(Almost Full) EFX Exists for Four Agents (and Beyond)

Ben Berger, Avi Cohen, Michal Feldman and Amos Fiat

Abstract

The existence of EFX allocations 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 EFX is known to exist for (i) agents with identical valuations, (ii) 2 agents, (iii) 3 agents with additive valuations, (iv) agents with one of two additive valuations and (v) agents with two valued instances. 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 results are (i) every setting with 4 additive agents admits an EFX allocation that leaves at most a single item unallocated, (ii) every setting with n additive valuations has an EFX allocation with at most $n-2$ unallocated items. Moreover, all of our results extend beyond additive valuations to all *nice cancelable* valuations (a new class, including additive, unit-demand, budget-additive and multiplicative valuations, among others). Furthermore, using our new techniques, we show that previous results for additive valuations extend to nice cancelable valuations.

The full version of the paper can be found at `https://arxiv.org/abs/2102.10654`

1 Introduction

The question of justness, fairness and division of resources and commitments dates back to Aristotle [Chr42]. *Distributional justice*, the "just" allocation of limited resources, is fundamental in the work of Rawls [Raw99]. 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 [Ste49] 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 [GS58] for cake cutting, and in the context of general resource allocation by Foley [Fol67]. 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.* [LMMS04] and Budish [Bud11] 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 want to swap 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.* [CKM$^+$16, CKM$^+$19]. 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 Alices'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, [PR20], 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 [CGM20]), in settings with only one of two types of additive valuations (Mahara [Mah20]), or when the value of every agent to every item can take one of two permissible values (Amanatidis *et al.* [ABFR$^+$21]).
Indeed, Procaccia [Pro20] 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 only a small amount of items can be unallocated? The idea of partial allocations for EF and EFX allocations has appeared in multiple papers, e.g. [BKK13, CGG13, CGH19]. Caragiannis, Gravin and Huang, [CGH19] show that discarding some items gives good EFX allocations for the rest (achieving 1/2 of the maximum Nash Welfare). Chaudhury *et. al* [CKMS20] 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 [CGM20, CKMS20]. [CGM20] 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 techniques are powerful enough to allow us to (i) push the boundaries of EFX existence beyond known results, and (ii) present substantially simpler proofs for previously known results. Our results are described below, and are summarized in Table 1.

Our main result concerns EFX allocation for four agents. Extending EFX existence from three to four agents is highly non-trivial. Indeed, [CGM$^+$21] discovered an instance with 4 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 settings with 4 agents.

Theorem 1 (Main Result): Every setting with 4 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.

Setting	Prior results	Our results
EFX for 3 agents	Additive [CGM20]	Beyond additive*
EFX for n agents, one of 2 valuations	Additive [Mah20]	Beyond additive*
Partial EFX for n agents	$\leq n-1$ unallocated items [CKMS20]	$\leq n-2$ unallocated items
Partial EFX for 4 agents	≤ 3 unallocated items [CKMS20]	≤ 1 unallocated item

Table 1: Our results hold for (*)*nice cancelable* valuations, generalizing additive valuations

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.

Our second result makes an additional progress in settings with arbitrarily many agents.

Theorem 2: Every setting with n additive agents admits an EFX allocation with at most $n-2$ unallocated items.

Here too, the unallocated items are not envied by anyone. To prove Theorem 2, we show that for any EFX allocation with at least $n-1$ unallocated items, one can reshuffle bundles and reallocate them in a way that results in a Pareto-dominating EFX allocation. This means that one can find an EFX allocation with at most $n-2$ unallocated items.

In addition to these results, we establish the following extensions and simplifications.

Beyond additive valuations. Our results, Theorems 1 and 2 above, apply beyond additive valuations to a broader class that we term *nice cancelable* valuations. Moreover, we extend the results of [CGM20] (EFX for 3 additive agents) and [Mah20] (EFX for n agents with one of two additive valuations) to this new class.

Intuitively, nice cancelable valuations allow "cancelling out" of equal terms in inequalities, and as such they are a direct generalization of additive valuations[1]. Besides additive, this class also includes budget additive, unit-demand and multiplicative valuations, among others.

We stress that in our extension of the results by [CGM20, Mah20], one can have different nice cancelable valuations for different agents[2]. We further remark that the original proofs, as written, do not directly generalize to this more general class of valuations. It is our new techniques that allow us to generalize the results to this class.

Theorem 3: EFX existence for 3 agents and EFX existence for two types of valuations (any number of agents) extends beyond additive valuations to all nice cancelable valuations.

Simplification of proofs for known results. Our new techniques greatly simplify existing proofs of EFX existence for 3 agents [CGM20] and for the case of 2 types of additive valuations [Mah20]. Admittedly, simplicity is a matter of subjective judgment, but at least in terms of character count, the proof for the case of 3 agents with no envy in [CGM20] (some 5 pages) drops to half a page using our new techniques. Similarly, the proof for settings with 2 types of additive valuations in [Mah20] (some 8 pages) drops to one page using our new techniques. Moreover, in both settings, the simplified proofs apply beyond additive valuations to all nice cancelable valuations.

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.

[1] *E.g.*, if v is additive, then $v(\{a,b\}) > v(\{a,c\})$ implies $v(\{b\}) > v(\{c\})$.

[2] *E.g.*, an EFX allocation exists for three agents when agent a has a multiplicative valuation, b has a budget additive valuation, and c has a unit demand valuation. Another example, fix any two nice cancelable valuations, some agents have the 1st and others have the 2nd, an EFX allocation still exists.

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 $i \longrightarrow j$, champion edges $i \xrightarrow{g} j$, where g is an unallocated item, and generalized champion edges $i \xrightarrow{H \mid S} j$, 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 [CGM20]. 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 [CGM20].

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 $i \longrightarrow j$ 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 $i \xrightarrow{H \mid S} j$ 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 reallocation the cycle suggests.

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.

In several of our proofs we consider the case where $M_{\mathbf{X}}$ has envy edges separately from the case where it has not. For our applications, it turns out that if there are no envy edges, one can Pareto improve the EFX allocation (See Section 5.1 in [BCFF21]), in particular this advances the potential function. If there are envy edges it may no longer be possible to Pareto improve (as pointed out by [CGM20]) but one can nonetheless advance the potential function itself (see Section 5.2 in [BCFF21]).

1.3 Other Related Work

Lipton *et al.* [LMMS04] give a greedy algorithm for producing an EF1 allocation, this adds items to an agent that is not envied, or — alternately — switches bundles around an envy cycle, the so-called envy-cycles procedure. Caragiannis *et al.* [CKM$^+$19] show that maximizing Nash welfare (maximizing the geometric mean of agent utilities) gives an EF1 allocation that is also Pareto optimal. Approximations for maximizing Nash welfare were presented by Cole and Gkatzelis [CG15] and subsequent papers.

Varian [Var74] introduced the notion of *competitive equilibria from equal incomes (CEEI)*, which guarantees envy freeness. Envy free allocations that discard few items were considered in [BT00, BKK13]. In [BKK13] the resulting allocation is Pareto optimal, envy free, and maximal (no EF allocation allocates more items). There have been some papers on truthful mechanisms for proportional fairness (Mosel and Tamuz [MT10] and Cole, Gkatzelis and Goel [CGG13]). The latter mechanism uses no money and provides good guarantees — the mechanism discards a fraction of the resources to achieve truthfulness.

1.4 Paper Roadmap

Due to space limitations, some of the proofs appear in the various appendices and not in the body of the paper. The next section, Section 2, presents the model, introduces "nice cancelable valuations", gives definitions from [CGM20] as well as new definitions, notations, and proofs. Some proofs are deferred to Appendix A.

Section 3 defines "generalized championship". It may be helpful to consider Figure 1 while reading the section. How to find generalized championship edges is described in Section 3.1. Some proofs are deferred to Appendix B.

Section 4 proves that it suffices to discard at most $n-2$ items and yet guarantee the existence of an EFX allocation.

In Section 5 we show that for 4 agents it suffices to discard one item and yet still guarantee an EFX allocation. A roadmap of this proof itself appears in Figure 2. The proof exceeds the scope of this manuscript and is therefore deferred to the full version of the paper (See [BCFF21]).

Finally, Appendix C simplifies and extends prior results for 3 agents and for one of two additive valuations.

2 Preliminaries

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

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 $\mathbf{v} = (v_1, \ldots, v_n)$. An *allocation* is a vector $\mathbf{X} = (X_1, \ldots, X_n)$ 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 $i \longrightarrow j$, 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.

Nice cancelable valuations. We consider a class of valuation functions that can be viewed as a generalization of additive valuations. These include additive, unit-demand and budget-additive valuations, among others.

Definition 2.1. A valuation v is *cancelable* if for any two bundles $S, T \subseteq M$ and an item $g \in M \setminus (S \cup T)$, $v(S \cup \{g\}) > v(T \cup \{g\}) \Rightarrow v(S) > v(T)$.

A valuation v is *non-degenerate* if $v(S) \neq v(T)$ for any two different bundles S, T. A valuation v' is said to *respect* another valuation v if for every two bundles $S, T \subseteq M$ such that $v(S) > v(T)$ it also holds that $v'(S) > v'(T)$. A cancelable valuation v is *nice* if there is a non-degenerate cancelable valuation v' that respects v. In particular, any non-degenerate cancelable valuation is a nice valuation (by setting $v' \equiv v$).

We can show that in order to prove the existence of an EFX allocation for a given valuation profile $\mathbf{v} = (v_1, \ldots, v_n)$ of nice cancelable valuations, it is without loss of generality to assume that all of the valuations are non-degenerate (Lemma A.1 in Appendix A). Thus, for the remainder of this paper we assume that all valuations are cancelable and non-degenerate. Under this assumption it is easy to verify that for any valuation v and bundles S, T, R such that $R \subseteq M \setminus (S \cup T)$ we have

$$v(S \cup R) > v(T \cup R) \Leftrightarrow v(S) > v(T).$$

Other useful claims on cancelable valuations are deferred to the appendix.

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 a 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 $Y_i \geq_i X_i$ for every $i \in [n]$, and there exists some i for which the inequality is strict. Chaudhury *et al.* [CGM20] have shown that there need not exist a Pareto-dominating EFX allocation when $n = 3$ and $k \geq 1$. To overcome this obstacle they introduce an alternative "lexicographic" progress measure which we also use:

Definition 2.2 ([CGM20])**.** Fix some arbitrary ordering of the agents $a_1, \ldots, a_n$. The allocation $\mathbf{Y}$ *dominates* $\mathbf{X}$ if for some $k \in [n]$, we have that $Y_{a_j} =_{a_j} X_{a_j}$ for all $1 \leq j < k$, and $Y_{a_k} >_{a_k} X_{a_k}$.

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

Lemma 2.3. If for every EFX allocation $\mathbf{X}$ with k unallocated items, there exists a partial EFX allocation $\mathbf{Y}$ that dominates $\mathbf{X}$, then there exists an 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 general goal is to find a dominating EFX allocation $\mathbf{Y}$. In fact, in our results 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 [CGM20], [CKMS20].

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 2.4 ([CGM20])**.** 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 $D^g_{i,j}$. Note that i envies the set $(X_j \cup \{g\}) \setminus D^g_{i,j}$, but no agent strongly envies it.

An important case considered frequently in the paper is where $g \notin D^g_{i,j}$. In this case $X_j = (X_j \setminus D^g_{i,j}) \dot\cup D^g_{i,j}$. Following [CGM20], if $i \xrightarrow{g} j$ and $g \notin D^g_{i,j}$, then we say that i *g-decomposes* j into *top* and *bottom half-bundles* $(X_j \setminus D^g_{i,j})$ and $D^g_{i,j}$, 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 $(X_j \cup \{g\}) \setminus D^g_{i,j} = T_j \cup \{g\}$.

In the following observations from [CGM20], $i \xrightarrow{\not g} j$ and $i \not\rightarrow j$ are the respective negations of $i \xrightarrow{g} j$, $i \longrightarrow j$.

Observation 2.5. For every agent i, there exists an agent j who champions i with respect to g.

Observation 2.6. If $i \xrightarrow{g} j$ and $i \not\rightarrow j$, then $g \notin D^g_{i,j}$, i.e., i g-decomposes j.

Observation 2.7. If $i \xrightarrow{\not g} j$ and j is g-decomposed: $X_j = T_j \dot\cup B_j$, then $X_i >_i T_j \cup \{g\}$.

Observation 2.8. If i g-decomposes j, $i \xrightarrow{\not g} k$ and k is g-decomposed, then $T_k <_i T_j$.

3 Generalized Championship

A crucial component in our techniques is the extension of Definition 2.4 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 $D_{i,j}^{H|S}$.

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 $T \subseteq S \cup D_{i,j}^{H|S}$. We denote the negation of $i \xrightarrow{H \mid S} j$ by $i \xrightarrow{\cancel{H \mid S}} j$.

Definition 3.2. The *champion graph* $M_{\mathbf{X}} = ([n], E)$ 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: $i \longrightarrow j$ iff i envies j.
2. Champion edges: $i \xrightarrow{g} j$ iff i champions j w.r.t. g, where g is an unallocated good.
3. Generalized champion edges: $i \xrightarrow{H \mid S} j$ iff i champions j with respect to $H \mid S$.

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. Futhermore, it is not hard to verify that $i \xrightarrow{g} j$ iff $i \xrightarrow{\{g\} \mid \emptyset} j$ and that $i \longrightarrow j$ iff $i \xrightarrow{\emptyset \mid \emptyset} j$. Thus, we can treat basic edges in $M_{\mathbf{X}}$ as a generalized champion edges.

Example 3.3. Consider the instance given in Table 2, and let $\mathbf{X}$ be the partial EFX allocation where $X_1 = \{a, b, c\}$, $X_2 = \{d\}$, $X_3 = \{e, f\}$. Figure 1 depicts the graph $M_{\mathbf{X}}$. We haven't drawn all edges; rather, we chose a subset of the edges that illustrate the different types of edges. Item g is unallocated in $\mathbf{X}$, thus $U = \{g\}$. Since $\{a, b, c\} <_1 \{d\}$, $1 \longrightarrow 2$. Moreover, combined with the fact that no one strongly envies $\{d\}$, it also means that $1 \xrightarrow{g} 2$. Since $\{d\} <_2 \{b, g\}$ and no one strongly envies $\{b, g\}$, $2 \xrightarrow{g} 1$. Similarly, since $\{d\} <_2 \{f, g\}$ and no one strongly envies $\{f, g\}$, $2 \xrightarrow{g} 3$, and since $\{e, f\} <_3 \{c, g\}$ and no one strongly envies $\{c, g\}$, $3 \xrightarrow{g} 1$. Finally, it holds that $2 \xrightarrow{\{a,b\} \mid \{e\}} 3$ since $\{d\} <_2 \{a, b, f\}$ and no one strongly envies $\{a, b, f\}$.

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

Definition 3.4. A cycle $C = a_1 \xrightarrow{H_1 \mid S_1} a_2 \xrightarrow{H_2 \mid S_2} \cdots \xrightarrow{H_{k-1} \mid S_{k-1}} a_k \xrightarrow{H_k \mid S_k} a_1$ 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 $a_\ell \xrightarrow{H_\ell \mid S_\ell} \mathsf{succ}(a_\ell)$, for $\ell \in [k]$.[3]

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

[3]We 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.

	a	b	c	d	e	f	g
1	1	2	3	7			
2	1	3		6	1	3	4
3			5		3	3	2

Table 2: A profile with 7 items and 3 additive agents. Unspecified values are zeros.

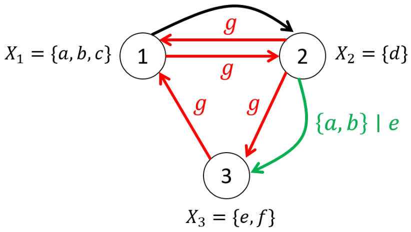

Figure 1: The graph $M_{\mathbf{X}}$ corresponding to the setting described in Table 2 and the allocation $X_1 = \{a, b, c\}$, $X_2 = \{d\}$, $X_3 = \{e, f\}$. Only a subset of the edges is shown.

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

Lemma 3.5. 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 3.6 (Following [CKMS20])**.** 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 cycles[4].

Remark 3.7. Lemma 3.5 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 $\{a_i \xrightarrow{H_i \mid S_i} \mathsf{succ}(a_i)\}_{i \in [k]}$, and if the conditions in the definition of a Pareto-improvable cycle are satisfied, then Lemma 3.5 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 $k \xrightarrow{S \mid B_j} j$ where $B_j \subseteq X_j$ is some bottom half-bundle induced by a g-decomposer of j (see discussion below Definition 2.4). Therefore, to facilitate readability we use the following convention:

Convention 3.8. For any two agents k, j and any set S disjoint from X_j, we write $k \xrightarrow{S \mid \circ} j$ (resp. $(S \mid \circ)$-champion) as shorthand for $k \xrightarrow{S \mid B_j} j$ (resp. $(S \mid B_j)$-champion), where the half-bundle B_j will be clear from the context.

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

Definition 3.9. A cycle $C = a_1 \xrightarrow{g} a_2 \xrightarrow{g} \cdots \xrightarrow{g} a_k \xrightarrow{g} a_1$ 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.

[4]Envy cycles, the simplest form of basic PI cycles, were considered in [LMMS04]. Basic PI-cycles were considered in [CKMS20] using different terminology — championship was only defined in [CGM20]. Our definition of a PI-cycle captures and generalizes these notions.

2. There are no parallel envy edges, i.e., $a_i \not\rightarrow \mathsf{succ}(a_i)$ for all i.

3. There are no internal g-champion edges, *i.e.*, $\forall i, j \in [k]$: $a_i \xrightarrow{g} a_j$ iff $a_j = \mathsf{succ}(a_i)$.

Observation 3.10. Agents j on a good g-cycle are g-decomposed by $\mathsf{pred}(j)$ into $X_j = T_j \uplus B_j$.

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 3.11. If $i \not\rightarrow j$ then $i \xrightarrow{\cancel{B_j \mid \circ}} j$.

Observation 3.12. For any two agents i, j along a good g-cycle, $\mathsf{pred}(i) \xrightarrow{\cancel{B_i \mid \circ}} j$.

Lemma 3.13. Let C be a good g-cycle. For any agent i along the cycle, there exists an agent a such that $a \xrightarrow{B_i \mid \circ} \mathsf{succ}(i)$.

Proof. C is a good cycle, hence i g-decomposes $\mathsf{succ}(i)$ into $X_{\mathsf{succ}(i)} = T_{\mathsf{succ}(i)} \cup B_{\mathsf{succ}(i)}$. Furthermore, $i \xrightarrow{\cancel{g}} i$ and i is g-decomposed into $X_i = T_i \cup B_i$. Thus by Observation 2.8 (with $j = \mathsf{succ}(i)$ and $k = i$ in the Observation statement) we have $T_i <_i T_{\mathsf{succ}(i)}$. Hence, by cancelability, $X_i = T_i \cup B_i <_i T_{\mathsf{succ}(i)} \cup B_i$. Since the set $T_{\mathsf{succ}(i)} \cup B_i$ is envied by some agent, it must have a most envious agent and the claim follows. □

Lemma 3.14. Let C be a good g-cycle. Let i, j, k be agents along the cycle. If $k \xrightarrow{B_i \mid \circ} j$, then there exists an agent a (not necessarily in the cycle) such that $a \xrightarrow{B_i \mid \circ} \mathsf{succ}(k)$.

Proof. If $k = \mathsf{pred}(j)$, then $j = \mathsf{succ}(k)$ and we are done (take $a = k$). Assume otherwise. C is a good cycle, hence k g-decomposes $\mathsf{succ}(k)$ into $X_{\mathsf{succ}(k)} = T_{\mathsf{succ}(k)} \cup B_{\mathsf{succ}(k)}$. Furthermore, $k \xrightarrow{\cancel{g}} j$ since $k \neq \mathsf{pred}(j)$, and j is g-decomposed into $X_j = T_j \cup B_j$. Thus by Observation 2.8 we have $T_j <_k T_{\mathsf{succ}(k)}$. Hence, $X_k <_k T_j \cup B_i <_k T_{\mathsf{succ}(k)} \cup B_i$, where the first inequality holds by $k \xrightarrow{B_i \mid \circ} j$ and the second by cancelability. Since the set $T_{\mathsf{succ}(k)} \cup B_i$ is envied by some agent, it must have a most envious agent and the claim follows. □

For every bottom half-bundle B_i along a good g-cycle C, applying Lemma 3.13 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 3.14 to discover a new B_i-champion edge. Once again, if the new edge is internal to the cycle, then we can reapply Lemma 3.14. 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 3.15. Let C be a good g-cycle. Let i, j, k be three agents along C. If $i \xrightarrow{B_k \mid \circ} j$ and k is on the path from j to i in C, then we say that the edge $i \xrightarrow{B_k \mid \circ} j$ is a *good edge* (or good B_k-edge). If $\ell \xrightarrow{B_k \mid \circ} j$ for some agent ℓ outside the cycle C, then we say that the edge $\ell \xrightarrow{B_k \mid \circ} j$ is an *external edge* (or external B_k-edge).

The figure on the right illustrates Definition 3.15. 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 $2 \xrightarrow{B_1 \mid B_4} 4$ is a *good edge*, since 1 is on the path from 4 to 2 in C. On the other hand, $3 \xrightarrow{B_4 \mid B_2} 2$ 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. $5 \xrightarrow{B_2 \mid B_3} 3$ is an external edge.

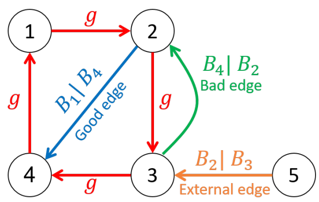

Theorem 3.16. Let C be a good g-cycle. For every agent j along the cycle, there exists either a good B_j-edge in C, or an external B_j-edge in C.

Proof. Assume without loss of generality that $C = 1 \xrightarrow{g} 2 \xrightarrow{g} \cdots \xrightarrow{g} k \xrightarrow{g} 1$ and $j = 1$, i.e., we try to find B_1-champion edges. By Lemma 3.13 there exists an edge $\ell_1 \xrightarrow{B_1 \mid \circ} 2$ 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 3.14 there exists an edge $\ell_2 \xrightarrow{B_1 \mid \circ} \mathsf{succ}(\ell_1)$, for some agent ℓ_2 which can be equal to ℓ_1. As long as the result of Lemma 3.14 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 $\ell_{i+1} \xrightarrow{B_1 \mid \circ} \mathsf{succ}(\ell_i)$ for every $i \geq 1$.

If for some $i \geq 1$, we have $\ell_{i+1} \leq \ell_i$ then the edge $\ell_{i+1} \xrightarrow{B_1 \mid \circ} \mathsf{succ}(\ell_i)$ is a good edge (since the path from $\mathsf{succ}(\ell_i)$ 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. □

The following observation and its corollary allow us to narrow down the possible configurations of B_j-edges obtained from Theorem 3.16.

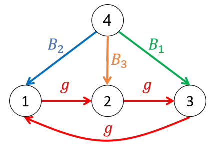

Observation 3.17. If $i \xrightarrow{B_j \mid \circ} k$ and $i \not\rightarrow k$ then $B_k <_i B_j$.

Corollary 3.18. Let C be a good g-cycle. Consider the set of B_j-edges guaranteed by Theorem 3.16 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 on the right demonstrates an impossible configuration).

4 EFX with at most $n - 2$ unallocated goods

In this section we prove the following:

Theorem 4.1. For every profile of n additive valuations (and more generally, nice cancelable valuations), there exists an EFX allocation $\mathbf{X}$ with at most $n - 2$ unallocated items. Moreover, in $\mathbf{X}$ no agent envies the set of unallocated items.

The following lemma shows that the basic edges of the graph $M_{\mathbf{X}}$ follow a very particular structure. This lemma is used in the proof of Theorem 4.1.

Lemma 4.2. Let $\mathbf{X}$ be an EFX allocation with at least $n-1$ unallocated items, and let G be $M_{\mathbf{X}}$ restricted to basic edges, *i.e.*, envy and basic champion edges as per Definition 2.4.

If G does not admit a basic PI cycle, then the number of unallocated goods is exactly $n - 1$, and G is a union of $n - 1$ parallel Hamiltonian cycles. Every such cycle consists of g-champion edges for some unallocated item g.

Proof. Recall that U denotes the set of unallocated items, and let $U = \{g_1, \ldots, g_k\}$ for some $k \geq n-1$. Let e_1^1 be an arbitrary incoming g_1-champion edge of agent 1 (such an edge exists by Observation 2.5).

If the source of this edge is agent 1 we are done (we have a self champion). Assume w.l.o.g. that the source of e_1^1 is agent 2, and consider its incoming g_2-champion edge, denoted e_2^2. If the source of e_2^2 is agent 2 or agent 1, we are done (in the first case we have a self g_2-champion, and in the second case we have a basic PI cycle of size 2). Assume w.l.o.g. that the source of e_2^2 is agent 3. We can continue this way to conclude that w.l.o.g we have a directed path $n \xrightarrow{g_{n-1}} n-1 \xrightarrow{g_{n-2}} \cdots \xrightarrow{g_1} 1$. If $k \geq n$, then consider the incoming g_n-champion edge of agent n. No matter what the source of this edge is, it must close a basic PI cycle.

Ergo, $k = n - 1$. Consider the incoming g_1-champion edge of agent n, denoted e_n^1. The source of this edge must be agent 1 (every other option closes a basic PI cycle). Now consider the incoming g_2-champion edge of agent 1, e_1^2. Similarly, the source of this edge must be agent 2. We can again continue this way until we get to the incoming g_{n-1}-champion edge of agent $n-2$, denoted e_{n-2}^{n-1}, and conclude that its source is agent $n-1$. We consider the incoming g_1-champion edge of agent $n-1$, and continue with the same reasoning, to finally conclude that the source of the incoming g_{n-1}-champion edge of agent n is agent 1.

At this point there is a Hamiltonian cycle $n \xrightarrow{g_i} n-1 \xrightarrow{g_i} \cdots \xrightarrow{g_i} 1 \xrightarrow{g_i} n$ for every i. Any other basic champion or envy edge must close a basic PI cycle and thus does not exist by assumption. This concludes the second part of the lemma. $\square$

We are now ready to prove Theorem 4.1.

Proof. By Lemma 2.3, it suffices to prove that if $\mathbf{X}$ has at least $n-1$ unallocated items, then there exists a Pareto dominating EFX allocation $\mathbf{Y}$. By Lemma 3.5, it suffices to find a PI cycle in $M_{\mathbf{X}}$.

By Lemma 4.2, there are exactly $n-1$ unallocated goods, and G is a union of $n-1$ parallel Hamiltonian cycles $1 \xrightarrow{g_i} 2 \xrightarrow{g_i} \cdots n \xrightarrow{g_i} 1$, one for each unallocated good $g_1, \ldots, g_{n-1}$.

Furthermore, it follows from Lemma 4.2 that we can assume that we have no other champion edges and no envy edges. Thus, all these cycles are good.

By Theorem 3.16 for all agents k and every unallocated good g there exists a good $D_{k-1,k}^{g}$-champion edge, from some agent j to some agent j'. (Indeed, since G is a union of parallel Hamiltonian cycles, no external edges exists.) Choose some arbitrary agent k and unallocated good g and corresponding agents j, j', and let $Z = D_{k-1,k}^{g}$. We show that $j \xrightarrow{Z \mid \circ} j'$ closes a Pareto improvable cycle in $M_{\mathbf{X}}$.

By definition of a good edge, there is a unique path P consisting of g-champion edges from agent j' to j, passing through agent k. By Lemma 4.2, for every $q \in U$, every edge in P has a parallel q-champion edge. Note that P has at most $n-1$ edges.

Rename the agents so that agent j is agent $|P|$, agent j' is agent 1, let k' be the new index for agent k, and let $P = 1, 2, \ldots, k'-1, k', \ldots |P|$. Note that Z is now the discard set of the champion edge $k'-1 \xrightarrow{g} k'$. Let $U' = U \setminus \{g\}$ ($|U'| = n-2$), and rename the items in U' where $U' = \{r_1, r_2, \ldots, r_{k'-2}, r_{k'}, \ldots, r_{n-1}\}$. We now describe a PI cycle:

$$1 \xrightarrow{r_1} 2 \xrightarrow{r_2} 2 \xrightarrow{r_3} \cdots \xrightarrow{r_{k'-2}} k'-1 \xrightarrow{g} k' \xrightarrow{r_{k'}} k'+1 \xrightarrow{r_{k'+1}} \cdots$$
$$\cdots \xrightarrow{r_{|P|-1}} |P| \xrightarrow{Z \mid \circ} 1$$

Since Z is the discard set of the champion edge $k'-1 \xrightarrow{g} k'$, it is discarded along the cycle. All other edges in the cycle are with respect to distinct unallocated goods. Therefore, this is a Pareto-improvable cycle. $\square$

5 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 1 unallocated good. We prove this for the class of nice cancelable valuations which contains additive. By Lemma 2.3 it suffices to prove:

Theorem 5.1. Let $\mathbf{X}$ be an EFX allocation on 4 nice cancelable agents with at least 2 unallocated items. Then, there exists an EFX allocation $\mathbf{Y}$ that dominates $\mathbf{X}$.

While this is indeed our main result, its proof exceeds the scope of this manuscript. We therefore omit the proof in its entirety, and refer the interested reader to the full version of the paper (see [BCFF21]). 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.

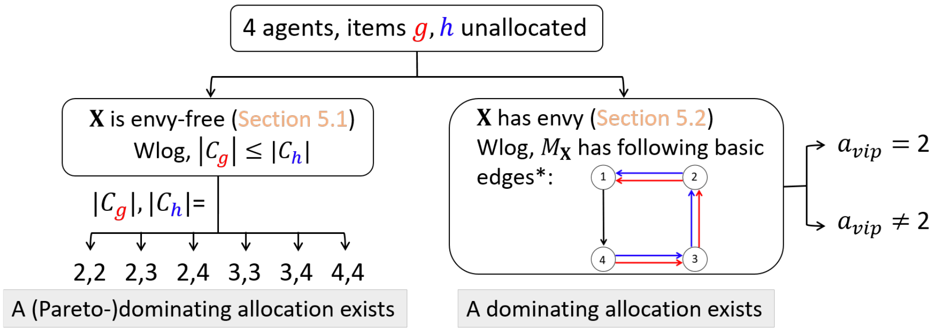

Figure 2: High-level roadmap of the proof of Theorem 5.1. $|C_g|$ and $|C_h|$ are the lengths of some good cycles of items g and h, respectively. $(*)$ The g and h champions of agent 4 can be either agent 1 or 2.

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 (Section 5.1 of [BCFF21]) or not (Section 5.2 of [BCFF21]). 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 must exist in the champion graph $M_{\mathbf{X}}$.

When $\mathbf{X}$ has envy, we argue that the only interesting case is where the basic edges 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 splits to cases accordingly, where the case $a_{\mathsf{vip}} = 2$ is treated separately from the case where a_{vip} is one of the other three agents. Our approach here is heavily inspired by and follows a similar high-level structure to that of [CGM20] in their analysis of the envy case in their 3 agent result. Our proof structure is depicted in Figure 2.

Acknowledgment

We are very grateful to Karin Egri for her great help in producing the figures used in this paper. We are deeply in debt to Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta and Misra Pranabendu for their wonderful papers and for sharing unpublished ongoing results.

Ben Berger, Avi Cohen and Michal Feldman have been supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 866132), and by the Israel Science Foundation (grant number 317/17). Amos Fiat has been supported by the Israel Science Foundation (grant number 1841/14) and the Blavatnik fund.

References

[ABFR+21] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A. Voudouris. Maximum Nash welfare and other stories about EFX. *Theoret. Comput. Sci.*, 863:69–85, 2021.

[BCFF21] Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. (almost full) EFX exists for four agents (and beyond), 2021. `http://arxiv.org/abs/2102.10654`.

[BKK13] Steven J. Brams, D. Marc Kilgour, and Christian Klamler. Two-person fair division of indivisible items: An efficient, envy-free algorithm. MPRA Paper 47400, University Library of Munich, Germany, June 2013.

[BT00] Steven Brams and Alan Taylor. *The win-win solution - guaranteeing fair shares to everybody.* WW Norton & Company, 2000.

[Bud11] Eric Budish. The combinatorial assignment problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

[CG15] Richard Cole and Vasilis Gkatzelis. Approximating the Nash social welfare with indivisible items. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, page 371–380, New York, NY, USA, 2015. Association for Computing Machinery.

[CGG13] Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Mechanism design for fair division: Allocating divisible items without payments. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, page 251–268, New York, NY, USA, 2013. Association for Computing Machinery.

[CGH19] Ioannis Caragiannis, Nick Gravin, and Xin Huang. 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*, pages 527–545, 2019.

[CGM20] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. EFX exists for three agents. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 1–19, 2020.

[CGM+21] Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta, and Misra Pranabendu. Personal communication. 2021.

[Chr42] Anton-Hermann Chroust. Aristotle's conception of equity (epieikeia). *Notre Dame Law Review*, 18:119–128, 1942.

[CKM+16] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, page 305–322, New York, NY, USA, 2016. Association for Computing Machinery.

[CKM+19] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.

[CKMS20] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. A little charity guarantees almost envy-freeness. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2658–2672. SIAM, 2020.

[Fol67] Duncan K. Foley. Resource allocation and the public sector. *Yale economic essays*, 7:45–98, 1967.

[GS58] George Gamow and Marvin Stern. *Puzzle-math.* Viking Press, 1958.

[LMMS04] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, EC '04, page 125–131, New York, NY, USA, 2004. Association for Computing Machinery.

[Mah20] Ryoga Mahara. Existence of EFX for two additive valuations, 2020. `http://arxiv.org/abs/2008.08798`.

[MT10] Elchanan Mossel and Omer Tamuz. Truthful fair division. In *Proceedings of the Third International Conference on Algorithmic Game Theory*, SAGT'10, page 288–299, Berlin, Heidelberg, 2010. Springer-Verlag.

[PR20] Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2):1039–1068, 2020.

[Pro20] Ariel D. Procaccia. Technical perspective: An answer to fair division's most enigmatic question. *Commun. ACM*, 63(4):118, March 2020.

[Raw99] John Rawls. *A Theory of Justice.* Harvard University Press, 1999.

[Ste49] Hugo Steinhaus. Sur la division pragmatique. *Econometrica*, 17:315–319, 1949.

[Var74] Hal Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.

Ben Berger
Tel Aviv University
Tel Aviv, Israel
Email: `benberger1@tauex.tau.ac.il`

Avi Cohen
Tel Aviv University
Tel Aviv, Israel
Email: `avicohen2@mail.tau.ac.il`

Michal Feldman
Tel Aviv University
Tel Aviv, Israel
Email: `michal.feldman@cs.tau.ac.il`

Amos Fiat
Tel Aviv University
Tel Aviv, Israel
Email: `fiat@tau.ac.il`

A Proofs and Claims from Section 2

Lemma A.1. To prove the existence of an EFX allocation for a given valuation profile $\mathbf{v} = (v_1, \ldots, v_n)$ of nice cancelable valuations, it is without loss of generality to assume that all of the valuations are non-degenerate.

Proof. For each $i \in [n]$, v_i is a nice cancelable valuation. Hence, for every $i \in [n]$ there exists a non-degenerate cancelable valuation v'_i that respects v_i. Let $\mathbf{v'} = (v'_1, \ldots, v'_n)$ be the corresponding valuation profile.

We show that any allocation which is EFX for the profile $\mathbf{v'}$ is also EFX for the profile $\mathbf{v}$. Let $\mathbf{X} = (X_1, \ldots, X_n)$ be an EFX allocation for the profile $\mathbf{v'}$ and assume towards contradiction that $\mathbf{X}$ is not an EFX allocation for the profile $\mathbf{v}$. Hence, under the profile $\mathbf{v}$ there exists an agent i that strongly envies another agent j, *i.e.*, $v_i(X_i) < v_i(X_j \setminus \{h\})$ *for some* $h \in X_j$. *Since* v'_i *respects* v_i, *it follows that* $v'_i(X_i) < v'_i(X_j \setminus \{h\})$, *in contradiction to* $\mathbf{X}$ *being EFX over the profile* $\mathbf{v'}$. □

The family of nice cancelable valuations contains some well-known classes of valuations. Additive valuations are clearly cancelable and are shown to be nice in [CGM20]. The following lemma shows that this class contains many other classes of valuations, including unit-demand, budget-additive and multiplicative.

Lemma A.2. Unit-demand, budget-additive and multiplicative valuations are nice cancelable.

Proof. **Budget-Additive:** Let v be a budget-additive valuation, *i.e.*, for every $S \subseteq M$,

$$v(S) = \min \left\{ \sum_{g \in S} v(g),\ B \right\},$$

for some $B > 0$. We start by proving v is cancelable. Consider $S, T \subseteq M$ and $g \in M \setminus (S \cup T)$ such that

$$v(S \cup \{g\}) > v(T \cup \{g\}).$$

First, $v(S \cup \{g\}) \leq B$, by definition of v. Therefore, $v(T \cup \{g\}) < B$, so v is additive over $T \cup \{g\}$. It follows that $v(T) = v(T \cup \{g\}) - v(\{g\})$. Second, since budget-additive valuations are sub-additive, $v(S) \geq v(S \cup \{g\}) - v(\{g\})$. Combining these two observations we get

$$v(S) \geq v(S \cup \{g\}) - v(\{g\}) > v(T \cup \{g\}) - v(\{g\}) = v(T).$$

This proves that v is cancelable.

We next prove that v is nice. Define the valuation $v' : 2^M \to \mathbb{R}_{\geq 0}$ as the underlying additive valuation of v, *i.e.*, for every $S \subseteq M$,

$$v'(S) = \sum_{g \in S} v(g).$$

We now show that v' respects v.

Suppose $v(S) > v(T)$ for some $S, T \subseteq M$. Since $v(S) \leq B$, it follows that $v(T) < B$, and thus v is additive over T. By definition of v', this implies that $v(T) = v'(T)$. Furthermore, notice that $v'(S) \geq v(S)$. Thus,

$$v'(S) \geq v(S) > v(T) = v'(T).$$

This proves that v' respects v.

Finally, since v' is an additive valuation it is nice and cancelable (as shown in [CGM20]). Therefore, there exists a non-degenerate cancelable valuation v'' that respects v'. Because v' respects v, it follows by transitivity that v'' respects v as well. Since v'' is non-degenerate and cancelable, this proves that v is nice.

Multiplicative: Let v be a multiplicative valuation, *i.e.*, for every $S \subseteq M$,

$$v(S) = \prod_{g \in S} v(g).$$

Multiplicative valuations are trivially cancelable. Since additive valuations are nice, and since taking the log of a multiplicative valuation gives us an additive function, similar arguments as in [CGM20] imply that multiplicative is also nice.

Unit-Demand: Let v be a unit-demand valuation, *i.e.*, for every $S \subseteq M$,

$$v(S) = \max_{g \in S} v(g).$$

We first show v is cancelable. Consider $S, T \subseteq M$ and $g \in M \setminus (S \cup T)$ such that

$$v(S \cup \{g\}) > v(T \cup \{g\}).$$

Clearly, g is not the maximal element in $S \cup \{g\}$, otherwise we would have $v(S \cup \{g\}) = v(\{g\}) \leq v(T \cup \{g\})$. Therefore, $v(S \cup \{g\}) = v(S)$. We get

$$v(S) = v(S \cup \{g\}) > v(T \cup \{g\}) \geq v(T).$$

This proves that v is cancelable.

We next prove that v is nice. Define

$$\delta = \min_{S,T \subseteq M,\ v(S) \neq v(T)} |v(S) - v(T)|.$$

That is, δ is the minimal difference between the value of any two non-equal valued sets of items. Let $g_0, \ldots, g_{m-1}$ be the items in M ordered in non-decreasing value, ties broken arbitrarily. Let $\varepsilon = 2^{-(m+1)}\delta$. Define the valuation $v' : 2^M \to \mathbb{R}_{\geq 0}$ as follows:

$$v'(S) = v(S) + \varepsilon \sum_{i : g_i \in S} 2^i.$$

To complete the proof we need to show that v' is a non-degenerate cancelable valuation and that v' respects v. We begin with the latter. Suppose $v(S) > v(T)$ for some $S, T \subseteq M$. This implies that $v(S) > v(T) + \delta/2$, by definition of δ. Moreover, since $\sum_{i=0}^{m-1} 2^i < 2^m$, we have

$$v'(T) = v(T) + \varepsilon \sum_{i : g_i \in T} 2^i < v(T) + \varepsilon \cdot 2^m = v(T) + \frac{\delta}{2}.$$

We get

$$v'(S) > v(S) > v(T) + \frac{\delta}{2} > v'(T),$$

as required.

We next prove v' is non-degenerate. For all $S, T \subseteq M$ such that $v(S) \neq v(T)$, we have shown above that $v'(S) \neq v'(T)$. So it remains to show that $v'(S) \neq v'(T)$ whenever $v(S) = v(T)$ and $S \neq T$. Since $v(S) = v(T)$, to prove that $v'(S) \neq v'(T)$ it suffices to show that

$$\varepsilon \sum_{i : g_i \in S} 2^i \neq \varepsilon \sum_{i : g_i \in T} 2^i,$$

which clearly holds for every $S \neq T$.

Finally, we prove v' is cancelable. Consider $S, T \subseteq M$ and $g_j \in M \setminus (S \cup T)$ such that

$$v'(S \cup \{g_j\}) > v'(T \cup \{g_j\}).$$

It is impossible that $v(S \cup \{g_j\}) < v(T \cup \{g_j\})$, since v' respects v. If $v(S \cup \{g_j\}) > v(T \cup \{g_j\})$, then $v(S) > v(T)$ since v is cancelable. Since v' respects v, this implies $v'(S) > v'(T)$, so we are done. We are left with the case where $v(S \cup \{g_j\}) = v(T \cup \{g_j\})$. Since $v'(S \cup \{g_j\}) > v'(T \cup \{g_j\})$, this implies

$$\varepsilon \cdot \sum_{i: g_i \in S \cup \{g_j\}} 2^i > \varepsilon \cdot \sum_{i: g_i \in T \cup \{g_j\}} 2^i.$$

Eliminating $\varepsilon \cdot 2^j$ from both sides,

$$\varepsilon \sum_{i: g_i \in S} 2^i > \varepsilon \sum_{i: g_i \in T} 2^i. \tag{1}$$

If $v(S) \geq v(T)$ then Equation (1) implies $v'(S) > v'(T)$ and we are done. We complete the proof by showing that the case $v(S) < v(T)$ is impossible. Assume towards contradiction that $v(S) < v(T)$. Let g_s and g_t be the highest valued items in S and T, respectively (breaking ties according to the ordering we defined on the items). Since $v(S) < v(T)$ and v is unit-demand, $v(g_s) < v(g_t)$. Due to our ordering, it follows that $s < t$. Therefore,

$$\sum_{i: g_i \in S} 2^i \leq \sum_{i=0}^{s} 2^i < 2^{s+1} \leq 2^t \leq \sum_{i: g_i \in T} 2^i,$$

in contradiction to Equation (1). This shows that v' is cancelable. □

On the other hand, not all cancelable valuations are nice, as shown in Proposition A.3.

Proposition A.3. Not all cancelable valuations are nice. Even when restricted to submodular cancelable valuations it need not be nice.

Proof. We define a valuation v over the set of items $M = \{a, b, c, d, e, f\}$ and show that it is cancelable and submodular but not nice. The following table defines the value of v over each singleton.

	a	b	c	d	e	f
v	101	102	102	103	103	104

Notice that $v(a) < v(b) = v(c) < v(d) = v(e) < v(f)$, i.e., the values are non-decreasing from left to right.

The next table defines the value of v over each pair of items. For convenience, we depict this table as a matrix, where the coordinate (x, y) contains the value $v(\{x, y\})$ (the matrix is symmetric).

	a	b	c	d	e	f
a	-	152	152	153	153	154
b	152	-	152	155	155	156
c	152	152	-	155	155	156
d	153	155	155	-	155	156
e	153	155	155	155	-	156
f	156	156	156	156	156	-

Finally, for any set $S \subseteq M$ containing three or more items we define $v(S) = 200$, and we set $v(\emptyset) = 0$. This completes the definition of v over all subsets of M.

Cancelability: Consider a pair of sets $S, T \subseteq M$ and an item $g \in M \setminus (S \cup T)$ such that $v(S \cup \{g\}) > v(T \cup \{g\})$. First consider the cases where $|S| = |T| = 1$. Therefore, $S \cup \{g\}$ and $T \cup \{g\}$ are both pairs of items whose values are located in row g of the matrix above. It is not hard to verify that within each row of the matrix the values are non-decreasing left to right with equality only between the columns of equal valued items (the columns b and c and the columns d and e contain equal values). Therefore, the fact that $v(S \cup \{g\}) > v(T \cup \{g\})$ implies that $v(S \cup \{g\})$ is located to the right of $v(T \cup \{g\})$ in row g of the matrix. Therefore, $v(S) \geq v(T)$. The case $v(S) = v(T)$ implies $v(S \cup \{g\}) = v(T \cup \{g\})$ since the columns of equal valued singletons are identical in the matrix of pairs. Thus, $v(S) > v(T)$. This completes the case $|S| = |T| = 1$.

If $|T| > 2$, then $v(T \cup \{g\}) = 200$ and there exists no set S such that $v(S \cup \{g\}) > 200$. Therefore, we may assume $|T| \leq 1$. The case $|T| = 0$ is easy. In this case, $v(S \cup \{g\}) > v(T \cup \{g\})$ implies that S is non-empty. Since $T = \emptyset$ and $S \neq \emptyset$, we obtain $v(S) > v(T)$, so cancelability is maintained.

We are left with the case $|T| = 1$. The case $|S| = 0$ is not possible since no singleton has a higher value than a pair of items. The case $|S| = |T| = 1$ has been handled above. So it is left to consider $|S| > 1$ and $|T| = 1$. In this case, notice that $v(S) > 150$ while $v(T) < 150$, therefore $v(S) > v(T)$, as desired. This proves that v is cancelable.

Submodularity: It suffices to show that for every item g and every pair of sets S and T such that $S \subseteq T \subseteq M \setminus \{g\}$ we have that $v(S \cup \{g\}) - v(S) \geq v(T \cup \{g\}) - v(T)$. This condition clearly holds if $S = T$, so assume $S \subsetneq T$. Therefore, $|S| < |T|$. Notice that the following holds for any set $Z \subseteq M$ and any item $x \in M \setminus Z$:

$$\begin{aligned}
&\text{if } |Z| = 0, && 100 < v(Z \cup \{x\}) - v(Z) \leq 104;\\
&\text{if } |Z| = 1, && 50 \leq v(Z \cup \{x\}) - v(Z) < 60;\\
&\text{if } |Z| = 2, && 40 < v(Z \cup \{x\}) - v(Z) < 50;\\
&\text{if } |Z| \geq 3, && v(Z \cup \{x\}) - v(Z) = 0.
\end{aligned}$$

Therefore, the fact that $|S| < |T|$ directly implies that $v(S \cup \{g\}) - v(S) \geq v(T \cup \{g\}) - v(T)$. This proves that v is submodular.

Not nice: Assume towards contradiction that there exists a non-degenerate cancelable valuation v' that respects v. Notice that the following inequalities hold:

$$154 = v(\{a, f\}) < v(\{b, e\}) = 155, \tag{2}$$

$$155 = v(\{d, e\}) < v(\{c, f\}) = 156, \tag{3}$$

$$152 = v(\{b, c\}) < v(\{a, d\}) = 153. \tag{4}$$

Therefore, the analogous inequalities must hold for v'. Consider the comparison between $\{a, d, f\}$ and $\{b, d, e\}$ in v'. v' is non-degenerate, so $v'(\{a, d, f\}) \neq v'(\{b, d, e\})$. If $v'(\{a, d, f\}) > v'(\{b, d, e\})$, then by cancelability $v'(\{a, f\}) > v'(\{b, e\})$, which contradicts Equation (2), so

$$v'(\{a, d, f\}) < v'(\{b, d, e\}). \tag{5}$$

Similarly, due to Equation (3) we obtain

$$v'(\{b, d, e\}) < v'(\{b, c, f\}), \tag{6}$$

and from Equation (4) we get

$$v'(\{b, c, f\}) < v'(\{a, d, f\}). \tag{7}$$

However, combining Equations (5) and (6) we get $v'(\{a, d, f\}) < v'(\{b, c, f\})$, in contradiction to Equation (7). This proves that there exists no non-degenerate cancelable valuation v' such that v' respects v, and thus v is not nice. $\square$

If an additive agent strongly envies some bundle S, then, iterative removal of the least valued item until strong envy is eliminated results in a smallest size subset of S that the agent envies. The next lemma shows that this property extends to cancelable valuations.

Lemma A.4. Let v be a cancelable valuation. Let T be some bundle, and let S a subset of T. Let Z be the subset obtained from T by iteratively removing the item with least marginal contribution until the leftover bundle has size $|S|$. Then $v(Z) \geq v(S)$.

Proof. Define $T_0 = T$, and for $j \geq 0$ define $T_{j+1} = T_j \setminus \{c\}$, where $c \in T_j$ is the item with least marginal contribution to T_j. It suffices to prove that for every $0 \leq j \leq |T|$ we have

$$T_j = \arg \max_{S \subseteq T : |S| = |T| - j} v(S).$$

We prove by induction on j. For $j = 0$ the claim is immediate. Assume that the claim is true for j and we prove for $j + 1$. Let $c \in T_j$ be the item with least marginal contribution to T_j, hence by definition we have $T_{j+1} = T_j \setminus \{c\}$. Let $S \subseteq T$ such that $|S| = |T_{j+1}|$. We need to show that $v(S) \leq v(T_{j+1})$. If $S = T_{j+1}$, then this is immediate. Therefore, assume $S \neq T_{j+1}$. Since S and T_{j+1} have the same size, this means that there is some item $b \in T_{j+1} \setminus S$, and thus $S \cup \{b\}$ and T_j have the same size. By the induction hypothesis we get

$$v(S \cup \{b\}) \leq v(T_j) = v(T_j \setminus \{b\} \cup \{b\}),$$

implying $v(S) \leq v(T_j \setminus \{b\}) \leq v(T_j \setminus \{c\}) = v(T_{j+1})$ where the first inequality holds by cancelability, and the second by definition of c. The claim follows. $\square$

Proof of Lemma 2.3. The first part of the lemma follows directly from the fact that the number of possible allocations is finite and domination is a total order relation.

We now show that if in a given partial EFX allocation $\mathbf{X}$ some agent envies the set of unallocated items then there exists a partial EFX allocation $\mathbf{Y}$ that Pareto dominates $\mathbf{X}$. Analogous to the above, this proves the second part of the lemma.

If some agent envies U, then there exists a subset T of U that some agent i envies and T is a smallest subset of U that some agent envies. We obtain $\mathbf{Y}$ by replacing X_i with T. $\mathbf{Y}$ is EFX, since i did not strongly envy anyone before and now she is better off, and no one strongly envies i by minimality of T. $\square$

Proof of Observation 2.5. Since valuations are non-degenerate and monotone, $X_i <_i X_i \cup \{g\}$. Since i envies $X_i \cup \{g\}$, this set has a most envious agent, and by definition that agent is a champion of i with respect to g. $\square$

Proof of Observation 2.6. By definition of $D^g_{i,j}$, agent i envies $(X_j \cup \{g\}) \setminus D^g_{i,j}$. If $g \in D^g_{i,j}$ then $(X_j \cup \{g\}) \setminus D^g_{i,j}$ is a subset of X_j, implying that agent i envies a subset of X_j and thus the set X_j as well. Hence, i envies j, in contradiction to our assumption. $\square$

Proof of Observation 2.7. Assume that j is g-decomposed by some agent k into $X_j = T_j \uplus B_j$. By definition of $D^g_{k,j}$, no agent strongly envies $(X_j \cup \{g\}) \setminus D^g_{k,j} = T_j \cup \{g\}$. Therefore, any agent that envies this set is a most envious agent of $X_j \cup \{g\}$, and thus a g-champion of j. Hence, agent i does not envy that set. $\square$

Proof of Observation 2.8. We have, $T_k \cup \{g\} <_i X_i <_i T_j \cup \{g\}$, where the first inequality is by Observation 2.7 since $i \xrightarrow{\not{g}} k$ and the second inequality is by definition of basic championship (since i g-decomposes j). $\square$

B Proofs from Section 3

Proof of Lemma 3.5. Let C be a Pareto-improvable cycle in $M_{\mathbf{x}}$, and assume w.l.o.g. that $C = 1 \xrightarrow{H_1 \mid S_1} 2 \xrightarrow{H_2 \mid S_2} \cdots \xrightarrow{H_{k-1} \mid S_{k-1}} k \xrightarrow{H_k \mid S_k} 1$. Define the allocation $\mathbf{Y}$ as follows: for every agent i,

$$Y_i = \begin{cases} \big(\big(X_{\mathsf{succ}(i)} \setminus S_i\big) \cup H_i\big) \setminus D^{H_i \mid S_i}_{i,\mathsf{succ}(i)} & i \in [k] \text{ } (i.e., i \text{ on cycle } C) \\ X_i & \text{otherwise} \end{cases}$$

First, note that the assumptions on the sets H_i, S_i ensure that the sets Y_i are disjoint. That is, $\mathbf{Y}$ is indeed an allocation. For every $i \in [k]$, i is the $(H_i \mid S_i)$-champion of $\mathsf{succ}(i)$. Thus $X_i <_i Y_i$, by definition of generalized championship. Since the bundles did not change for agents outside the cycle, we conclude that $\mathbf{Y}$ Pareto-dominates $\mathbf{X}$.

It remains to show that $\mathbf{Y}$ is EFX. Since no agent becomes worse off in the transition from $\mathbf{X}$ to $\mathbf{Y}$, it suffices to show that in allocation $\mathbf{X}$ no agent strongly envies the set Y_i for all $i \in [n]$. Indeed, if i is an agent outside the cycle, by the fact that $\mathbf{X}$ is EFX, no agent strongly envies $Y_i = X_i$ in $\mathbf{X}$. If i is an agent in the cycle, no agent strongly envies $Y_i = \big(\big(X_{\mathsf{succ}(i)} \setminus S_i\big) \cup H_i\big) \setminus D^{H_i \mid S_i}_{i,\mathsf{succ}(i)}$ in $\mathbf{X}$, by definition of generalized championship. $\square$

Proof of Observation 3.10. This holds by Observation 2.6 since $\mathsf{pred}(j) \xrightarrow{g} j$ and $\mathsf{pred}(j) \not\rightarrow j$ by definition of a good g-cycle. $\square$

Proof of Observation 3.11. Recall that a generalized championship relation $i \xrightarrow{H \mid S} j$ is required to satisfy $H \cap S = \emptyset$. Here $H = S = B_j$. Thus, if $B_j \neq \emptyset$ then the requirement clearly does not hold. Otherwise $B_j = \emptyset$ and $i \xrightarrow{\emptyset \mid \emptyset} j$ implies that $i \longrightarrow j$, which we assume does not hold. In any case, the relation $i \xrightarrow{B_j \mid \circ} j$ cannot hold. $\square$

Proof of Observation 3.12. If $i = j$ the statement holds by Observation 3.11. Assume otherwise. By Observation 3.10 we have that $\mathsf{pred}(i)$ g-decomposes i. Further, by definition of a good g-cycle $\mathsf{pred}(i) \not\rightarrow i$ and $\mathsf{pred}(i) \xrightarrow{\cancel{g}} j$. Therefore, $X_{\mathsf{pred}(i)} >_{\mathsf{pred}(i)} T_i \cup B_i >_{\mathsf{pred}(i)} T_j \cup B_i$, where the second inequality is by Observation 2.8 and cancelability. Thus, $\mathsf{pred}(i)$ does not envy $T_j \cup B_i$ and the claim follows. $\square$

Proof of Observation 3.17. Since $i \not\rightarrow k$ and $i \xrightarrow{B_j \mid \circ} k$ we have $T_k \cup B_k = X_k \leq_i X_i <_i (X_k \setminus B_k) \cup B_j = T_k \cup B_j$, and by cancelability this implies $B_k <_i B_j$. $\square$

Proof of Corollary 3.18. Assume towards contradiction that there is some agent a which is the source of all B_i-edges given by Theorem 3.16, and a does not envy any agent along the cycle. Let $j = \arg\min_i \{v_a(B_i) \mid i \text{ is an agent along } C\}$. By assumption, there exists some j' along C such that $a \xrightarrow{B_j \mid \circ} j'$, and $a \not\rightarrow j'$. By Observation 3.17, $B_j >_a B_{j'}$. Hence, $v_a(B_j) > v_a(B_{j'})$, in contradiction to the definition of j. $\square$

C Simplification and Extension of Known Results

In this section we simplify the proofs of full EFX existence for 3 additive agents [CGM20] and for n agents with one of two fixed additive valuations [Mah20]. Moreover, our proofs extend beyond additive to all nice cancelable valuations. Our proofs demonstrate the versatility of our techniques.

C.1 EFX for 3 nice cancelable valuations

By Lemma 2.3 it is sufficient to prove the following theorem.

Theorem C.1. Let $\mathbf{X}$ be an EFX allocation for 3 agents with nice cancelable valuations, with at least one unallocated item. Then, there exists an EFX allocation $\mathbf{Y}$ that dominates $\mathbf{X}$.

By assumption there is an item g that is unallocated in $\mathbf{X}$. The original proof in [CGM20] distinguishes between two main cases according to whether $\mathbf{X}$ is envy-free or not. In the case where $\mathbf{X}$ is not envy-free the original proof extends almost immediately to nice cancelable valuations (with one exception, see below). The property of an additive valuation v that is applied there is that for any bundles S, T, R such that R is disjoint from both S and T, $v(S \cup R) > v(T \cup R) \Leftrightarrow v(S) > v(T)$, and this property also holds for nice cancelable valuations as pointed out in Section 2.

The one exception is in Section 4.2 in [CGM20] (right after Observation 16), in which their proof requires a subtle adjustment, as follows. In their sub-case "$a = 2$" they define the set Z_i, for $i = 1, 3$, to be a smallest subset of X_3 such that $Z_i >_i \max_i(X_1 \setminus G_{21} \cup g, X_2)$. In our case, the set Z_i needs to be defined as the subset of X_3 obtained by iteratively removing the item of least marginal value to agent i, as long as the leftover bundle is still better than $\max_i(X_1 \setminus G_{21} \cup g, X_2)$ from agent i's point of view. Lemma A.4 then guarantees that the new Z_i is indeed a smallest-size subset of X_3 as in the original proof. The rest of the argument then follows as in the original proof.

We now turn to the envy-free case. In the original proof the following property of an additive valuation v is used to prove the sub-case where $M_{\mathbf{X}}$ contains a g-cycle of size 2: if $v(S) > v(T)$ and $v(S) > v(R)$, where R and T are disjoint, then $2 \cdot v(S) > v(T \cup R)$. This property does *not* hold in general for nice cancelable valuations (*e.g.*, v is multiplicative and $v(S) = 5, v(R) = v(T) = 4$), thus their proof does not extend to nice cancelable valuations.

In what follows we present a proof, based on our new techniques, that applies to all nice cancelable valuations, and is also substantially simpler than the original proof.

Lemma C.2. Let $\mathbf{X}$ be a partial envy-free allocation on 3 agents. Then there exists an EFX allocation $\mathbf{Y}$ that Pareto dominates it.

Proof. By Lemma 3.5, it suffices to find a Pareto-improvable (PI) cycle in the champion graph $M_{\mathbf{X}}$. Since every agent has an incoming g-champion edge (Observation 2.5), $M_{\mathbf{X}}$ has a cycle of g edges. Let C be such a cycle of minimal length. If C is a self loop, then we are done by Corollary 3.6. Hence, it remains to consider the cases where C has length two or three. Since $\mathbf{X}$ is envy-free, C is a good g-cycle, and as such it induces a g-decomposition of X_j into top and bottom half-bundles $X_j = T_j \cup B_j$ for every agent j along the cycle.

Case 1: $|C| = 2$. Assume w.l.o.g. that $C = 1 \xrightarrow{g} 2 \xrightarrow{g} 1$. By Theorem 3.16, $M_{\mathbf{X}}$ contains a good or external B_1-edge and a good or external B_2-edge. By Corollary 3.18 they cannot both be external (since their source would be the same, agent 3). Thus, w.l.o.g. $M_{\mathbf{X}}$ contains a good B_1-edge, which can only be $1 \xrightarrow{B_1 \mid \circ} 2$, and we obtain the PI cycle $1 \xrightarrow{B_1 \mid \circ} 2 \xrightarrow{g} 1$.

Case 2: $|C| = 3$. W.l.o.g. let $C = 1 \xrightarrow{g} 2 \xrightarrow{g} 3 \xrightarrow{g} 1$. As C contains all 3 agents, there are no external edges going into C. Thus by Theorem 3.16, $M_{\mathbf{X}}$ contains a good B_i-edge, for each $i \in [3]$.

If for some $i \in [3]$ the good B_i-edge is $i \xrightarrow{B_i \mid \circ} \mathsf{pred}(i)$, then we get the PI cycle $i \xrightarrow{B_i \mid \circ} \mathsf{pred}(i) \xrightarrow{g} i$, and we are done. Thus we can assume that all good edges are parallel to the edges of C, i.e., from j to $\mathsf{succ}(j)$. W.l.o.g., assume there is a good edge from agent 1 to agent 2. This good edge cannot be the good B_2-edge, because $1 \xrightarrow{\cancel{B_2 \mid \circ}} 2$ by Observation 3.11. Thus, this good edge is $1 \xrightarrow{B_i \mid \circ} 2$, for some $i \in \{1, 3\}$, and the good

B_2-edge must be either $2 \xrightarrow{B_2 \mid \circ} 3$ or $3 \xrightarrow{B_2 \mid \circ} 1$. The former case admits the PI cycle $1 \xrightarrow{B_i \mid \circ} 2 \xrightarrow{B_2 \mid \circ} 3 \xrightarrow{g} 1$; the latter admits the PI cycle $1 \xrightarrow{B_i \mid \circ} 2 \xrightarrow{g} 3 \xrightarrow{B_2 \mid \circ} 1$. $\square$

C.2 EFX for agents with one of two valuations

Consider a setting with n agents, where any agent has one of two valuations v_a, v_b. Let $a_0, \ldots, a_t$ denote the agents with valuation v_a, and $b_0, \ldots, b_\ell$ denote the agents with valuation v_b, ordered such that

$$X_{a_0} \leq_a X_{a_1} \leq_a \ldots \leq_a X_{a_t} \qquad \text{and} \qquad X_{b_0} \leq_b X_{b_1} \leq_b \ldots \leq_b X_{b_\ell}.$$

The following theorem shows that if v_a and v_b are nice cancelable valuations, then given any partial EFX allocation, there exists an EFX allocation that Pareto dominates it. This implies (by Lemmata 2.3, A.1) that every instance in this setting admits a full EFX allocation.

Theorem C.3. In every setting with two nice cancelable valuations, given any partial EFX allocation, there exists an EFX allocation that Pareto dominates it.

Before presenting the proof, we present a useful observation. We say that envy (resp., most envious) *propagates backward* within the valuation class a if whenever some agent a_i envies a set S (resp., is most envious of a set S), then for every $j < i$, agent a_j envies S (resp., is most envious of S) as well. We say that championship propagates backward within the valuation class a in an analogous way. We define backward propagation within the valuation class b analogously. One can easily verify that envy propagates backward. The following observation shows that so does championship.

Observation C.4. Championship propagates backward within the same valuation class.

Proof. We prove the claim for valuation class a. The proof for valuation class b is analogous. We show that the relation most envious propagates backward; by extension, championship propagates backward as well. Suppose that for some $i \in \{0, \ldots, \ell\}$, agent a_i is most envious of a set S, and let D be the discard set of S with respect to a_i. This means that a_i envies $S \setminus D$. Since envy propagates backward, so does agent a_j. By definition of a discard set, no agent strongly envies $S \setminus D$. Therefore, a_j is most envious of S. $\square$

We are now ready to prove Theorem C.3.

Proof. Fix a partial EFX allocation, and let g be an unallocated item. By Corollary 3.6 we may assume that no agent is a g-self champion. We first claim that

$$a_0 \xrightarrow{g} b_0 \text{ and } b_0 \xrightarrow{g} a_0. \tag{8}$$

Indeed, if $b_j \xrightarrow{g} b_0$ for some j, then b_0 is a g-self champion by backward propagation. Thus, since b_0 must have a g-champion (Observation 2.5), then $a_j \xrightarrow{g} b_0$ for some j. Since championship propagates backward, $a_0 \xrightarrow{g} b_0$. By symmetry, $b_0 \xrightarrow{g} a_0$.

We may also assume that no agent envies a_0 (and similarly, b_0). Clearly, no a_j envies a_0. It remains to show that no b_j envies a_0. Indeed, if some agent b_j envies a_0, then b_0 envies a_0, and together with the fact that $a_0 \xrightarrow{g} b_0$, we have a Pareto-improvable cycle, so we are done by Lemma 3.5. Similarly, if any agent envies b_0, we have a Pareto-improvable cycle.

By Equation (8) and the assumption that no agent envies a_0 or b_0, $a_0 \xrightarrow{g} b_0 \xrightarrow{g} a_0$ is a good g-cycle, thus by Observation 2.6 the bundles of a_0 and b_0 decompose into top and bottom half-bundles. Let T_{a_0} and B_{a_0} (resp., T_{b_0} and B_{b_0}) be the top and bottom half-bundles of a_0 (resp., b_0), respectively.

We next argue that $a_0 \xrightarrow{B_{a_0} \mid \circ} b_0$. Since $a_0 \xrightarrow{g} b_0 \xrightarrow{g} a_0$ is a good cycle, by Theorem 3.16 there exists a good or external B_{a_0}-edge that goes into b_0. If this is a good B_{a_0}-edge then it can only be $a_0 \xrightarrow{B_{a_0} \mid \circ} b_0$. If this is an external B_{a_0}-edge, it can't be $b_j \xrightarrow{B_{a_0} \mid \circ} b_0$ since $b_0 \xrightarrow{\cancel{B_{a_0} \mid \circ}} b_0$ (by Observation 3.12) and championship propagates backward. Hence, the external B_{a_0}-edge must be $a_j \xrightarrow{B_{a_0} \mid \circ} b_0$ for some $j \in \{1, \ldots, t\}$. Again, since championship propagates backward, $a_0 \xrightarrow{B_{a_0} \mid \circ} b_0$.

It follows that we have a Pareto-improvable cycle consisting of $a_0 \xrightarrow{B_{a_0} \mid \circ} b_0$ and $b_0 \xrightarrow{g} a_0$. We may now apply Lemma 3.5 to conclude the proof.

□