI cycle:


$$1 \xrightarrow{g} 3 \xrightarrow{B_3^g} 4 \xrightarrow{h} 1, 1 \xrightarrow{g} 3 \xrightarrow{B_3^g} 1, \\ 1 \xrightarrow{g} 3 \xrightarrow{h} 4 \xrightarrow{B_3^g} 1,$$

respectively. Thus, we can assume that there is an external B_3^g -edge, which can be either $2 \xrightarrow{B_3^g} 1$ or $2 \xrightarrow{B_3^g} 4$, and thus the structure is one of the following:




We now ask which agent i satisfies $i \xrightarrow{B_1^g} 3$ (such exists by Lemma 3.12, since $3 = \text{succ}(1)$ in C_g). We cannot have $4 \xrightarrow{B_1^g} 3$, since $4 = \text{pred}(1)$ in C_g . If $1 \xrightarrow{B_1^g} 3$, we are done: $1 \xrightarrow{B_1^g} 3 \xrightarrow{h} 4 \xrightarrow{g} 1$. If $3 \xrightarrow{B_1^g} 3$, we are done regardless of whether $2 \xrightarrow{B_3^g} 1$ or $2 \xrightarrow{B_3^g} 4$; in the first case we have the PI edge set $1 \xrightarrow{h} 2 \xrightarrow{B_3^g} 1$, $3 \xrightarrow{B_1^g} 3$, and in the second case we have the PI edge set $1 \xrightarrow{h} 2 \xrightarrow{B_3^g} 4 \xrightarrow{g} 1$, $3 \xrightarrow{B_1^g} 3$. Thus, we must have $2 \xrightarrow{B_1^g} 3$.

Therefore, the external B_3^g -edge cannot be $2 \xrightarrow{B_3^g} 1$, as otherwise by Observation 3.16 we get $B_1^g <_2 B_3^g <_2 B_1^g$, contradiction. Hence we have $2 \xrightarrow{B_3^g} 4$, and we obtain the following structure:



We now ask which agent i satisfies $i \xrightarrow{B_4^g} 1$ (such exists by Lemma 3.12, since $1 = \text{succ}(4)$ in C_g). We cannot have $3 \xrightarrow{B_4^g} 1$, since $3 = \text{pred}(4)$ in C_g . We cannot have $2 \xrightarrow{B_4^g} 1$, as otherwise we get a contradiction to Corollary 3.17, since there are already external B_1^g and B_3^g edges whose source is agent 2. If $4 \xrightarrow{B_4^g} 1$, we are done: $1 \xrightarrow{h} 2 \xrightarrow{B_1^g} 3 \xrightarrow{g} 4 \xrightarrow{B_4^g} 1$. Thus we must have $1 \xrightarrow{B_4^g} 1$ and we have the structure:



By Lemma 3.13, $1 \xrightarrow{B_4^g} 1$ implies that there is an agent i that satisfies $i \xrightarrow{B_4^g} 3$ ($3 = \text{succ}(1)$ in C_g). We cannot have $2 \xrightarrow{B_4^g} 3$, as again we would get a contradiction to Corollary 3.17. We cannot have $3 \xrightarrow{B_4^g} 3$ since $3 = \text{pred}(4)$ in C_g . If $1 \xrightarrow{B_4^g} 3$ we are done: $1 \xrightarrow{B_4^g} 3 \xrightarrow{g} 4 \xrightarrow{h} 1$. Thus we must have $4 \xrightarrow{B_4^g} 3$, and we are done: $3 \xrightarrow{g} 4 \xrightarrow{B_4^g} 3$.

Acknowledgment

We are very grateful to Karin Egri for her great help in producing the figures used in this paper.

This research was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 866132), by the Israel Science Foundation (grant number 317/17), and by the NSF-BSF (grant number 2020788).

References

- Amanatidis, G.; Bimpas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2021. Maximum Nash welfare and other stories about EFX. *Theoret. Comput. Sci.*, 863: 69–85.
- Berger, B.; Cohen, A.; Feldman, M.; and Fiat, A. 2021. Almost Full EFX Exists for Four Agents. <http://arxiv.org/abs/2102.10654>, arXiv:2102.10654.
- Brams, S. J.; Kilgour, D. M.; and Klamler, C. 2013. Two-Person Fair Division of Indivisible Items: An Efficient, Envy-Free Algorithm. MPRA Paper 47400, University Library of Munich, Germany.
- Budish, E. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy*, 119(6): 1061–1103.
- Caragiannis, I.; Gravin, N.; and Huang, X. 2019. Envy-freeness up to any item with high Nash welfare: The virtue of donating items. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 527–545.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, 305–322. New York, NY, USA: Association for Computing Machinery. ISBN 9781450339360.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3): 1–32.
- Chaudhury, B. R.; Garg, J.; and Mehlhorn, K. 2020. EFX exists for three agents. In *Proceedings of the 21st ACM Conference on Economics and Computation*, 1–19.
- Chaudhury, B. R.; Garg, J.; Mehlhorn, K.; Mehta, R.; and Misra, P. 2021. Improving EFX Guarantees through Rainbow Cycle Number. In Biró, P.; Chawla, S.; and Echenique, F., eds., *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, 310–311. ACM.
- Chaudhury, B. R.; Kavitha, T.; Mehlhorn, K.; and Sgouritsa, A. 2020. A little charity guarantees almost envy-freeness. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2658–2672. SIAM.
- Chroust, A.-H. 1942. Aristotle's Conception of Equity (Epieikeia). *Notre Dame Law Review*, 18: 119–128.
- Cole, R.; Gkatzelis, V.; and Goel, G. 2013. Mechanism Design for Fair Division: Allocating Divisible Items without Payments. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, 251–268. New York, NY, USA: Association for Computing Machinery. ISBN 9781450319621.
- Foley, D. K. 1967. Resource allocation and the public sector. *Yale economic essays*, 7: 45–98.
- Gamow, G.; and Stern, M. 1958. *Puzzle-math*. Viking Press. ISBN 9780670583355.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, EC '04, 125–131. New York, NY, USA: Association for Computing Machinery. ISBN 1581137710.
- Mahara, R. 2020. Existence of EFX for Two Additive Valuations. <http://arxiv.org/abs/2008.08798>, arXiv:2008.08798.
- Mahara, R. 2021. Extension of Additive Valuations to General Valuations on the Existence of EFX. In *29th Annual European Symposium on Algorithms (ESA 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- Plaut, B.; and Roughgarden, T. 2020. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2): 1039–1068.
- Procaccia, A. D. 2020. Technical Perspective: An Answer to Fair Division's Most Enigmatic Question. *Commun. ACM*, 63(4): 118.
- Rawls, J. 1999. *A Theory of Justice*. Harvard University Press. ISBN 9780674000773.
- Steinhaus, H. 1949. Sur la division pragmatique. *Econometrica*, 17: 315–319.