

Constrained Fair and Efficient Allocations

Benjamin Cookson Soroush Ebadian Nisarg Shah

University of Toronto
{bcookson,soroush,nisarg}@cs.toronto.edu

Abstract

Fairness and efficiency have become the pillars of modern fair division research, but prior work on achieving both simultaneously is largely limited to the *unconstrained* setting. We study fair and efficient allocations of indivisible goods under additive valuations and various types of allocation feasibility constraints, and demonstrate the unreasonable effectiveness of the maximum Nash welfare (MNW) solution in this previously uncharted territory.

Our main result is that MNW allocations are $\frac{1}{2}$ -envy-free up to one good (EF1) and Pareto optimal under the broad family of (arbitrary) matroid constraints. We extend these guarantees to *complete* MNW allocations for base-orderable matroid constraints, and to a family of non-matroidal constraints (which includes balancedness) using a novel “alternate worlds” technique. We establish tightness of our results by providing counterexamples for the satisfiability of certain stronger desiderata, but show an improved result for the special case of goods with copies [Gafni et al., 2023]. Finally, we also establish novel best-of-both-worlds guarantees for goods with copies and balancedness.

1 Introduction

Fair division of resources among agents is a primitive that has applications, both to multiagent systems [Chevalyere et al., 2006] and to everyday problems such as estate division and divorce settlement [Shah, 2017]. Over the last decade, the fair division literature has undergone a dramatic transformation. The pioneering work of Caragiannis et al. [2019] established that, under additive valuations, the so-called maximum Nash welfare (MNW) allocations, which (informally) maximize the product of agent utilities, simultaneously satisfy two appealing guarantees: a fairness criterion known as *envy-freeness up to one good* (EF1), which demands that no agent prefer the allocation of another agent (modulo a single good) to her own, and an efficiency criterion known as *Pareto optimality* (PO), which demands that no alternative allocation be able to make an agent happier without making any agent worse off. These provable fairness and efficiency guarantees have been critical to their use in the real world via the not-for-profit website Spliddit.org [Shah, 2017].

Ever since then, the combination of fairness and efficiency, in the form of approximate envy-freeness and Pareto optimality, has become the guiding principle for seeking fair division solutions, e.g., for subclasses of additive valuations [Hosseini et al., 2021], for non-additive valuations [Benabbou et al., 2021, Barman and Suzuki, 2024], when addressing manipulations [Psomas and Verma, 2022], or for allocating chores [Ebadian et al., 2022, Garg et al., 2022], or for allocating public goods [Fain et al., 2018, Ebadian et al., 2024].

However, in many real-world fair division problems, there are feasibility constraints on the bundle that each agent can receive. Examples include course allocation [Budish et al., 2017], public housing assignment [Benabbou et al., 2020], or allocation of conference submissions to reviewers [Garg et al., 2010]. Unfortunately, the literature on *constrained* fair division has been largely limited to seeking only fairness guarantees.

- Biswas and Barman [2018] show the existence of an EF1 allocation subject to cardinality constraints, where the goods are partitioned into categories and each agent must be allocated at most a prescribed maximum number of goods from each category.
- Biswas and Barman [2019] extend this to any base-orderable matroid constraint, where the bundle of goods allocated to each agent must be an independent set of a given base-orderable matroid (cardinality constraints form a partition matroid, which is a special case), when agents have identical additive valuations.¹
- Gafni et al. [2023] study a special case of cardinality constraints, which they refer to as *goods with copies*, motivated by the fact that it in turn subsumes chore division as a special case. They define an appealing strengthening of EF1, termed EF1_{WC}, and establish its existence for restricted valuation classes.
- The popular round robin algorithm yields an EF1 allocation subject to *balancedness*, where all agents must be assigned bundles of roughly equal cardinality (differing by at most one) [Caragiannis et al., 2019].
- A famous non-matroidal constraint is where the goods are vertices of an undirected graph, and the bundle allocated to each agent must form a connected subset; when the graph is a line, an EF1 allocation is known to exist [Igarashi, 2023], and for general graphs, a $\frac{1}{2}$ -EF1 allocation with up to $n - 1$ unallocated goods is known to exist under restricted preferences [Caragiannis et al., 2022].

In addition to the above summary of related work, we provide comparisons to other pieces of related work throughout the paper, and also refer the reader to the extensive survey on constraints in fair division by Suksompong [2021]. Quite surprisingly, the existence of an allocation that satisfies both (even approximate) EF1 and PO remains severely understudied in these constrained domains. The only exception is the work on budget constraints, where each good has a size and the total size of goods allocated to any agent must be at most a threshold. Gan et al. [2023] show that an EF1 allocation always exists, and Wu et al. [2021] show that any MNW allocation subject to such a constraint is $\frac{1}{4}$ -EF1 and PO. Our main research question is to expand on this line of work:

Under which types of feasibility constraints do (approximately) fair and efficient allocations exist?

1.1 Our Results

Our main result is that every MNW allocation is $\frac{1}{2}$ -EF1 and PO under the broad class of (arbitrary) matroid constraints (see Section 2 for a formal definition). While these allocations are efficient,

¹Biswas and Barman [2019] incorrectly state their result for an arbitrary matroid constraint in the original paper, but later versions correctly state that the result holds for base-orderable matroids. The existence of an EF1 allocation here remains open for general matroid constraints, even for identical additive valuations.

they can be *incomplete*, i.e., leave some goods unallocated, which may be undesirable in settings such as allocation of shifts to nurses and assignment of conference submissions to reviewers. To that end, we show that for base-orderable matroid constraints, even allocations that are MNW among the set of *complete* and feasible allocations are $\frac{1}{2}$ -EF1 and PO. Base-orderable matroids subsume the case of cardinality constraints [Biswas and Barman, 2018].

Then, using a novel technique of constructing “alternate worlds”, we show that the $\frac{1}{2}$ -EF1 and PO guarantees can be extended to MNW allocations subject to a broad class of non-matroidal constraints, which includes balancedness as a special case. We also show that certain strengthenings of EF1 and PO are unachievable in the realm of constrained allocations, but show an improvement from EF1 to EF1_{WC} for the case of goods with copies [Gafni et al., 2023].

Finally, we expand the recent work on “best of both worlds” (BoBW) guarantees [Aziz et al., 2024] to the realm of constrained allocations. Building on the work of Echenique et al. [2021], we prove that randomized allocations that are ex ante EF and PO along with ex post EF₁¹, Prop1, and PO exist for the case of goods with copies, and the same result except for the EF₁¹ guarantee holds when adding the balancedness constraint. For formal definitions, see Section 4.

2 Preliminaries

For any $r \in \mathbb{N}$, define $[r] := \{1, 2, \dots, r\}$. Let $N = [n]$ be a set of *agents*, and M be a set of m (*indivisible*) goods. Each agent i has a valuation function $v_i : M \rightarrow \mathbb{R}_{\geq 0}$, where $v_i(g)$ is her value for good g . We assume additive valuations: with slight abuse of notation, the value of agent i for a set of goods $S \subseteq M$ is $v_i(S) := \sum_{g \in S} v_i(g)$. Let $v = (v_1, \dots, v_n)$ be the *valuation profile*. In an *allocation* $A = (A_1, \dots, A_n)$, $A_i \subseteq M$ is the bundle of goods assigned to agent i and $A_i \cap A_j = \emptyset$ for all distinct $i, j \in N$. We say that A is *complete* if $\cup_{i \in N} A_i = M$. The utility to agent i under this allocation is $v_i(A_i)$.

2.1 Feasibility Constraints

We are interested in *constrained* allocation problems, where we are allowed to choose an allocation only from a given set of allocations \mathcal{F} . One of the most general family of constraints uses matroids.

Matroid constraints. We are given a matroid over the goods (M, \mathcal{I}) satisfying: (a) $\emptyset \in \mathcal{I}$; (b) (*hereditary property*) if $S \in \mathcal{I}$ and $S' \subseteq S$, then $S' \in \mathcal{I}$; and (c) (*exchange property*) if $S, T \in \mathcal{I}$ and $|T| > |S|$, then there exists $g \in T \setminus S$ such that $S \cup \{g\} \in \mathcal{I}$. Then, an allocation A is called *feasible* (i.e., $A \in \mathcal{F}$) if $A_i \in \mathcal{I}$ for all agents $i \in N$. Bases of a matroid are its maximal independent sets. The exchange property implies that all bases have the same cardinality (known as the rank of the matroid), and for any two bases S and T , there exist $g \in S \setminus T$ and $g' \in T \setminus S$ such that $S \cup \{g'\} \setminus \{g\}$ and $T \cup \{g\} \setminus \{g'\}$ are both bases as well. We will reference the following popular families of matroids in this paper.

- *Base-orderable matroids.* These are matroids with a strengthened base exchange property: for any two bases S and T , there exists a *bijection* $f : S \rightarrow T$ such that, for every $g \in S$, $S \cup \{f(g)\} \setminus \{g\}$ and $T \cup \{g\} \setminus \{f(g)\}$ are both bases as well.
- *Laminar matroids.* This is a special case of base-orderable matroids, where we are given a laminar family of subsets of goods (termed *categories*) $\mathcal{C} \subseteq 2^M$. Here, laminar means that for

every $C, C' \in \mathcal{C}$, we have $C \cap C' \in \{\emptyset, C, C'\}$. For each category $C \in \mathcal{C}$, we are also given an upper bound h_C . Then, \mathcal{I} consists of all sets S such that $|S \cap C| \leq h_C$ for every category $C \in \mathcal{C}$. For feasibility, we assume $h_C \geq \lceil |C|/n \rceil$ for each $C \in \mathcal{C}$.

- *Partition matroids.* This is a special case of laminar matroids, where \mathcal{C} is a partition of the set of goods M , i.e., $C \cap C' = \emptyset$ for all $C, C' \in \mathcal{C}$ and $\cup_{C \in \mathcal{C}} C = M$. This is also known in the fair division literature as *cardinality constraints* [Biswas and Barman, 2018].
- *Goods with copies.* Introduced by Gafni et al. [2023], this is a special case of partition matroids, where $h_C = 1$ and $|C| \leq n$ for each category $C \in \mathcal{C}$. The typical motivation for this is when the goods in each category are copies (i.e., perfect substitutes) of each other, i.e., $v_i(g) = v_i(g')$ for all $i \in N, C \in \mathcal{C}$, and $g, g' \in C$. This is what we will assume when referring to the case of goods with copies.

Non-matroidal constraints. We also study the following constraints that do not form a matroid.

- *Partition matroids with lower bounds.* Here, we have a partition matroid induced by the partition \mathcal{C} of M and upper bounds $(h_C)_{C \in \mathcal{C}}$. Additionally, we also have lower bounds $(\ell_C)_{C \in \mathcal{C}}$. Then, \mathcal{F} consists of the set of allocations A where $\ell_C \leq |A_i \cap C| \leq h_C$ for every $i \in N$ and $C \in \mathcal{C}$. For feasibility, we assume $\ell_C \leq \lfloor |C|/n \rfloor \leq \lceil |C|/n \rceil \leq h_C$ for each $C \in \mathcal{C}$.
- *Balanced.* This is a special case of partition matroids with lower bounds with a singleton $\mathcal{C} = \{M\}$, $\ell_M = \lfloor m/n \rfloor$ and $h_M = \lceil m/n \rceil$. That is, balanced allocations A satisfy $|A_i| \in \{\lfloor m/n \rfloor, \lceil m/n \rceil\}$ for all $i \in N$, so the number of goods allocated to any two agents differ by at most one.

2.2 Maximum Nash Welfare

Caragiannis et al. [2019] introduced the *maximum Nash welfare* rule and proved that, given any unconstrained fair division instance with additive valuations, it always returns an EF1+PO allocation (thus settling the open question of the existence of such allocations). The following is the natural adaptation of their rule to the *constrained* case.

Definition 1 ((Constrained) Maximum Nash Welfare (MNW)). For an allocation A , define $P(A) = \{i \in N : v_i(A_i) > 0\}$ to be the set of agents receiving a positive utility. An allocation $A \in \mathcal{F}$ is a *(constrained) maximum Nash welfare* (MNW) allocation if it maximizes the number of agents receiving positive utility, $|P(A)|$, and, subject to that, maximizes the product of positive agent utilities (Nash welfare), $NW(A) := \prod_{i \in P(A)} v_i(A_i)$.²

When the above definition is applied to the set of *complete and feasible* allocations, we refer to the resulting allocations as *complete MNW allocations*. Note that a complete MNW allocation may not necessarily be an MNW allocation (among the set of all feasible allocations); see [Example 1](#).

²Note that in the typical case where there exists an allocation yielding a positive utility to every agent, this simplifies to saying that A is an MNW allocation if it maximizes $\prod_{i \in N} v_i(A_i)$.

2.3 Fairness and Efficiency Desiderata

We are interested in allocations that satisfy various desiderata. The two that play a key role are the following.

Definition 2 (Approximate envy-freeness up to one good (EF1)). For $\alpha \in [0, 1]$, an allocation $A \in \mathcal{F}$ is α -approximate *envy-free up to one good* (α -EF1) if, for every pair of agents $i, j \in N$, either $A_j = \emptyset$ or $v_i(A_i) \geq \alpha \cdot v_i(A_j \setminus \{g\})$ for some $g \in A_j$. In words, agent i should not envy agent j , up to a factor of α , after excluding some good from agent j 's bundle. When $\alpha = 1$, we simply call it an EF1 allocation.

Definition 3 (Approximate (constrained) Pareto optimality (PO)). For $\alpha \in [0, 1]$, an allocation $A \in \mathcal{F}$ is α -approximate (*constrained*) *Pareto optimal* (α -PO) if there is no other allocation $B \in \mathcal{F}$ such that $v_i(B_i) \geq \alpha \cdot v_i(A_i)$ for each $i \in N$ and at least one inequality is strict. In words, no other feasible allocation should be able to make an agent happier without making any agent less happy. When $\alpha = 1$, we simply call it a PO allocation.

We also study relaxations and strengthenings of these two desiderata, but we define them in their respective sections.

3 The Unreasonable Fairness of (Constrained) Maximum Nash Welfare

A common technique to deal with feasibility constraints is to reduce the problem to an unconstrained instance by modifying the valuation of each agent for any bundle S to be her highest value for any feasible subset $T \subseteq S$, computing a desirable allocation A for the unconstrained instance, and then giving to each agent i her most valuable feasible subset of A_i . [Biswas and Barman \[2018\]](#) observed that such a reduction for partition matroid constraints induces submodular valuations.³ However, this actually holds true for any matroid constraint due to known results from matroid theory. In addition to this technique, we use a result by [Caragiannis et al. \[2019\]](#) that every MNW allocation for an unconstrained instance with submodular valuations satisfies a relaxation of EF1 that they term *marginal envy-freeness up to one good* (MEF1), and observe that MEF1 implies $\frac{1}{2}$ -EF1.

Theorem 1. *Under any matroid constraint, every MNW allocation is $\frac{1}{2}$ -EF1 and PO.*

Proof. For this, we can reduce the constrained problem with additive valuations to an unconstrained problem with non-additive valuations using a popular trick (see, e.g., the works of [Biswas and Barman \[2018\]](#) and [Dror et al. \[2023\]](#)). Consider an instance with additive valuations $v = (v_1, \dots, v_n)$ and matroid constraint (M, \mathcal{I}) . In the new unconstrained instance, the valuation function of agent i is given by $\tilde{v}_i(S) = \max_{T \subseteq S, T \in \mathcal{I}} v_i(T)$; that is, each agent values S by the best feasible bundle within S . This construction is known as a *weighted rank function* of a matroid with $v_i(g)$ being the weight of element g , and it is known that any weighted rank function of any matroid is monotone submodular [[Schrijver, 2003](#)].

Next, consider any MNW allocation A in the constrained instance. We consider A as an allocation in the unconstrained instance. Next, we show that A is also an MNW allocation in the unconstrained instance. Suppose this is not true. Then, there exists an MNW allocation \tilde{B} in the unconstrained instance such that one of two things happens:

³Valuation v is submodular if $v(S \cup T) + v(S \cap T) \leq v(S) + v(T)$ for all $S, T \subseteq M$.

- (1) $|P(\tilde{B})| = |\{i \in N : \tilde{v}_i(\tilde{B}_i) > 0\}| > |\{i \in N : \tilde{v}_i(A_i) > 0\}| = |P(A)|$
- (2) $P(\tilde{B}) = P(A) = P$ and $\prod_{i \in P} \tilde{v}_i(\tilde{B}_i) > \prod_{i \in P} \tilde{v}_i(A_i)$.

Generate an allocation B in the constrained instance by giving to each agent i her most valuable feasible subset of \tilde{B}_i . Notice that $\tilde{v}_i(A_i) = v_i(A_i)$ and $\tilde{v}_i(\tilde{B}_i) = v_i(B_i)$ for each $i \in N$. Now, consider allocations A and B in the constrained instance. In the first case above, we would get $|\{i \in N : v_i(B_i) > 0\}| > |\{i \in N : v_i(A_i) > 0\}|$. And in the second case above, we would get $\{i \in N : v_i(B_i) > 0\} = \{i \in N : v_i(A_i) > 0\} = P$ (say) and $\prod_{i \in P} v_i(B_i) > \prod_{i \in P} v_i(A_i)$. Both would violate the fact that A is an MNW allocation in the constrained instance.

Finally, [Caragiannis et al. \[2019\]](#) show that every MNW allocation in an unconstrained instance with submodular valuations $\tilde{v} = (\tilde{v}_1, \dots, \tilde{v}_n)$, including allocation A above, satisfies marginal envy-freeness up to one good (MEF1): for every pair of agents $i, j \in N$, either $A_j = \emptyset$ or there exists $g \in A_j$ such that $\tilde{v}_i(A_i \cup A_j \setminus \{g\}) - \tilde{v}_i(A_i) \leq \tilde{v}_i(A_j)$. For monotone submodular valuations, it is easy to see that MEF1 implies $\frac{1}{2}$ -EF1: $\tilde{v}_i(A_j \setminus \{g\}) \leq 2\tilde{v}_i(A_j)$. Noticing that $\tilde{v}_i(A_i) = v_i(A_i)$ and $\tilde{v}_i(A_j) = v_i(A_j)$ yields that A is $\frac{1}{2}$ -EF1 in the constrained instance too.

The fact that A is PO is established similarly. If there exists an allocation B of the constrained instance that Pareto dominates A , then, even in the unconstrained instance, B would Pareto dominate A , contradicting the fact that A , being an MNW allocation in the unconstrained instance, is PO [[Caragiannis et al., 2019](#)]. \square

An MNW allocation from [Theorem 1](#) may be incomplete. At first, one may worry that such an allocation can be highly inefficient. However, note that not even a complete allocation can Pareto dominate an MNW allocation, nor have a higher Nash welfare. As a consequence, every agent i must have zero value for every unallocated good g that can be feasibly added to A_i (otherwise adding g to A_i would yield a Pareto improvement), and no agent i can envy any feasible bundle from the unallocated goods (otherwise swapping such a bundle with A_i would yield a Pareto improvement).

Nonetheless, as motivated in [Section 1](#), there are settings in which complete allocations are desirable because free disposal of goods may not be allowed. *Does every matroid constrained instance admit a complete allocation that is $\frac{1}{2}$ -EF1 and PO?* The only reason [Theorem 1](#) does not already answer this positively is that there may exist an instance in which *every* MNW allocation is incomplete. The following example shows that such instances indeed exist!

Example 1. Consider an instance with 2 agents, 8 goods (g_1 through g_8), and the following laminar matroid constraint:

- Category $C_1 = \{g_1, g_2, g_3, g_4\}$ has upper bound $h_{C_1} = 2$.
- Category $C_2 = \{g_1, g_2, g_3, g_4, g_5, g_6, g_7, g_8\}$ has upper bound $h_{C_2} = 4$.

The valuations of the agents are as follows.

| | g_1 | g_2 | g_3 | g_4 | g_5 | g_6 | g_7 | g_8 |
|---------|-------|-------|-------|-------|-------|-------|-------|-------|
| Agent 1 | 0 | 1 | 0 | 0 | 1 | 1 | 1 | 0 |
| Agent 2 | 0 | 0 | 1 | 1 | 0 | 0 | 0 | 1 |

Note that there is a unique MNW allocation:

$$A_1 = \{g_2, g_5, g_6, g_7\}, A_2 = \{g_3, g_4, g_8\}.$$

In the above example, notice that every *complete* allocation is Pareto dominated by the incomplete MNW allocation. A similar result is observed by [Benabbou et al. \[2021\]](#), where they show that for binary submodular valuations, some instances have no allocation that maximizes utilitarian welfare that are both complete and “clean” (where clean means that no agent’s bundle includes a good which that agent has 0 marginal utility for). Under the paradigm of reducing instances under matroid constraints to submodular valuations, one can easily see that in order for the allocation in the original instance to be complete, the allocation over the reduced submodular instance must be complete and clean (with respect to all goods that do not have 0 additive value in the original instance).

This is why completeness should only be sought if free disposal is indeed impossible. Also, due to the above example, in what follows, Pareto optimality of complete allocations would mean only that they are not Pareto dominated by other *complete* and feasible allocations.

We now show that at least for the broad family of *base-orderable* matroid constraints, *complete MNW allocations* — recall that these are MNW allocations among the set of complete and feasible allocations — are also $\frac{1}{2}$ -EF1 and PO.

Theorem 2. *Under any base-orderable matroid constraints, every complete MNW allocation is $\frac{1}{2}$ -EF1 and PO.*

Proof. Consider any instance with additive valuations $v = (v_1, \dots, v_n)$ and a base-orderable matroid constraint (M, \mathcal{I}) . Let A be any complete MNW allocation.

To see that A is PO, suppose for contradiction that a complete allocation B Pareto dominates A . Then, B must either give a positive utility to strictly more agents ($|P(B)| > |P(A)|$) or give a positive utility to the same set of agents ($P(B) = P(A)$) while increasing their product of utilities ($NW(B) > NW(A)$). Both possibilities contradict the fact that A is a complete MNW allocation.

Next, we show that A is $\frac{1}{2}$ -EF1. Suppose for contradiction that there exist agents $i, j \in N$ such that $A_j \neq \emptyset$ and $2 \cdot v_i(A_i) < v_i(A_j) - v_i(g)$ for all $g \in A_j$.

Due to Observation 2 from [Dror et al. \[2023\]](#), we can assume without loss of generality that A_i is a basis of the matroid (M, \mathcal{I}) for each $i \in N$. This is achieved by performing a preprocessing step of adding dummy goods to the instance. The details of this step along with a proof that complete MNW allocations in the original instance and in the preprocessed instance are equivalent is provided in [Appendix B](#).

Now, since A_i and A_j are both bases, let $f : A_j \rightarrow A_i$ be the bijection from the definition of base-orderability of (M, \mathcal{I}) . Define $A_j^* := \{g \in A_j \mid v_i(g) > v_i(f(g))\}$ to be the set of all goods $g \in A_j$ such that agent i ’s utility would increase if g and $f(g)$ were swapped between A_i and A_j . Define its projection $A_i^* := \{f(g) \mid g \in A_j^*\}$. Then:

$$\begin{aligned} & \sum_{g \in A_j^*} (v_i(g) - v_i(f(g))) \\ &= v_i(A_j^*) - v_i(A_i^*) \geq v_i(A_j) - v_i(A_i) \\ &> v_i(A_i) + \max_{g' \in A_j^*} (v_i(g') - v_i(f(g'))), \end{aligned} \tag{1}$$

where the second transition follows because

$$v_i(A_j) - v_i(A_i) = (v_i(A_j^*) - v_i(A_i^*)) + (v_i(A_j \setminus A_j^*) - v_i(A_i \setminus A_i^*)),$$

and

$$v_i(A_j \setminus A_j^*) - v_i(A_i \setminus A_i^*) = \sum_{g \in A_j \setminus A_j^*} v_i(g) - v_i(f(g)) \leq 0;$$

and the third transition follows due to the assumed violation of $\frac{1}{2}$ -EF1.

Now, we can observe that

$$\sum_{g \in A_j^*} (v_j(g) - v_j(f(g))) = v_j(A_j^*) - v_j(A_i^*) \leq v_j(A_j^*) \leq v_j(A_j). \quad (2)$$

Note that for every $g \in A_j^*$, not only is $v_i(g) > v_i(f(g))$ (due to the definition of A_j^*), but also $v_j(g) > v_j(f(g))$. If this were not true, then swapping g and $f(g)$ between A_i and A_j , which yields a complete and feasible allocation due to $A_i \cup \{g^*\} \setminus \{f(g^*)\}$ and $A_j \setminus \{g^*\} \cup \{f(g^*)\}$ also being bases of (M, \mathcal{I}) , would be a Pareto improvement over A . Choosing $g^* \in \arg \min_{g \in A_j^*} \frac{v_j(g) - v_j(f(g))}{v_i(g) - v_i(f(g))}$, we have

$$\begin{aligned} \frac{v_j(A_j)}{v_i(A_i) + v_i(g^*) - v_i(f(g^*))} &> \frac{\sum_{g \in A_j^*} (v_j(g) - v_j(f(g)))}{\sum_{g \in A_j^*} (v_i(g) - v_i(f(g)))} \\ &\geq \frac{v_j(g^*) - v_j(f(g^*))}{v_i(g^*) - v_i(f(g^*))} \\ &= \frac{v_j(A_j) - v_j(A_j \setminus \{g^*\} \cup \{f(g^*)\})}{v_i(A_i \cup \{g^*\} \setminus \{f(g^*)\}) - v_i(A_i)}, \end{aligned}$$

which can be rearranged to get:

$$v_i(A_i \cup \{g^*\} \setminus \{f(g^*)\}) \cdot v_j(A_j \setminus \{g^*\} \cup \{f(g^*)\}) > v_i(A_i) \cdot v_j(A_j).$$

Consider allocation B where $B_i = A_i \cup \{g^*\} \setminus \{f(g^*)\}$, $B_j = A_j \setminus \{g^*\} \cup \{f(g^*)\}$, and $B_k = A_k$ for $k \in N \setminus \{i, j\}$. Note that B is also a complete allocation. If either i or j had zero utility under A , then B gives a positive utility to strictly more agents ($|P(B)| > |P(A)|$), and if both had a positive utility under A , then B has a strictly higher product of positive agent utilities ($NW(B) > NW(A)$). In either case, it contradicts A being a complete MNW allocation. This proves that A must be $\frac{1}{2}$ -EF1 too. \square

3.1 Beyond Matroidal Constraints

One may be tempted to apply the trick from [Theorem 1](#) of reducing to unconstrained instances for even non-matroidal constraints. Even if this induces valuations from the broader class of subadditive valuations,⁴ one can use a recent result by [Barman and Suzuki \[2024\]](#) to establish the existence of EF1 and $\frac{1}{2}$ -PO allocations. However, even for very simple non-matroidal constraints, it is easy to see that the induced valuations would not even be subadditive.

⁴Valuation v is subadditive if $v(S \cup T) \leq v(S) + v(T)$ for all $S, T \subseteq M$.

For example, consider partition matroids with lower bounds, where the number of goods allocated to every agent from each category C must be in between a lower bound ℓ_C and an upper bound h_C for category C . Now, an agent's value for a bundle of goods can stay zero while the bundle contains fewer than ℓ_C goods from any category C , and suddenly jump to a positive value once it contains at least ℓ_C goods from every category C , violating subadditivity.

Nonetheless, for the case of partition matroids with lower bounds, we develop a novel approach of *alternate worlds construction* to establish that complete MNW allocations are $\frac{1}{2}$ -EF1 and PO. The following lemma lays out the approach.

Lemma 1. *Let $\{\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k\}$ be a collections of sets of feasible allocations (“alternate worlds”), defined over the same set of goods M . Fix any $\alpha \in [0, 1]$. If every MNW allocation among \mathcal{F}_t is α -EF1 and PO for each $t \in [k]$, then every MNW allocation among $\mathcal{F} = \cup_{t \in [k]} \mathcal{F}_t$ will be α -EF1 and PO.*

Proof. Consider any allocation A that is MNW among $\mathcal{F} = \cup_{t \in [k]} \mathcal{F}_t$. The fact that it is PO among \mathcal{F} follows through the same argument as for the unconstrained case (see the proof of [Theorem 2](#)). The fact that A is α -EF1 follows from the fact that A must be an MNW allocation among \mathcal{F}_t for some $t \in [k]$ (i.e., in one of the alternate worlds). For contradiction, assume that this is false. Since $\mathcal{F} = \cup_{t \in [k]} \mathcal{F}_t$, we must have $A \in \mathcal{F}_t$ for some $t \in [k]$. Then, the fact that A is not an MNW allocation among \mathcal{F}_t implies that there exists another allocation $B \in \mathcal{F}_t$ which either gives positive utility to strictly more agents or gives positive utility to the same set of agents as A but has a higher product of their utilities. Since $B \in \mathcal{F}$, this would also contradict A being an MNW allocation among \mathcal{F} . \square

We can apply [Lemma 1](#) to [Theorems 1](#) and [2](#) to obtain $\frac{1}{2}$ -EF1 and PO for a broad class of non-matroidal constraints obtained as “unions of matroid constraints”. Note that, usually, the union of matroids (M, \mathcal{I}_1) and (M, \mathcal{I}_2) is defined as (M, \mathcal{I}) , where $\mathcal{I} = \{I_1 \cup I_2 \mid I_1 \in \mathcal{I}_1, I_2 \in \mathcal{I}_2\}$, which is always a matroid [[Oxley, 2011](#)]. The union we are referencing is different: if \mathcal{F}_1 and \mathcal{F}_2 are the sets of feasible allocations induced by matroid constraints (M, \mathcal{I}_1) and (M, \mathcal{I}_2) respectively, then the “union” of these constraints would have the set of feasible allocations $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2$, which may not be induced by any matroid constraint.

Corollary 1. *Given a set of feasible allocations $\mathcal{F} = \cup_{t=1}^k \mathcal{F}_t$, where \mathcal{F}_t is the set of feasible allocations under a matroid constraint for each $t \in [k]$, every MNW allocation from \mathcal{F} is $\frac{1}{2}$ -EF1 and PO. If \mathcal{F}_t is the set of complete and feasible allocations under a base-orderable matroid constraint for each $t \in [k]$, then every complete MNW allocation from \mathcal{F} is $\frac{1}{2}$ -EF1 and PO.*

Next, we show that for *complete* allocations, the non-matroidal constraint imposed by a partition matroid with lower bounds can be represented as a union of partition matroid constraints (which are base-orderable).

Lemma 2. *Let \mathcal{F} be the set of complete and feasible allocations under a partition matroid constraint with lower bounds. Then, $\mathcal{F} = \cup_{t=1}^k \mathcal{F}_t$, where \mathcal{F}_t is the set of complete and feasible allocations under a partition matroid constraint for each $t \in [k]$.*

Proof. Let \mathcal{C} be the set of categories with associated lower bounds $(\ell_C)_{C \in \mathcal{C}}$ and upper bounds $(h_C)_{C \in \mathcal{C}}$ in a partition matroid constraint with lower bounds. Let \mathcal{F} be the set of complete and feasible allocations under this constraint. Note that in any allocation in \mathcal{F} , from each category C , some subset of $n \cdot \ell_C$ goods are divided evenly among the agents (so each agent receives ℓ_C goods),

and the rest are divided so that each agent receives at most $h_C - \ell_C$ of them. This inspires the following construction.

We create $\prod_{C \in \mathcal{C}} \binom{|C|}{n \cdot \ell_C}$ many partition matroids, one for each possible way of selecting $Q_C \subseteq C$ of size $n \cdot \ell_C$ from each $C \in \mathcal{C}$. In the partition matroid corresponding to a given choice of $Q = (Q_C)_{C \in \mathcal{C}}$, there are two categories for every category $C \in \mathcal{C}$: $C_Q^1 = Q_C$ with an upper bound of $h_{C_Q^1} = \ell_C$ and $C_Q^2 = C \setminus Q_C$ with an upper bound of $h_C - \ell_C$. The key observation is that under any *complete* and feasible allocation subject to this partition matroid, goods in C_Q^1 will be divided exactly evenly (so that each agent receives ℓ_C of them) while goods in C_Q^2 will be divided so that each agent receives at most $h_C - \ell_C$ of them. This immediately yields that any complete and feasible allocation under any such partition matroid is a complete and feasible allocation in \mathcal{F} .

To see the converse, take any $A \in \mathcal{F}$. For each $C \in \mathcal{C}$, define $Q_C^1 = \cup_{i \in N} Q_{C,i}$, where $Q_{C,i}$ is an arbitrary subset of $A_i \cap C$ of size ℓ_C . Note that the $n \cdot \ell_C$ goods in Q_C^1 are divided evenly among the agents, with each receiving exactly ℓ_C of them. Also, since $|A_i \cap C| \leq h_C$, each agent must receive at most $h_C - \ell_C$ goods from $Q_C^2 = C \setminus Q_C^1$. Thus, A is a complete and feasible allocation under the partition matroid corresponding to $Q = (Q_C)_{C \in \mathcal{C}}$. \square

Plugging in [Lemma 2](#) into [Corollary 1](#) yields $\frac{1}{2}$ -EF1 and PO of complete MNW allocations subject to partition matroids with lower bounds. It is not clear if we can decompose the set of all (possibly incomplete) feasible allocations subject to partition matroids with lower bounds in the same manner, which would allow us to directly apply [Corollary 1](#) for MNW allocations, as our proof of [Lemma 2](#) crucially uses completeness. Nonetheless, we show that $\frac{1}{2}$ -EF1 and PO guarantees can be extended to MNW allocations too.

Theorem 3. *Under any partition matroid constraint with lower bounds, every MNW and complete MNW allocation is $\frac{1}{2}$ -EF1 and PO.*

Proof. The result for complete MNW allocations (with Pareto optimality only within the set of all complete and feasible allocations) follows from [Lemma 2](#) and [Corollary 1](#).

Let \mathcal{C} be the set of categories with associated lower bounds $(\ell_C)_{C \in \mathcal{C}}$ and upper bounds $(h_C)_{C \in \mathcal{C}}$ in any partition matroid constraint with lower bounds, and let A be any (possibly incomplete) MNW allocation in that instance. The fact that A is PO (within the set of all feasible allocations) follows exactly as in the proof of [Theorem 1](#) (as it does not use the matroid property of the constraint).

Let us now show that A is $\frac{1}{2}$ -EF1. First, we extend A to a complete allocation \tilde{A} by arbitrarily allocating the unallocated goods while keeping the allocation feasible. This is possible for partition matroids with lower bounds because our assumption of $\ell_C \leq \lfloor |C|/n \rfloor \leq \lceil |C|/n \rceil \leq h_C$ for each $C \in \mathcal{C}$ means that whenever there is an unallocated good from any category, there must be an agent who can feasibly receive it.

Note that \tilde{A} must be a complete MNW allocation. This uses the same argument as in the proof of [Theorem 1](#) too: if \tilde{A} was not a complete MNW allocation, we could find a complete MNW allocation B that would either give positive utility to more agents or to the same set of agents but with a higher Nash welfare than that of \tilde{A} . Since \tilde{A} and A induce the same utilities, this would also contradict A being an MNW allocation. Hence, $\frac{1}{2}$ -EF1 of \tilde{A} follows from above.

To see that A is $\frac{1}{2}$ -EF1, fix any pair of agents $i, j \in N$. We want to show that either $A_j = \emptyset$ or $2v_i(A_i) \geq v_i(A_j) - v_i(g)$ for some $g \in A_j$. From $\frac{1}{2}$ -EF1 of \tilde{A} , we have that either $\tilde{A}_j = \emptyset$ or

$2v_i(\tilde{A}_i) \geq v_i(\tilde{A}_j) - v_i(g^*)$ for some $g^* \in \tilde{A}_j$. In the former case, $A_j \subseteq \tilde{A}_j$ would imply $A_j = \emptyset$. And in the latter case, we have

$$2v_i(A_i) = 2v_i(\tilde{A}_i) \geq v_i(\tilde{A}_j \setminus \{g^*\}) \geq v_i(A_j \setminus \{g^*\}),$$

where the first transition holds because \tilde{A} cannot be a Pareto improvement over A (A is PO) and the last transition holds again because $A_j \subseteq A_j^*$. \square

3.2 Implications

Implication for partition matroids. For partition matroids (cardinality constraints), [Biswas and Barman \[2018\]](#) show that an EF1 allocation always exists, whereas [Theorem 1](#) shows that MNW allocations achieve $\frac{1}{2}$ -EF1 and PO, sacrificing some fairness for added efficiency.

Implications for balancedness. Balancedness is the special case of partition matroids with lower bounds, where all bundle sizes are almost equal: $|A_i| \in \{\lfloor m/n \rfloor, \lceil m/n \rceil\}$ for all $i \in N$. It is known that round robin achieves exact EF1 subject to balancedness [[Caragiannis et al., 2019](#)], but we provide an example in [Appendix A.1](#) showing that it is inefficient. In contrast, [Theorem 3](#) shows that MNW allocations subject to balancedness are $\frac{1}{2}$ -EF1 and PO, sacrificing some fairness in favor of increased efficiency, as in the case of partition matroids. To the best of our knowledge, this is the first result about balanced MNW allocations.

Implication for goods with copies. As this is a special case of partition matroids, the same implication holds for MNW allocations. However, for this case, [Gafni et al. \[2023\]](#) argue that a more suitable notion of fairness is EF1_{WC} , which is a strengthening of EF1. This is because a chore division instance with additive costs can be modeled as a goods with copies instance with additive values and $n - 1$ copies of a good corresponding to every chore: EF1_{WC} in the goods with copies instance would imply EF1 for the underlying chore division instance, while EF1 for the goods with copies instance would not.

Definition 4 ((Approximate) Envy-Freeness Up to One Good Without Commons (EF1_{WC})). For $\alpha \in [0, 1]$, an allocation $A \in \mathcal{F}$ is α -approximate *envy-freeness up to one good without commons* (α - EF1_{WC}) if, for every pair of agents $i, j \in N$, either $v_i(A_i) \geq \alpha \cdot v_i(A_j)$ or there exists a category $C \in \mathcal{C}$ and a good $g \in C$ such that $A_i \cap C = \emptyset$ and $v_i(A_i) \geq \alpha \cdot v_i(A_j \setminus \{g\})$. In words, agent i should not envy agent j , up to a factor of α , after excluding some good from agent j 's bundle *that agent i does not also have a copy of*. When $\alpha = 1$, we simply call it an EF1_{WC} allocation.

The fact that we can only exclude a good from A_j of which agent i does not already have a copy makes α - EF1_{WC} stronger than α -EF1. We already know from [Theorems 1](#) and [2](#) that every MNW and complete MNW allocation for a goods with copies instance achieves $\frac{1}{2}$ -EF1. We can strengthen this to show that they in fact achieve $\frac{1}{2}$ - EF1_{WC} .

First, we need the following lemma, showing that for goods with copies (in fact, even for partition matroids), completeness comes without a price for MNW allocations.

Lemma 3. *Under any partition matroid constraint, every complete MNW allocation is an MNW allocation (among the set of all feasible allocations).*

Proof. As part of the proof of [Theorem 3](#), we proved that every MNW allocation under a partition matroid constraint with lower bounds can be extended to a complete MNW allocation. This proof is similar, but in the opposite direction.

For contradiction, assume this is false, and there exists some complete MNW allocation A for a instance constrained by the partition matroid constraints \mathcal{C} that is not MNW over all allocations. In this case, we know that there can be no complete allocations that maximize Nash Welfare for this instance, or that would contradict the fact that A is complete MNW. Therefore, all MNW allocations must be incomplete, let A' be the MNW allocation that minimizes the number of unallocated goods.

Since we know that \mathcal{C} admits at least 1 complete allocation, it must be the case that for each constraint $C_t \in \mathcal{C}$, we have that $h_t \geq \lceil |C_t|/n \rceil$. Since A' is not complete, there must be some constraint $C' \in \mathcal{C}$ and some good $g \in C'$ such that $g \notin A'_i$ for all $i \in N$. Further, it must be the case that for some agent $j \in N$, $|A'_j \cap C'| < \lceil |C_t|/n \rceil$. If not then there would be $\lceil |C_t|/n \rceil n \geq |C_t|$ allocated goods from C_t , contradicting the existence of g . Finally, note that because of this, $A'_j \cup \{g\}$ must be a feasible bundle, and we must have that $v_j(A'_j) \leq v_j(A'_j \cup \{g\})$. Consider what happens if we replace A'_j with $A'_j \cup \{g\}$, this allocation would not decrease the number of agents with positive utility, or the product of positive utilities, meaning it must also be an MNW allocation. However, it will strictly decrease the number of unallocated goods, causing a contradiction, and proving that A must be MNW over all allocations. \square

Now, we are ready to prove that MNW allocations in fact get $\frac{1}{2}$ -EF1_{WC} for goods with copies.

Theorem 4. *In the case of goods with copies, every MNW and complete MNW allocation is $\frac{1}{2}$ -EF1_{WC} and PO.*

Proof. For some instance constrained by a partition matroid $\mathcal{M} = (M, \mathcal{I})$, and any agent $i \in N$, let $B_i : 2^M \rightarrow \mathcal{I}$ be the function that maps any set $S \subseteq M$ to i 's most preferred feasible allocation I such that $I \subseteq S$.

First notice that in any MNW allocation A in a matroid constrained instance, and for any pair of agents $i, j \in N$, either $A_j = \emptyset$, or there exists some good $g \in A_j$ such that $2v_i(A_i) \geq v_i(B_i(A_i \cup A_j \setminus \{g\}))$. One can see that this condition is derived from the definition of ‘‘Marginal EF1’’ that Caragiannis et al. [2016] prove holds for unconstrained instances with submodular valuation functions, for the purposes of this proof we will refer to it as ‘‘Constrained Marginal EF1’’. In the proof of Theorem 1, we show that every MNW allocation over a matroid constraint corresponds to an allocation over an unconstrained setting where each agent $i \in N$ has the valuation function \tilde{v}_i , equaling the the ‘‘weighted rank function’’ of the matroid with their original valuation function v_i being used as the weights. It is known that this unconstrained allocation will be Marginal EF1, and for every $i, j \in N$, we will have $\tilde{v}_i(A_j) = v_i(A_j)$. This directly implies that the MNW allocation over the original constrained instance will be Constrained Marginal EF1.

It is clear that if an allocation is constrained marginal EF1, then it will also be $\frac{1}{2}$ -EF1. What we will show next is that if the property holds for an allocation A over a goods with copies instance, then A will be $\frac{1}{2}$ -EF1_{WC}.

For contradiction, assume that this is false, and for some goods with copies instance constrained by \mathcal{C} , there exists an allocation A that is Constrained Marginal EF1 but not $\frac{1}{2}$ -EF1_{WC}.

Since we know that A is Constrained Marginal EF1, we know there exists some constraint $C \in \mathcal{C}$, and some $g \in A_j \cap C$, such that $2v_i(A_i) \geq v_i(B_i(A_i \cup A_j \setminus \{g\}))$. We will analyze two cases:

Case 1: $|A_i \cap C| = 0$. In this case, $\frac{1}{2}$ -EF1_{WC} directly follows from the fact that $2v_i(A_i) \geq v_i(B_i(A_i \cup A_j \setminus \{g\}))$ implies $2v_i(A_i) \geq v_i(A_j \setminus \{g\})$.

Case 2: $|A_i \cap C| = 1$. In this case, by the way that goods with copies constraints are structured, it must be true that $B_i(A_i \cup A_j \setminus \{g\}) = B_i(A_i \cup A_j)$. We know that agent i has the same valuation

for all goods in C . We also know that there can only be 1 good from C in any feasible bundle, and that $A_i \cup A_j$ contains at least two goods from C (g as well as some good $g' \in A_i$). Since $A_i \cup A_j$ has multiple goods from C that both have equivalent value, it is easy to see that removing one of them will have no affect on the best possible bundle i can construct using the goods. Therefore, in this case, Constrained Marginal EF1 implies that $2v_i(A_i) \geq v_i(B_i(A_i \cup A_j)) \geq v_i(A_j)$, which gives us $\frac{1}{2}$ -EF1_{WC}.

These two cases exhaust all possible scenarios (either A_i has a copy from C or it does not), giving us a contradiction.

This proves that for goods with copies instances, all MNW allocations will be $\frac{1}{2}$ -EF1_{WC}. From [Lemma 3](#), we can also conclude that the set of all complete MNW allocations for a given instance will be a subset of the MNW allocations, and they will therefore be $\frac{1}{2}$ -EF1_{WC} as well. \square

Because the primary motivation of goods with copies is that it can be used to model chore division as a special case, this is one example where free disposal of goods is not possible (otherwise, some chores may be assigned to multiple agents). Hence, it is useful that [Theorem 4](#) handles not only MNW allocations, but also complete MNW allocations.

To the best of our knowledge, this is the first known approximation for EF1_{WC} under general additive preferences. [Gafni et al. \[2023\]](#) leave the existence of EF1_{WC} open, only proving it for special cases such as that of leveled preferences. The fact that MNW allocations, in addition to being $\frac{1}{2}$ -EF1_{WC}, also satisfy PO is an added bonus.

However, there is one significant downside to [Theorem 4](#). When a chore division instance is modeled as a goods with copies instance, while EF1_{WC} for the latter implies EF1 for the former, even $\frac{1}{2}$ -EF1_{WC} for the latter does not imply any reasonable approximation of EF1 for the former, as the following example shows.

Example 2. Consider the following instance with 3 agents and 3 goods labeled g_1 through g_3 , each with 2 copies.

| | g_1 | g_2 | g_3 |
|---------|-------|-------|-------|
| Agent 1 | 1 | 1/100 | 1/100 |
| Agent 2 | 1 | 1 | 1 |
| Agent 3 | 1 | 1 | 1 |

The MNW allocation A (unique, up to symmetry between agents 2 and 3) is given by $A_1 = \{g_1\}$, $A_2 = \{g_1, g_2, g_3\}$, and $A_3 = \{g_2, g_3\}$. One can verify that this allocation is $\frac{1}{2}$ -EF1_{WC}. However, for the underlying chore division instance, where the values of each agent for each chore is the negation of the value for the corresponding good given above, the corresponding allocation A' of chores is given by $A'_1 = \{g_2, g_3\}$, $A'_2 = \emptyset$, and $A'_3 = \{g_1\}$. This allocation provides zero approximation of EF1 because we have $v_1(A'_1 \setminus \{g_2\}) = v_1(A'_1 \setminus \{g_3\}) = -1/100$ but $v_1(A'_2) = 0$.

Recently, [Garg et al. \[2024\]](#) show that $\frac{1}{2}$ -EF2 and PO can be achieved for chore division through a Fisher market based algorithm. It would be interesting to extend this to goods with copies.

Search for EF1+PO allocations. The primary goal of our work is to initiate a systematic study of the existence of *fair and efficient* allocations under feasibility constraints. The ultimate goal of seeking (exactly) EF1 and PO allocations still remains a major open question in all cases.

Open Question: Does an allocation that is EF1 and PO always exist subject to an arbitrary matroid constraint and heterogeneous additive valuations? What about just the special case of goods with copies?

It is worth remarking that subject to matroid constraints (resp., budget constraints), MNW allocations achieve $\frac{1}{2}$ -EF1 (resp., $\frac{1}{4}$ -EF1) with PO (Theorem 1 and Wu et al. [2021]), while the open question is whether exact EF1 (resp., $\frac{1}{2}$ -EF1) with PO is achievable. Curiously, in both cases, MNW allocations achieve half the fairness level of the known upper bound subject to PO.

In fact, even the existence of an EF1 allocation remains open beyond partition (resp., base-orderable) matroids for heterogeneous (resp., identical) additive valuations [Biswas and Barman, 2018, 2019], and the existence of an EF1_{WC} allocation remains open for goods with copies.

3.3 Impossibility of Stronger Guarantees

In this section, we explore whether the fairness and efficiency guarantees established in Theorem 1 (and in the subsequent results above) can be strengthened. We explore several directions, establishing a negative result in almost all cases but also identifying interesting open questions.

3.3.1 Tightness of the Fairness Guarantee of MNW

First, we show that MNW does not achieve a better guarantee than $\frac{1}{2}$ -EF1 under matroid constraints.

Theorem 5. *For any $\epsilon > 0$, there exist instances with partition matroid constraints and the balancedness constraint in which no MNW (or complete MNW) allocation is $(\frac{1}{2} + \epsilon)$ -EF1.*

Proof. Let $k \in \mathbb{N}$. Consider the partition-matroid-constrained instance with $n = 2$ agents and $m = 2k$ goods, with the goods being partitioned into two groups, $S = \{g_1, \dots, g_k\}$ and $T = \{g_{k+1}, \dots, g_{2k}\}$. For each good $g_s \in S$, we have $v_1(g_s) = v_2(g_s) = 1$. For each good $g_t \in T$, we have $v_1(g_t) = 0$ and $v_2(g_t) = \frac{1}{2}$. The partition matroid constraints dictate that no agent can receive more than k goods from the entire set of goods $M = S \cup T$.

We argue that the only MNW allocation is the complete allocation A with $A_1 = S$ and $A_2 = T$. To see this, consider any allocation A . Let $x = v_1(A_1) = |A_1 \cap S|$. Then, $v_2(A_2) \leq (k - x) + \frac{1}{2}x$, as agent 2 can receive at most k goods in total, and her utility is maximized by receiving the remaining $k - x$ goods of S along with some x goods from T . Hence, $NW(A) \leq x \cdot (k - \frac{1}{2}x)$. The right hand side is uniquely maximized at $x = k$, resulting in the allocation A with $A_1 = S$ and $A_2 = T$ uniquely attaining the highest Nash welfare of $\frac{k^2}{2}$.

In this allocation, agent 1 is not envious, but agent 2 has $v_2(A_2) = \frac{1}{2}k$ whereas $v_2(A_1 \setminus \{g\}) = k - 1$ for any $g \in A_1$. Hence, this allocation is $\frac{k}{2(k-1)}$ -EF1, i.e., $(\frac{1}{2} + \frac{1}{2k-2})$ -EF1. For any $\epsilon > 0$, one can choose $k > 1 + \frac{1}{2\epsilon}$ to ensure that the unique MNW allocation is not $(\frac{1}{2} + \epsilon)$ -EF1.

Since the unique MNW allocation in this instance is in fact balanced, changing the partition matroid constraint (each agent receives at most k goods) to the balancedness constraint (each agent receives exactly k goods) results in the same allocation being the unique MNW allocation, thus establishing tightness of the $\frac{1}{2}$ -EF1 guarantee with respect to the balancedness constraint too. \square

The tightness result above clearly extends to all families of constraints more general than partition matroid constraints (in particular, to laminar and base-orderable matroid constraints) or balancedness (in particular, to partition matroids with lower bounds). Thus, the only family of constraints we study that is not covered is goods with copies, for which we provide a similar example in [Appendix A.2](#).

3.3.2 Strengthening Fairness Criteria

One common strengthening of EF1 is “Stochastically Dominant EF1” (SD-EF1). An allocation satisfies SD-EF1 if it can be inferred that the allocation is EF1 by only looking at the agents’ ordinal preferences over the goods. We provide the formal definition of SD-EF1 below:

Define \succeq_i (resp., \succ_i) as the weak (resp., strict) ordering over the goods in M induced by v_i , where, for all $g, g' \in M$, $g \succeq_i g'$ if and only if $v_i(g) \geq v_i(g')$ and $g \succ_i g'$ if and only if $v_i(g) > v_i(g')$.

Definition 5 (SD-EF1). An allocation A of a set of goods S is stochastic-dominance envy-free up to one good (SD-EF1) if for all $i, j \in N$ with $A_j \neq \emptyset$, there exists a $g^* \in A_j$ such that for all $g \in S$, $|\{g' \in A_i : g' \succeq g\}| \geq |\{g' \in A_j \setminus \{g^*\} : g' \succeq g\}|$.

In unconstrained instances, SD-EF1 is known to be achievable through the “round-robin” mechanism. Since round-robin always produces a balanced allocation, this means that SD-EF1 is achievable under balancedness constraints as well. However, under more general partition matroid constraint, we can show that SD-EF1 is no longer able to be guaranteed.

Theorem 6. *Under partition matroid constraints, SD-EF1 cannot be guaranteed.*

Proof. Consider the following instance, in which there are two agents, and eight goods partitioned into 4 subsets of 2 goods, represented by the red groupings. We place an upper-bound of 1 on each grouping, meaning that each agent must receive exactly 1 good from each group.

| | g_1 | g_2 | g_3 | g_4 | g_5 | g_6 | g_7 | g_8 |
|---------|-------|-------|-------|-------|-------|-------|-------|-------|
| Agent 1 | 8 | 4 | 7 | 5 | 6 | 1 | 3 | 2 |
| Agent 2 | 6 | 2 | 8 | 5 | 7 | 3 | 4 | 1 |

In any SD-EF1 allocation A over this instance, it follows from the definition that each agent must get at least 1 of their top 2 goods. This means that A_1 must contain one of $\{g_1, g_3\}$ and A_2 must contain at least one of $\{g_3, g_5\}$. Similarly, it must be the case that each agent receives at least 2 of their top 4 goods. This means that both A_1 and A_2 must contain exactly 2 of $\{g_1, g_3, g_4, g_5\}$.

Notice that under these restrictions, in any SD-EF1 allocation where agent 1 receives g_3 , they cannot receive g_5 (or they would have both of agent 2’s top 2 goods) or g_4 (due to the constraints). Since the allocation over the top 4 goods must be balanced (each agent receiving exactly two of them), the only possible allocation over $\{g_1, g_3, g_4, g_5\}$ in this scenario would be $(\{g_1, g_3\}, \{g_4, g_5\})$. Using the same logic, when agent 2 receives g_3 , the only possible allocation is $(\{g_1, g_4\}, \{g_3, g_5\})$. Since one of the agents must be given g_3 , these are the only two ways that the top 4 goods can be allocated. Note that in both of these allocations, agent 1 gets g_1 and agent 2 gets g_5 .

Finally, consider how the remaining goods $\{g_2, g_6, g_7, g_8\}$ must be allocated to guarantee SD-EF1. Each agent must receive at least 3 of their top 6 goods, and since we know that each agent has exactly 2 of their top 4 goods, that means that agent 1 must receive at least one of $\{g_2, g_7\}$, and agent 2 must receive at least one of $\{g_6, g_7\}$. Since we know that agent 1 must be allocated g_1 , the matroid constraints say they cannot receive g_2 , so they must receive g_7 . Similarly, agent 2 is known to have g_5 , so they cannot receive g_6 , which means they must also receive g_7 . Both agents cannot be simultaneously allocated g_7 , hence, in this instance, there is no complete and feasible allocation that is SD-EF1. \square

Notably, this settles the question of SD-EF1 existence for all of the main classes of constraints studied in this paper, except for goods with copies. We leave whether SD-EF1 can be achieved for goods with copies as an open question.

Open Question: Does there always exist an SD-EF1 allocation of goods with copies?

3.3.3 Strengthening Efficiency Criteria

In this section, we first consider a strengthening of PO, which we term PO^+ . Intuitively, an allocation $A \in \mathcal{F}$ is PO^+ if it is not Pareto dominated by any unconstrained allocation, i.e. any possible allocation that could be constructed from a base set of goods, regardless if it is in \mathcal{F} . We define the set of complete unconstrained allocations analogously.

Definition 6 (Unconstrained Pareto Optimality (PO^+)). Let \mathcal{U} be the set of unconstrained allocations. An allocation $A \in \mathcal{F}$ is PO^+ if there is no other allocation $B \in \mathcal{U}$ such that $v_i(B_i) \geq v_i(A_i)$ for each $i \in N$ and at least one inequality is strict. In words, no other allocation, feasible or infeasible, should be able to make an agent happier without making any agent less happy.

For goods with copies, and hence, for partition and laminar matroid constraints, PO^+ is easily seen to be too strong, even in the absence of any other desiderata.

Example 3. Consider an goods with copies instance with 3 agents, and 3 goods with 2 copies each. Agents have the following valuations for the goods, where $\epsilon < \frac{1}{2}$ is some very small value:

| | g_1 | g_2 | g_3 |
|---------|------------|------------|------------|
| Agent 1 | 1 | ϵ | ϵ |
| Agent 2 | ϵ | 1 | ϵ |
| Agent 3 | ϵ | ϵ | 1 |

If we ignored the restrictions on feasibility enforced by the goods with copies constraints, then clearly the most efficient way to allocate the goods would be to assign both copies of g_i to a_i for each $i \in \{1, 2, 3\}$. This would result in each agent having a valuation of 2. However, under the goods with copies constraints, each agent can only get a single copy of the one good they find valuable, so in any feasible allocation, there would be no agent with a utility higher than $1 + 2\epsilon < 2$.

However, the balanced setting is a curious special case. A result by [Conitzer et al. \[2017\]](#) implies the following.

Theorem 7 (Conitzer et al. 2017). *In the balanced case, when each agent has a strictly positive valuation for each good, there always exists a feasible allocation that is PO^+ .*

This raises the interesting question, in the balanced case where all agents have strictly positive utilities, does there always exist an allocation that is $EF1+PO^+$? It turns out this is not true. As we show in the below theorem, it is not possible to guarantee any approximation of $EF1$ along with PO^+ in the balanced case.

Theorem 8. *Under the balancedness constraint, it is impossible to guarantee α - $EF1$ and PO^+ for any $\alpha \in (0, 1]$.*

Proof. For contradiction, assume this is false, and for some $\alpha \in (0, 1]$, every balanced instance admits an allocation that is PO^+ and α - $EF1$.

Let $k = \lceil 2/\alpha \rceil$. Consider the following balanced-constrained instance, with 2 agents and $m = 2(k^2 + k)$ goods. The m goods are broken up into two types, a set L of $m/2$ low-valued goods, for which $v_1(g) = v_2(g) = 1$ for all $g \in L$, and a set H of high-valued goods, for which $v_1(g) = k + 1$ and $v_2(g) = k$ for each $g \in H$. Because this the allocation must be balanced, in every feasible instance both agents will receive exactly $k^2 + k$ goods.

By our assumption, we know that there exists an allocation for this instance that is α - $EF1$ and PO^+ , call this allocation A .

We first note that any instance in which agent 1 receives at least k goods from L will not be PO^+ . To see this, consider what happens if agent 1 has k goods from L and agent 2 has 1 good from H . It would be a Pareto improvement for the agents to swap these goods, as agent 1's utility would increase by 1, and agent 2's utility would remain the same.

Also note that if agent 1 has at most $k - 1$ goods from L , then due to the balancedness constraint that they must have $m/2 = k^2 + k$ total goods, agent 1 must have at least $k^2 + 1$ goods from H . Since $A_2 = M \setminus A_1$, the inverse of this argument would be that agent 2 has at most $k - 1$ goods from H and at least $k^2 + 1$ goods from L . Therefore, for agent 2, we have

$$\begin{aligned} v_2(A_1) &\geq (k^2 + 1) \cdot k + (k - 1) \cdot 1 > k^3 + k, \\ v_2(A_2) &\leq (k - 1) \cdot k + (k^2 + 1) \cdot 1 \leq 2k^2, \end{aligned}$$

Since A is α - $EF1$, we know the following statement must be true

$$v_2(A_2) \geq \alpha \cdot (v_2(A_1) - k).$$

Substituting the bounds from above, we have

$$2k^2 > \alpha k^3 \iff 2/\alpha > k,$$

which contradicts our assumption that $k = \lceil 2/\alpha \rceil$. Therefore, A cannot be α - $EF1$ and PO^+ . \square

4 Best-of-Both-Worlds Guarantees

Next, we present randomized allocations that achieve desirable fairness and efficiency guarantees both ex ante (i.e., in expectation) and ex post (i.e., in every allocation in the support of the probability distribution). Aziz et al. [2024] initiated this line of work in (unconstrained) fair division, now referred to as “best of both worlds” (BoBW) guarantees.

We denote the dot product of two vectors $x, y \in \mathbb{R}^d$ by $x \cdot y = \sum_{j \in [d]} x_j \cdot y_j$. A *fractional allocation* is denoted by $x \in [0, 1]^{N \times M}$, where $x_{i,g}$ is the “fraction” of good g allocated to agent i and $\sum_{i \in N} x_{i,g} = 1$ for each $g \in M$. The utility of agent i under a fractional allocation A is $v_i(A_i) = \sum_{g \in M} x_{i,g} \cdot v_i(g)$.

A randomized allocation $x = \sum_{\ell \in [L]} \lambda_\ell A^\ell$ is a probability distribution in which (integral) allocation A^ℓ is selected with probability λ_ℓ for each $\ell \in [L]$ and $\sum_{\ell \in [L]} \lambda_\ell = 1$. It induces a fractional allocation, also denoted x , in which $x_{i,g}$ is the marginal probability of agent i being allocated good g .

We say that a randomized allocation x satisfies a desideratum D *ex ante* if its induced fractional allocation satisfies D ,⁵ and that x satisfies a desideratum D *ex post* if every (integral) allocation in its support satisfies D .

4.1 EF+PO fractional allocations under linear constraints.

First, we define (a relevant special case of) the model of Echenique et al. [2021].

Definition 7 (Linearly-Constrained Divisible Economy). There is a set of agents $N = [n]$ and a set of divisible goods $M = [m]$, with each good g having a *supply* $q_g \in [1, n]$. For some $d \in \mathbb{N}$, there is a $d \times m$ matrix A and a $d \times 1$ vector b , both with non-negative entries, such that, for a fractional allocation x , $x \in \mathcal{F}$ if and only if $Ax_i \leq b$ for each $i \in N$.

The supply is useful for modeling the case of goods with copies. Note that linear constraints can be used to capture goods with copies and/or balancedness (actually, even fractional versions of laminar matroids, but our BoBW results hold for only these two cases). For goods with copies, we set q_g to be the number of copies of g available, and demand $x_{i,g} \leq 1, \forall g \in M$ for each $i \in N$, which is a set of linear constraints. To impose balancedness, we add $\sum_{g \in M} x_{i,g} \leq m/n$ for each $i \in N$; note that m/n does not need to be an integer and no lower bounds are required in this fractional case.

Definition 8 (Competitive Equilibrium From Equal Incomes (CEEI)). For a fractional allocation x and price vector $p \in \mathbb{R}_{\geq 0}^M$, (x, p) is a competitive equilibrium from equal incomes (CEEI) if the following conditions hold:

1. x is feasible, i.e., $Ax_i \leq b$ for all $i \in N$.
2. Each agent i is allocated a bundle x_i that maximizes her utility subject to a budget of 1, i.e.,

$$x_i \in \arg \max_x \{v_i(x) \mid p \cdot x \leq 1\}.$$

Further, x_i should be the *cheapest bundle* (minimizing $p \cdot x_i$) subject to achieving the maximum utility of $v_i(x_i)$.

3. The market clears, i.e., $\sum_{i \in N} x_{i,g} \leq q_g$ for all g , and $\sum_{i \in N} x_{i,g} < q_g$ only if $p_g = 0$.

⁵We will only refer to envy-freeness (EF) and Pareto optimality (PO) of fractional allocations, defined exactly as they are for integral allocations. PO for a fractional allocation requires that not even a *fractional* allocation Pareto dominate it.

Algorithm 1 CE Lottery for goods-with-copies

- 1: $x \leftarrow$ a competitive (fractional) allocation of [Theorem 9](#) due to [Echenique et al. \[2021\]](#)
- 2: For all $i \in N$ and $t \in [m]$, let $Q_{i,t} \leftarrow \sum_{j \in [t]} x_{i,\sigma_i(j)}$ be the total fraction of i 's top- t goods allocated to her ($\sigma_i(j)$ denotes agent i 's j -th most preferred good)
- 3: Define the bihierarchy set of constraints for an integral allocation y :

$$\mathcal{H}_1 = \{ \lfloor Q_{i,t} \rfloor \leq \sum_{j \in [t]} y_{i,\sigma_i(j)} \leq \lceil Q_{i,t} \rceil \mid i \in N, t \in [m] \},$$
$$\mathcal{H}_2 = \{ \sum_i y_{i,g} \leq q_g \mid g \in M \} \cup \{ y_{i,g} \leq 1 \mid i \in N, g \in M \}.$$

- 4: Decompose $x = \sum_{\ell=1}^L \lambda_\ell \cdot y^\ell$ into L integral allocations using the algorithm of [Budish et al. \[2013\]](#) (see their Appendix B) with input x and $\mathcal{H}_1, \mathcal{H}_2$
 - 5: **return** the randomized allocation $\sum_{\ell=1}^L \lambda_\ell \cdot y^\ell$
-

The following is a corollary of [Theorem 1](#) of [Echenique et al. \[2021\]](#).⁶

Theorem 9 ([Echenique et al. 2021](#)). *For a linearly-constrained divisible economy with strictly positive additive valuations,⁷ there always exists a fractional allocation x and a price vector p such that (x, p) is CEEI. Further, in any such CEEI, x is envy-free and Pareto optimal.*

Our goal is to implement a fractional allocation x that is part of a CEEI as a probability distribution over integral allocations that satisfy two relaxations of EF1 (satisfying them ex post): Prop1 [[Conitzer et al., 2017](#)] and EF₁¹ [[Barman and Krishnamurthy, 2019](#)].

Definition 9 (Proportionality up to one good (Prop1)). An integral allocation A is *proportional up to one good* (Prop1) if, for each agent $i \in N$, either $v_i(A_i) \geq v_i(M)/n$ or there exists a good $g \notin A_i$ such that $v_i(A_i \cup \{g\}) \geq v_i(M)/n$.

Definition 10 (Envy-freeness up to one good more-and-less (EF₁¹)). An integral allocation A is *envy-free up to one good more-and-less* (EF₁¹) if, for every pair of agents $i, j \in N$ such that $A_j \neq \emptyset$ and $A_i \neq M$, $v_i(A_i \cup \{g_i\}) \geq v_i(A_j \setminus \{g_j\})$ for some $g_i \notin A_i$ and $g_j \in A_j$.

We would also like our implementation of x to be PO ex post. Luckily, this requirement is straightforward, as it is known that the implementation of any ex ante PO fractional allocation will be ex post PO.

Proposition 1 ([\[Aziz et al., 2024\]](#) Proposition 1). *If a randomized allocation x is ex ante PO, then x is also ex post PO.*

4.2 Algorithm For Decomposing Constrained CEEIs

We decompose a fractional CEEI allocation using [Algorithm 1](#) to obtain the following.

Theorem 10. *For goods with copies (resp., goods with copies with an additional balancedness constraint) and strictly positive valuations, there always exists a randomized allocation that is ex ante EF and PO, and ex post EF₁¹, Prop1, and PO (resp., Prop1 and PO).*

⁶Their stronger result holds for semi-strictly quasi-concave utilities, which subsumes strictly positive additive utilities.

⁷That is, $v_i(g) > 0$ for all $i \in N$ and $g \in M$.

Before proving [Theorem 10](#), we will introduce several helpful tools, the first of which being the bihierarchy matrix decomposition framework of [[Budish et al., 2013](#)], which [Algorithm 1](#) (closely following the algorithm of [[Aziz et al., 2024](#)]) uses as a black-box.

The bihierarchy theorem states how a fractional allocation can be decomposed into a convex combinations of integral allocations in such a way that each integral allocation adheres to a set of constraints. In the language of [[Budish et al., 2013](#)], a constraint over an allocation X takes the form of $q_S \leq \sum_{(i,r) \in S} X_{i,r} \leq \bar{q}_S$, where S is a set containing agent-object pairs, and $q_S, \bar{q}_S \in \mathbb{N}$ are upper and lower quotas for the sum of the elements of S . An allocation X satisfies a constraint S if the statement $q_S \leq \sum_{(i,r) \in S} X_{i,r} \leq \bar{q}_S$ is true. A constraint structure \mathcal{H} is a collection of constraints, and some allocation X satisfies a constraint structure \mathcal{H} if it satisfies all constraints in \mathcal{H} . The bihierarchy theorem is specifically for collections of constraints that form a bihierarchy structure.

Definition 11 (Hierarchy and Bihierarchy [[Budish et al., 2013](#)]). A constraint structure \mathcal{H} is a hierarchy iff for any pair of sets $S, S' \in \mathcal{H}$, we have that $S \cap S' \in \{\emptyset, S, S'\}$. A constraint structure \mathcal{H} is a bihierarchy iff there exists two hierarchies $\mathcal{H}_1, \mathcal{H}_2$ such that $\mathcal{H} = \mathcal{H}_1 \cup \mathcal{H}_2$ and $\mathcal{H}_1 \cap \mathcal{H}_2 = \emptyset$.

Theorem 11 (Bihierarchy Theorem [[Budish et al., 2013](#)]). For any fractional allocation x and bihierarchy \mathcal{H} , if x satisfies \mathcal{H} , then it can be represented as a randomized allocation $x = \sum_{\ell \in [L]} \lambda_\ell A^\ell$, where for each $\ell \in [L]$, A^ℓ satisfies \mathcal{H} . This randomized allocation can be found in strongly polynomial time.

We will also introduce the Utility++ guarantee of [[Aziz et al., 2024](#)], which is what we use to ensure that the integral allocations produced by [Algorithm 1](#) have the desired fairness properties.

Through the bihierarchy presented in [Algorithm 1](#), [Aziz et al. \[2024\]](#) prove a utility guarantee that we restate below. Recall the definition of the two hierarchies $\mathcal{H}_1 \cup \mathcal{H}_2$,

$$\begin{aligned} \mathcal{H}_1 &= \left\{ \lfloor Q_{i,t} \rfloor \leq \sum_{j \in [t]} y_{i, \sigma_i(j)} \leq \lceil Q_{i,t} \rceil \mid i \in N, t \in [m] \right\}, \\ \mathcal{H}_2 &= \left\{ \sum_i y_{i,g} \leq q_g \mid g \in M \right\} \cup \left\{ y_{i,g} \leq 1 \mid i \in N, g \in M \right\}, \end{aligned}$$

where $\sigma_i(j)$ is the j th good in the preference ranking of i (ties broken arbitrarily) and $Q_{i,t} = \sum_{j \in [t]} x_{i, \sigma_i(j)}$ is the total fraction of goods allocated to i in x among her top t goods. The first hierarchy for i enforces that the number of goods i receives among her first top t goods in the integral allocations, i.e., $\sum_{j \in [t]} y_{i, \sigma_i(j)}$, must be as close as possible to $Q_{i,t}$ (the total fraction of i 's top t goods allocated to her in x). The second hierarchy \mathcal{H}_2 ensures that the integral allocations adhere to goods with copies constraints, where each agent receives no more than a single copy of each good, and no more than q_g total copies of any good $g \in M$ are allocated. By invoking [Theorem 11](#) with the above bihierarchy set, we have the following utility guarantee.

Lemma 4 (Utility Guarantee++ [[Aziz et al., 2024](#)]). For a fractional allocation x and the bihierarchy set of constraints in [Algorithm 1](#), there exists a randomized allocation implementing $x = \sum_{\ell \in [L]} \lambda_\ell A^\ell$ such that for all integral allocations A^ℓ the following holds:

1. If $v_i(A_i^\ell) < v_i(x_i)$, then there exists a good g such that $x_{i,g} > 0$ and $v_i(A_i^\ell \cup \{g\}) > v_i(x_i)$.
2. If $v_i(A_i^\ell) > v_i(x_i)$, then there exists a good g such that $x_{i,g} < 1$ and $v_i(A_i^\ell \setminus \{g\}) < v_i(x_i)$.

Finally, we introduce the concept of ‘‘Bang-per-Buck’’ in a CEEI. For an allocation x and price vector p such that (x, p) is a CEEI, the bang-per-buck ratio of good g for agent i is $\text{bb}_i(g) = v_i(g) / p_g$.

Bang-per-buck ratios have been found to be very useful when working with CEEIs of unconstrained instances with additive valuations, as each agent’s CEEI bundle will contain only goods with the maximum bang-per-buck for that agent (see [Barman et al., 2018] for example). By the following lemma, we can show that even under goods with copies constraints, examining the bang-per-buck of each good will still allow us to give a similar structure to a CEEI solution.

Lemma 5. *In the case of goods with copies, given an allocation x and price vector p such that (x, p) is CEEI, for any agent i , a good g_1 that is fully allocated to i (i.e., $x_{i,g_1} = 1$), goods g_2 and g_3 that are partially allocated to i (i.e., $x_{i,g_2}, x_{i,g_3} \in (0, 1)$), and g_4 that is not allocated to i (i.e., $x_{i,g_4} = 0$), the following relation of bang-per-buck ratios holds:*

$$\text{bb}_i(g_1) \geq \text{bb}_i(g_2) = \text{bb}_i(g_3) \geq \text{bb}_i(g_4).$$

Proof. Assume by contradiction that for a pair of goods $g, g' \in \{g_1, \dots, g_4\}$ the claim does not hold, i.e., $\text{bb}_i(g) < \text{bb}_i(g')$ and $x_{i,g} > 0$. It holds that g is either partially or fully allocated to i , while g' is either partially or not at all allocated to i . Let $\delta_{g'} \in (0, 1)$ be some positive value such that $\delta_{g'} \leq 1 - x_{i,g'}$ and $\delta_{g'} \cdot p_{g'} / p_g \leq x_{i,g}$. Take a bundle x'_i such that $x'_{i,g'} = x_{i,g'} + \delta_{g'}$, $x'_{i,g} = x_{i,g} - \delta_{g'} \cdot p_{g'} / p_g$, and $x'_{i,g''} = x_{i,g''}$ for all $g'' \notin \{g, g'\}$. That is, agent i “spends” a nonzero part of their budget on g' instead of g . The inequalities ensure that such a transfer is feasible. Then, since $\text{bb}_i(g') > \text{bb}_i(g)$ and the rest of the spending of agent i in x'_i is the same as in x_i , we have $v_i(x'_i) > v_i(x_i)$, which contradicts condition 2 of Definition 8. \square

Note that the proof above does not necessarily hold when a balancedness constraint is also enforced on top of the goods with copies constraint. For instance, the transfer made in the proof may lead to an unbalanced bundle x'_i which is inadmissible.

With these definitions in mind, we are ready to prove Theorem 10.

Proof of Theorem 10. Let A^1, \dots, A^ℓ be the set of allocations returned by Algorithm 1 based on the fractional allocation x and price vector p such that (x, p) is CEEI. The ex-ante PO guarantee follows from Theorem 9, whose implementation is also ex-post PO by Proposition 1.

First, note that the bihierarchy constraints \mathcal{H}_1 and \mathcal{H}_2 will ensure that the goods with copies (resp. balancedness) constraints will hold for each A in the support of x . Any fractional goods with copies allocations will clearly adhere to the constraints in \mathcal{H}_2 , and \mathcal{H}_2 will enforce that each integral allocation in the support will meet the goods with copies constraint. The constraints $\{\sum_i y_{i,g} \leq q_g \mid g \in M\}$ will ensure that no more than q_g copies are allocated to each agent, and the constraints $\{y_{i,g} \leq 1 \mid i \in N, g \in M\}$ will ensure that no single agent receives more than 1 copy of any good.

The fact that the fractional allocation x will adhere to the constraints in \mathcal{H}_1 is self-evident. Its primary purpose is to ensure our fairness guarantees hold ex post, however, note that \mathcal{H}_1 also guarantees that if x is a balanced fractional allocation, then each of the integral allocations in its support will also be balanced. This is due to the fact that the constraints enforced by \mathcal{H}_1 for $t = m$ will be $\{\lfloor Q_{i,m} \rfloor \leq \sum_{j \in [m]} y_{i,\sigma_i(j)} \leq \lceil Q_{i,m} \rceil \mid i \in N\}$. Here, when x is balanced, we have that $Q_{i,m} = \sum_{j \in M} x_{i,j} = m/n$ and $\sum_{j \in [m]} y_{i,\sigma_i(j)} = \sum_{j \in [m]} y_{i,j}$.

The Prop1 guarantee follows from Lemma 4, as in the proof of Aziz et al. [2024]. Either $v_i(A_i) \geq v_i(x_i) \geq v_i(M)/n$, where the last inequality holds because EF implies Prop, or by Lemma 4, there exists a good $g_i \notin A_i$ such that $v_i(A_i \cup \{g_i\}) \geq v_i(x_i) \geq v_i(M)/n$.

Next, we prove the EF_1^1 guarantee. Consider agents i and j such that $v_i(A_i) < v_i(A_j)$. By [Lemma 4](#), there exists a good g_i with $x_{i,g_i} < 1$ such that

$$v_i(A_i \cup \{g_i\}) > v_i(x_i), \quad (3)$$

and a good g_j with $x_{j,g_j} < 1$ (i.e., partially allocated to j) such that

$$v_j(A_j \setminus \{g_j\}) < v_j(x_j). \quad (4)$$

We show that $p(A_j \setminus \{g_j\}) < p(x_j)$. Let x_j^{full} be the bundle containing all the goods that are fully allocated to j . We also have $x_j^{\text{full}} \subseteq A_j \setminus \{g_j\}$, since the method of [Budish et al. \[2013\]](#) ensures that any good fully allocated to one agent remains allocated to the same agent in all integral allocations. Rewriting, we have $v_j(A_j \setminus \{g_j\}) = v_j(x_j^{\text{full}}) + v_j(A_j \setminus (x_j^{\text{full}} \cup \{g_j\}))$ and $v_j(x_j) = v_j(x_j^{\text{full}}) + v_j(x_j \setminus x_j^{\text{full}})$. Combined with Eq. (4), we obtain

$$v_j(A_j \setminus (x_j^{\text{full}} \cup \{g_j\})) < v_j(x_j \setminus x_j^{\text{full}}).$$

Both sides of the inequality consist of goods that are partially allocated to j in x . By [Lemma 5](#), all goods involved on both sides must have the same bang-per-buck ratio. Hence, by multiplying both sides by that ratio, we get

$$p(A_j \setminus (x_j^{\text{full}} \cup \{g_j\})) < p(x_j \setminus x_j^{\text{full}}),$$

and after adding $p(x_j^{\text{full}})$ to both sides, we have

$$p(A_j \setminus \{g_j\}) < p(x_j) \leq 1.$$

Since $A_j \setminus \{g_j\}$ fits within the budget of 1, by the best bundle condition of CE, $v_j(A_j \setminus \{g_j\}) \leq v_j(A_j)$. Along with Eq. (3), we conclude that $v_i(A_i \cup \{g_i\}) > v_i(A_j \setminus \{g_j\})$, as desired. \square

4.3 Equivalence of EF_1^1 for chores and goods-with-copies.

It is clear to see that the guarantees we receive for goods with copies constraints directly imply the same guarantees in the unconstrained world, one simply has to set the number of copies for all goods to 1. We also argue that this result implies the same guarantees for unconstrained instances with chores rather than goods.

For full context, the definition for EF_1^1 for chores is as follows

Definition 12 (EF_1^1 for Chores). An integral allocation A is *envy-free up to one good more-and-less* (EF_1^1) if, for every pair of agents $i, j \in N$ such that $A_i \neq \emptyset$ and $A_j \neq M$, $v_i(A_i \setminus \{c_i\}) \geq v_i(A_j \cup \{c_j\})$ for some $c_i \in A_i$ and $c_j \notin A_j$.

When [[Gafni et al., 2023](#)] introduced the goods with copies framework, they showed that when an allocation is EF_{1WC} in an instance where each good has $n - 1$ copies, the allocation over the dual chores instance induced by the goods with copies will be EF_1 . EF_{1WC} is a stronger notion than EF_1 , and one may note that we are only claiming to find EF_1^1 over goods with copies instances, rather than some analogous “WC” strengthening of it. In the below lemma, we will show that in fact, EF_1^1 is a unique case of fairness property, where finding it in the goods with copies instance directly implies its existence in the dual chores instance.

Lemma 6. *The existence of an EF_1^1 allocation in a goods with copies constrained instance with $n - 1$ copies of each good implies the existence of an EF_1^1 allocation in the duel chores instance.*

Proof. For contradiction, assume this is false. Let A be an EF_1^1 allocation over a goods with copies instance with $n - 1$ copies of each good, and let B be the allocation induced from A in the duel chores instance (For every good g with $n - 1$ copies in the original instance, there is a chore c_g in the duel instance with a value of $v_i(c_g) = -v_i(g)$ for all $i \in N$. The allocation B is constructed by allocating each chore c_g to the unique agent $j \in N$ such that $g \notin A_j$). Assume that B is not EF_1^1 , so there exists a pair of agents $i, j \in N$, such that for all $c_i \in B_i, c_j \notin B_j$, we have that $v_i(B_i \setminus \{c_i\}) < v_i(B_j \cup \{c_j\})$.

Let the set S represent a set containing a single copy of each good in the original goods with copies instance. From the way we constructed the duel chores instance, we must have that $v_i(B_i) = v_i(A_i) - v_i(S)$ for all $i \in N$. Also note that in A , we must have that $v_i(A_i \cup \{g_i\}) \geq v_i(A_i \setminus \{g_j\})$ for some goods $g_i \notin A_i, g_j \in A_j$.

This fact implies that we can take the two corresponding chores in the duel instance $c_{g_i} \in A_i, c_{g_j} \notin A_j$, such that $v_i(c_{g_i}) = -v_i(g_i)$ and $v_i(c_{g_j}) = -v_i(g_j)$, and from this, we can rewrite the i, j EF_1^1 guarantee from the goods with copies instance as the following:

$$v_i(B_i) + v_i(S) - v_i(c_{g_i}) \geq v_i(B_j) + v_i(S) + v_i(c_{g_j})$$

removing the $v_i(S)$ from both sides get us

$$v_i(B_i) - v_i(c_{g_i}) \geq v_i(B_j) + v_i(c_{g_j})$$

which means that $v_i(B_i \setminus \{c_{g_i}\}) \geq v_i(B_j \cup \{c_{g_j}\})$, giving a contradiction. \square

4.4 Extensions to Broader Constraints

Note the missing guarantee of EF_1^1 when balancedness is imposed in [Theorem 10](#). Our approach closely follows the decomposition of a fractional CEEI allocation in the *unconstrained* case due to [Aziz et al. \[2024\]](#); however, there are key differences that we highlight. While our Prop1 and PO guarantees follow directly from lemmas they prove ([Proposition 1](#) and [Lemma 4](#)), our EF_1^1 guarantee does not. They use a property of fractional MNW allocations, which are CEEI in the unconstrained case. We need to establish an appropriately relaxed version of this property ([Lemma 5](#)) for CEEI allocations in our constrained case, which may be of independent interest. Finally, our algorithm is similar to theirs, with a slightly different \mathcal{H}_2 hierarchy in the bihierarchy rounding of [Budish et al. \[2013\]](#) due to working with goods with copies. As stated in the proof of [Theorem 10](#), \mathcal{H}_2 will enforce that each integral allocation in the support of the CEEI will meet the goods with copies constraints, while \mathcal{H}_1 enforces that if the CEEI is also balanced, then each integral allocation in the support will be balanced as well.

Goods with copies and balancedness constrains both have simple structures that allow their enforcement to be integrated easily into the existing structure of the hierarchies from the Utility++ guarantee. Unfortunately, this does not seem to be the case for more general constraints. Take partition matroids for example. The CEEI theorem of [Echenique et al. \[2021\]](#) is general enough to allow us to find a fractional CEEI that meets a given set of partition matroid constraints, however, there is no way that we know of to break down this fractional allocation into a convex combination of integral allocations which also meet the given constraints.

The easiest way to attempt to do this would be to add additional constraints to \mathcal{H} to ensure that the partition matroid constraints are met. For a given partition matroid \mathcal{C} , such a constraint set would look like:

$$\mathcal{H}_3 = \left\{ 0 \leq \sum_{j \in \mathcal{C}} y_{i,j} \leq h_{\mathcal{C}} \mid i \in N, \mathcal{C} \in \mathcal{C} \right\},$$

But unlike balancedness or goods with copies, this new constraint set does not fit into one of the existing hierarchies \mathcal{H}_1 and \mathcal{H}_2 . Enforcing them in addition to the other two hierarchies (which are both necessary for the Utility Guarantee++ lemma to work) would create a tri-hierarchy constraint, something which is not covered by the results of [Budish et al. \[2013\]](#).

5 Discussion

Our work is only the first step towards a systematic study of fair and efficient allocations under feasibility constraints. In addition to resolving the major open questions listed in [Section 3.2](#), the field pursue several exciting research directions.

Computation. We did not discuss how to compute MNW allocations subject to any of the various feasibility constraints we study. It is likely a hard problem in each case, as it is for the unconstrained case [[Lee, 2015](#)]. For the unconstrained case, the literature offers polynomial-time constant-approximations of the Nash welfare objective (the current best being a 1.45-approximation due to [Barman et al. \[2018\]](#)), often together with (approximate) fairness guarantees [[McGlaughlin and Garg, 2020](#), [Chaudhury et al., 2022](#)]. Such a result under feasibility constraints would be exciting.

Pushing for stronger BoBW guarantees. In contrast to our deterministic guarantees from [Section 3](#), our BoBW guarantees from [Section 4](#) hold only for the special cases of goods with copies and/or balancedness. Can such guarantees be established even for partition matroid constraints? We also understand the limits of fractional allocations a lot less in the constrained setting. To that end, we make an observation for linearly-constrained divisible economies ([Definition 7](#)). While CEEI allocations exist and achieve EF and PO ([Theorem 9](#)), fractional MNW allocations, which may not be CEEI, are still $\frac{1}{2}$ -EF and PO. This is because [Ebadian et al. \[2024\]](#) establish that fractional MNW allocations among a convex constraint set have an individual harm ratio of 1, and this immediately implies $\frac{1}{2}$ -EF, generalizing a recent result by [Tröbst and Vazirani \[2024\]](#) that a fractional MNW *matching* (i.e., when $m = n$ and balancedness is imposed) is $\frac{1}{2}$ -EF and PO. It is also worth noting CEEI allocations can achieve the appealing guarantees of EF and PO, there are desirable fairness and efficiency guarantees that cannot be achieved simultaneously, even by fractional allocations under a partition matroid constraint [[Kawase et al., 2023](#)].

Extending EF1_{WC} beyond goods with copies. EF1_{WC} is an appealing strengthening of EF1, but defined narrowly for the case of goods with copies. Can we define viable strengthenings of EF1 for more general settings (e.g., partition matroid constraints, or even arbitrary matroid constraints), which reduce to EF1_{WC} for the special case of goods with copies?

Acknowledgements

This research was partially supported by an NSERC Discovery grant.

References

- Haris Aziz, Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of both worlds: Ex ante and ex post fairness in resource allocation. *Operations Research*, 72(4):1674–1688, 2024.
- Siddharth Barman and Sanath Kumar Krishnamurthy. On the proximity of markets with integral equilibria. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1748–1755, 2019.
- Siddharth Barman and Mashbat Suzuki. Compatibility of fairness and nash welfare under subadditive valuations. arXiv:2407.12461, 2024.
- Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 557–574, 2018.
- Nawal Benabbou, Mithun Chakraborty, Xuan-Vinh Ho, Jakub Sliwinski, and Yair Zick. The price of quota-based diversity in assignment problems. *ACM Transactions on Economics and Computation (TEAC)*, 8(3):1–32, 2020.
- Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, and Yair Zick. Finding fair and efficient allocations for matroid rank valuations. *ACM Transactions on Economics and Computation*, 9(4): 1–41, 2021.
- Arpita Biswas and Siddharth Barman. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 91–97, 2018.
- Arpita Biswas and Siddharth Barman. Matroid constrained fair allocation problem. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 9921–9922, 2019.
- Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, 2013.
- Eric Budish, Gérard P. Cachon, Judd B. Kessler, and Abraham Othman. Course Match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research*, 65(2):314–336, 2017.
- 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 17th ACM Conference on Economics and Computation (EC)*, pages 305–322, 2016.
- 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*, 7(3): Article 12, 2019.

- Ioannis Caragiannis, Evi Micha, and Nisarg Shah. A little charity guarantees fair connected graph partitioning. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pages 4908–4916, 2022.
- Bhaskar Ray Chaudhury, Yun Kuen Cheung, Jugal Garg, Naveen Garg, Martin Hoefer, and Kurt Mehlhorn. Fair division of indivisible goods for a class of concave valuations. *Journal of Artificial Intelligence Research*, 74:111–142, 2022.
- Yann Chevaleyre, Paul E. Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaitre, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A. Rodríguez-Aguilar, and Paulo Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.
- Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 629–646, 2017.
- Amitay Dror, Michal Feldman, and Erel Segal-Halevi. On fair division under heterogeneous matroid constraints. *Journal of Artificial Intelligence Research*, 76:567–611, 2023.
- Soroush Ebadian, Dominik Peters, and Nisarg Shah. How to fairly allocate easy and difficult chores. In *Proceedings of the 21st International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 372–380, 2022.
- Soroush Ebadian, Rupert Freeman, and Nisarg Shah. Harm ratio: A novel and versatile fairness criterion. In *Proceedings of the 4th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*, 2024. Forthcoming.
- Federico Echenique, Antonio Miralles, and Jun Zhang. Constrained pseudo-market equilibrium. *American Economic Review*, 111(11):3699–3732, 2021.
- Brandon Fain, Kamesh Munagala, and Nisarg Shah. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 575–592, 2018.
- Yotam Gafni, Xin Huang, Ron Lavi, and Inbal Talgam-Cohen. Unified fair allocation of goods and chores via copies. *ACM Transactions on Economics and Computation*, 11(3): article 10, 2023.
- Jiarui Gan, Bo Li, and Xiaowei Wu. Approximation algorithm for computing budget-feasible $ef1$ allocations. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 170–178, 2023.
- Jugal Garg, Aniket Murhekar, and John Qin. Fair and efficient allocations of chores under bivalued preferences. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5043–5050, 2022.
- Jugal Garg, Aniket Murhekar, and John Qin. Fair division of indivisible chores via earning restricted equilibria. arXiv:2407.03318, 2024.
- Naveen Garg, Telikepalli Kavitha, Amit Kumar, Kurt Mehlhorn, and Julián Mestre. Assigning papers to referees. *Algorithmica*, 58:119–136, 2010.

- Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. Fair and efficient allocations under lexicographic preferences. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5472–5480, 2021.
- Ayumi Igarashi. How to cut a discrete cake fairly. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5681–5688, 2023.
- Yasushi Kawase, Hanna Sumita, and Yu Yokoi. Random assignment of indivisible goods under constraints. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2792–2799, 2023.
- Euiwoong Lee. APX-hardness of maximizing Nash social welfare with indivisible items. arXiv:1507.01159, 2015.
- Peter McGlaughlin and Jugal Garg. Improving nash social welfare approximations. *Journal of Artificial Intelligence Research*, 68:225–245, 2020.
- James G. Oxley. *Matroid Theory*. Oxford University Press, second edition, 2011.
- Alexandros Psomas and Paritosh Verma. Fair and efficient allocations without obvious manipulations. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 13342–13354, 2022.
- Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer, 2003.
- Nisarg Shah. Spliddit: two years of making the world fairer. *XRDS: Crossroads, The ACM Magazine for Students*, 24(1):24–28, 2017.
- Warut Suksompong. Constraints in fair division. *ACM SIGecom Exchanges*, 19(2):46–61, 2021.
- Thorben Tröbst and Vijay V Vazirani. Cardinal-utility matching markets: The quest for envy-freeness, pareto-optimality, and efficient computability. arXiv:2402.08851, 2024.
- Xiaowei Wu, Bo Li, and Jiarui Gan. Budget-feasible maximum nash social welfare is almost envy-free. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 465–471, 2021.

Appendix

A Omitted Details from Section 3

A.1 Inefficiency of Round Robin

Through the following example, we will show an instance where no round-robin allocation is PO.

Example 4. Consider an instance with two agents and 8 goods, where the agents have the following preferences:

| | g_1 | g_2 | g_3 | g_4 | g_5 | g_6 | g_7 | g_8 |
|---------|-------|-------|-------|-------|-------|-------|-------|-------|
| Agent 1 | 10 | 9 | 5 | 4 | 3 | 2 | 1 | 0 |
| Agent 2 | 10 | 9 | 8 | 7 | 6 | 5 | 1 | 0 |

Observe the best possible bundle that each agent can receive under a round-robin ordering. Since both of the agents have the same strict ordering of the goods ($v_i(g_1) > v_i(g_2) > \dots > v_i(g_8)$ for both $i \in [2]$), the best possible round-robin bundle for each agent will be the bundle $B = \{g_1, g_3, g_5, g_7\}$. For agent 1, we have $v_1(B) = 19$, and for agent 2, we have $v_2(B) = 25$.

However, consider the allocation $A = (\{g_1, g_2, g_7, g_8\}, \{g_3, g_4, g_5, g_6\})$. Agents preferences in A will be $v_1(A_1) = 20 > v_1(B)$ and $v_2(A_2) = 26 > v_2(B)$, showing that this allocation will Pareto Dominate any round-robin allocation.

A.2 Lower Bound for MNW Fairness Guarantees with Good with Copies Constraints

Theorem 12. For any $\epsilon > 0$, there exist instances with the goods-with-copies constraint in which no MNW (or complete MNW) allocation is $(1/2 + \epsilon)$ -EF1.

Proof. Let $k \in \mathbb{N}$ and fix $\delta < 1/k$. Consider the following instance with 2 agents and $2k$ goods. The goods are partitioned into S and T , each consisting of k goods. There are 2 copies of each of the k goods in S , and for each good $g_s \in S$, we have $v_1(g_s) = \delta$ and $v_2(g_s) = 1$. There is a single copy of each good $g_t \in T$, with agent valuations being $v_1(g_t) = v_2(g_t) = 1$.

Consider any allocation A . First, notice that since there are two copies of every good in S , the goods-with-copies constraint enforces that each agent receive one copy of each good in S . Suppose $x = |A_1 \cap T|$. Then, we have $v_1(A_1) = k\delta + x$ and $v_2(A_2) \leq k + (k - x) = 2k - x$. Thus, $NW(A) \leq (k\delta + x) \cdot (2k - x)$. The right hand side is uniquely maximized at $x = k - \frac{k\delta}{2}$, and, given the choice of δ , it is easy to see that the integral value of x that uniquely maximizes it is $x = k$: $x = k$ yields a value of $k^2 \cdot (1 + \delta)$, and $(k\delta + x) \cdot (2k - x) \geq k^2(1 + \delta)$ implies $x \in [k - k\delta, k]$, with k being the only integer in that range due to our choice of δ . Consider the allocation $A^* = (S \cup T, S)$, corresponding to $x = k$, in which agent 1 receives one copy of each good in S along with all of T while agent 2 receives one copy of each good in S . Since this achieves the maximum right hand side value, it must be the unique MNW (and complete MNW) allocation in this instance.

In this allocation, agent 1 is not envious, but, for agent 2, we have $v_2(A_2^*) = k$ and $v_2(A_1^* \setminus \{g\}) = 2k - 1$ for every $g \in A_1^*$. Hence, this allocation is $\frac{k}{2k-1}$ -EF1, i.e., $(\frac{1}{2} + \frac{1}{4k-2})$ -EF1. It is clear that for any $\epsilon > 0$, we can choose $k > \frac{1}{2} + \frac{1}{4\epsilon}$ to ensure that the allocation is not $(1/2 + \epsilon)$ -EF1. \square

B Matroid Preprocessing Step From Dror et al. [2023]

The process of Dror et al. [2023] involves performing a number of “Free Extensions” on the instance. For some matroid $\mathcal{M} = (M, \mathcal{I})$ with rank r , performing a free extension on \mathcal{M} means creating a new matroid $\mathcal{M}' = (M', \mathcal{I}')$ by adding a new good x , such that $M' = M \cup \{x\}$ and $\mathcal{I}' = \{I \cup \{x\} \mid I \in \mathcal{I}, |I| < r\}$, i.e., it adds x to all independent sets that have size less than the rank of the matroid. Note that a free-extension will never increase the rank of the matroid.

The preprocessing step will keep performing free extensions on \mathcal{M} , until we end up with a new matroid $\mathcal{M}' = (M', \mathcal{I}')$ where $|M'| = nr$. Each free extension will involve us adding a “dummy good” to M (a good that each agent has a valuation of 0 for). One can immediately see that since the rank of the of \mathcal{M}' is r , any complete and feasible allocation of M' must give exactly r goods to each agent.

Dror et al. [2023] prove that the matroid \mathcal{M}' constructed in the preprocessing step provides the following guarantees:

Observation 1 (Observation 2 of Dror et al. [2023]). The matroid \mathcal{M}' constructed in the preprocessing step provides the following guarantees:

- All bases of \mathcal{M} are bases of \mathcal{M}'
- For every complete allocation $A' \in \mathcal{F}'$, $|A'_i| = r$ for all $i \in N$. Therefore, every A'_i is a basis of \mathcal{M}' .
- For every complete allocation $A \in \mathcal{F}$, there is a complete allocation $A' \in \mathcal{F}'$, where for all $i \in N$, A'_i contains A_i plus zero or more dummy goods, Therefore, $v_j(A_i) = v_j(A'_i)$ for all $j \in N$.
- For every complete allocation $A' \in \mathcal{F}'$, there is a complete allocation $A \in \mathcal{F}$, where for all $i \in N$, A_i equals A'_i with all dummy goods removed. Therefore, $v_j(A'_i) = v_j(A_i)$ for all $j \in N$.
- \mathcal{M}' is base-orderable if and only if \mathcal{M} is base-orderable.

Let \mathcal{F} be the feasibility set of an instance constrained by a base-orderable matroid \mathcal{M} , and let \mathcal{F}' be the feasibility set of the instance constrained by \mathcal{M}' , the matroid constructed by performing the preprocessing step on \mathcal{M} .

We can say that an allocation $A' \in \mathcal{F}'$ will be a complete MNW allocation for \mathcal{F}' if and only if the corresponding allocation $A \in \mathcal{F}$, constructed by removing all the dummy goods from A' , is a complete MNW allocation for \mathcal{F} .

First, in the case where an MNW allocation $A' \in \mathcal{F}'$ has $NW(A') > 0$, this follows from the fact that all agents’ utilities remain the same after dummy goods are removed. Thus, for the allocation $A \in \mathcal{F}$ formed by removing the dummy goods from A' , we must have $NW(A) = NW(A')$. If there were another allocation $A^* \in \mathcal{F}$ such that $NW(A^*) > NW(A)$, then we could add dummy goods to each bundle of A^* until $|A^*_i| = r$ for all $i \in N$. This new allocation must be in \mathcal{F}' , and must have higher Nash Welfare than A' , giving a contradiction. Conversely, if $A \in \mathcal{F}$ maximizes Nash Welfare for \mathcal{F} , with $NW(A) > 0$, but its corresponding allocation $A' \in \mathcal{F}'$ does not, then the same contradiction can be arrived at by finding the MNW allocation for \mathcal{F}' and removing the dummy goods.

Next, in the case where an MNW allocation $A' \in \mathcal{F}'$ has $NW(A') = 0$, it must be the case that A' maximizes the number of agents that receive positive utility, and with respect to this, maximizes the product of positive utilities. Notice that after removing the dummy goods from A' , we get an allocation A where the exact same agents get positive utility, and the product of those agents' utilities are the same. If there is some $A^* \in \mathcal{F}$ where $NW(A^*) > 0$, then that gives a contradiction, as adding dummy goods to A^* would create an allocation in \mathcal{F}' with positive Nash Welfare. If there is an allocation in A^* where more agents receive positive utility, or the product of positive utilities is higher, then we can add dummy goods to create an allocation from \mathcal{F}' with the same guarantees, giving another contradiction. Conversely, if $A \in \mathcal{F}$ is the MNW allocation for \mathcal{F} and $NW(A) = 0$, but the corresponding allocation $A' \in \mathcal{F}'$ is not MNW, then the same contradictions can be arrived at by finding the MNW allocation for \mathcal{F}' and removing the dummy goods.