

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-matroid constraints (which includes balancedness). 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 (Chevalere 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, Peters, and Shah 2022; Garg, Murhekar, and Qin 2022), or for allocating public goods (Fain, Munagala, and Shah 2018; Ebadian, Freeman, and Shah 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).

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.

- A famous non-matroid 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, Micha, and Shah 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. One 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, Li, and Wu (2023) show that an EF1 allocation always exists, and Wu, Li, and Gan (2021) show that any MNW allocation subject to such a constraint is $\frac{1}{4}$ -EF1 and PO. Another exception is the work of Shoshan, Hazon, and Segal-Halevi (2023), which shows that in instances with partition matroid constraints and two agents, it is always possible to find an allocation that is EF1+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, 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, we show that the $\frac{1}{2}$ -EF1 and PO guarantees can be extended to MNW allocations subject to a class of non-matroid 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, Miralles, and Zhang (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 $i, j \in M$. 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 constraint. 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 $g \in S$, $(S \setminus \{g\}) \cup \{f(g)\}$ and $(T \setminus \{f(g)\}) \cup \{g\}$ 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-matroid constraints. We will also study the following constraints, which 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.

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 $\alpha \cdot v_i(B_i) \geq v_i(A_i)$ for each $i \in N$ and at least one inequality is strict. In words, no other feasible allocation should be able to make an agent happier

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 $\prod_{i \in N} v_i(A_i)$.

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. While the proof outline is somewhat straightforward, there are a number of subtleties involved; the detailed proof is provided in the full version of this paper.⁴

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

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.

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

⁴<https://arxiv.org/abs/2411.00133>

	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 allocations. 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, Feldman, and Segal-Halevi (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 the full version.

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:

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

where the second transition follows because

$$\begin{aligned} & v_i(A_j) - v_i(A_i) \\ &= (v_i(A_j^*) - v_i(A_i^*)) + (v_i(A_j \setminus A_j^*) - v_i(A_i \setminus A_i^*)) \end{aligned}$$

and

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

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

Now, we can observe that

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

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:

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

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

In the full version, we provide a matching lower-bound, showing that MNW cannot guarantee better than $\frac{1}{2}$ -EF1 for matroid-constrained instances.

3.1 Beyond Matroid Constraints

One may be tempted to apply the trick from Theorem 1 of reducing to unconstrained instances for even non-matroid constraints. Even if this induces valuations from the broader class of subadditive valuations,⁵ one can use a recent result by Barman and Suzuki (2024) to establish the existence

⁵Valuation v is subadditive if $v(S \cup T) \leq v(S) + v(T)$ for all $S, T \subseteq M$.

of EF1 and $\frac{1}{2}$ -PO allocations. However, many non-matroid constraints induce valuations that are not even subadditive.

For certain classes of non-matroid constraints, we introduce an approach of *alternate worlds construction* to establish how MNW can still be used to prove fairness and efficiency guarantees. The following lemma, whose proof is included in the full version, lays out the approach.

Lemma 1. *Let $\{\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k\}$ be a collection 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.*

We can apply Lemma 1 to Theorems 1 and 2 to obtain $\frac{1}{2}$ -EF1 and PO for a broad class of non-matroid 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.*

Consider partition matroids with lower bounds constraints, 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 . 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. We show that for *complete* allocations, the non-matroid constraint imposed by a partition matroid with lower bounds can be represented as a union of partition matroid constraints.

Note that for this particular class of constraint, the same fairness and efficiency guarantees can be achieved through a reduction to a laminar matroid, which is outlined in the full version of this paper. We include this to highlight the alternate worlds technique, and to show how this constraint can be thought of in terms of partition matroids, which are much simpler to work with than laminar matroids, and may prove to be useful for future work on this class of constraints.

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 $t \in [k]$.*

The proof of Lemma 2 appears in the full version. 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; the proof appears in the full version.

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

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.

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) suggest EF1_{WC} as an appealing strengthening of EF1, establishing its existence for certain restricted valuation classes.

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-free 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.

We show that $\frac{1}{2}$ -EF1 of MNW and complete MNW allocations can be improved to $\frac{1}{2}$ -EF1_{WC}, establishing the first (approximate) EF1_{WC} existence guarantee for general additive valuations. The proof appears in the full version.

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

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 $i \in N$. Round robin achieves exact EF1 under balancedness (Caragiannis et al. 2019), but we show in the full version that it can be 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.

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, Li, and Gan (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 Stronger Guarantees

We next consider strengthening EF1 to *stochastically-dominant envy-freeness up to one good* (SD-EF1) (Aziz et al. 2024), which demands that EF1 hold with respect to any valuations that induce the same ordinal preferences, and PO to *unconstrained Pareto optimality* (PO⁺), a new criterion we consider which demands that not even an *infeasible* allocation Pareto dominate the given feasible allocation. The formal definitions of these criteria along with the proofs of the following two results appear in the full version.

While PO⁺ by itself can be achieved subject to balancedness for unconstrained goods division (Conitzer, Freeman, and Shah 2017) and chore division (Garg, Murhekar, and Qin 2024) when all agents have non-zero valuations, it cannot be achieved together with any positive approximation of EF1. For the restricted class of bivalued preferences, though, Garg, Murhekar, and Qin (2023) show that balanced EF1 and PO⁺ allocations of chores exist.

Theorem 5. *Under the balancedness constraint, it is impossible to guarantee α -EF1 and PO⁺ for any $\alpha \in (0, 1]$.*

Similarly, the round robin algorithm produces balanced SD-EF1 allocations of goods and chores (Aziz et al. 2024), but SD-EF1 is unachievable for general partition matroid constraints.

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

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 ,⁶ and that x satisfies a desideratum D *ex post* if every (integral) allocation in its support satisfies D .

EF+PO fractional allocations under linear constraints.

First, we define (a relevant special case of) the model of Echenique, Miralles, and Zhang (2021).

Definition 5 (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 6 (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.,

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

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 $p_g = 0$.

The following is a corollary of Theorem 1 of Echenique, Miralles, and Zhang (2021).⁷

Theorem 7 (Echenique, Miralles, and Zhang 2021). *For a linearly-constrained divisible economy with strictly positive*

⁶We 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.

⁷Their stronger result holds for semi-strictly quasi-concave utilities, which subsumes strictly positive additive utilities.

Algorithm 1 CE Lottery for goods-with-copies

- 1: $x \leftarrow$ a competitive (fractional) allocation of Theorem 7 due to Echenique, Miralles, and Zhang (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 two sets of constraints for an integral allocation y (see full version for an explanation of this):

$$\mathcal{H}_1 = \{ \lfloor Q_{i,t} \rfloor \leq \sum_{j \in [t]} y_{i,\sigma_i(j)} \leq \lceil Q_{i,t} \rceil \mid i \in N, t \in [m] \},$$

$$\mathcal{H}_2 = \{ \sum_{i \in N} y_{i,g} \leq q_g \mid g \in M \}.$$

- 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 $\sum_{\ell=1}^L \lambda_\ell \cdot y^\ell$
-

additive valuations,⁸ 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 (EF)⁹ and Pareto optimal.

We 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, Freeman, and Shah 2017) and EF₁¹ (Barman and Krishnamurthy 2019).

Definition 7 (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 8 (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 decompose a fractional CEEI allocation x using Algorithm 1 to obtain the following guarantee.

Theorem 8. *For goods with copies (resp., goods with copies with an additional balancedness constraint), there always exists a randomized allocation that is ex ante EF and PO, and ex post EF₁¹, Prop1, and PO (resp., Prop1 and PO).*

Note the missing guarantee of EF₁¹ when balancedness is imposed. 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. First, since Aziz et al. (2024) work with a single copy of goods and chores whereas we work with an arbitrary number of copies of each good, Theorem 8 generalizes their results for both goods (by having a single copy of each good), and chores (by having $n - 1$ copies of each good). Next, while our Prop1 and PO guarantees follow directly from lemmas they prove (See the full version for a detailed explanation of these lemmas), our EF₁¹ guarantee does not.

⁸That is, $v_i(g) > 0$ for all $i \in N$ and $g \in M$.

⁹An allocation is envy-free if $v_i(x_i) \geq v_i(x_j)$ for all $i, j \in N$.

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 for CEEI allocations in our constrained case, which may be of independent interest, and is outlined in detail in the full version of this paper. Finally, our algorithm is similar to theirs, using the so called “bihierarchy rounding” of Budish et al. (2013), but with a slightly different \mathcal{H}_2 constraint due to working with goods with copies. We include a complete proof of Theorem 8, along with an introduction of the relevant concepts and techniques and an explanation of why EF₁¹ is missing for the case of balancedness, in the full version.

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, Krishnamurthy, and Vaish (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 5). While CEEI allocations exist and achieve EF and PO (Theorem 7), fractional MNW allocations, which may not be CEEI, are still $\frac{1}{2}$ -EF and PO. This is because Ebadian, Freeman, and Shah (2024) establish that fractional MNW allocations are 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 that while 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, Sumita, and Yokoi 2023).

Extending EF₁¹ beyond goods with copies. EF₁¹ 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 EF₁¹ for the special case of goods with copies?

Acknowledgements

This research was partially supported by an NSERC Discovery grant.

References

- Aziz, H.; Freeman, R.; Shah, N.; and Vaish, R. 2024. Best of both worlds: Ex ante and ex post fairness in resource allocation. *Operations Research*, 72(4): 1674–1688.
- Barman, S.; and Krishnamurthy, S. K. 2019. On the proximity of markets with integral equilibria. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 1748–1755.
- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Finding Fair and Efficient Allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 557–574.
- Barman, S.; and Suzuki, M. 2024. Compatibility of Fairness and Nash Welfare under Subadditive Valuations. arXiv:2407.12461.
- Benabbou, N.; Chakraborty, M.; Ho, X.-V.; Sliwinski, J.; and Zick, Y. 2020. The price of quota-based diversity in assignment problems. *ACM Transactions on Economics and Computation (TEAC)*, 8(3): 1–32.
- Benabbou, N.; Chakraborty, M.; Igarashi, A.; and Zick, Y. 2021. Finding fair and efficient allocations for matroid rank valuations. *ACM Transactions on Economics and Computation*, 9(4): 1–41.
- Biswas, A.; and Barman, S. 2018. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 91–97.
- Biswas, A.; and Barman, S. 2019. Matroid constrained fair allocation problem. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 9921–9922.
- Budish, E.; Cachon, G. P.; Kessler, J. B.; and Othman, A. 2017. Course Match: A Large-Scale Implementation of Approximate Competitive Equilibrium from Equal Incomes for Combinatorial Allocation. *Operations Research*, 65(2): 314–336.
- Budish, E.; Che, Y.-K.; Kojima, F.; and Milgrom, P. 2013. Designing Random Allocation Mechanisms: Theory and Applications. *American Economic Review*, 103(2): 585–623.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation*, 7(3): Article 12.
- Caragiannis, I.; Micha, E.; and Shah, N. 2022. A Little Charity Guarantees Fair Connected Graph Partitioning. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 4908–4916.
- Chaudhury, B. R.; Cheung, Y. K.; Garg, J.; Garg, N.; Hoefer, M.; and Mehlhorn, K. 2022. Fair division of indivisible goods for a class of concave valuations. *Journal of Artificial Intelligence Research*, 74: 111–142.
- Chevalleyre, Y.; Dunne, P. E.; Endriss, U.; Lang, J.; Lemaitre, M.; Maudet, N.; Padget, J.; Phelps, S.; Rodríguez-Aguilar, J. A.; and Sousa, P. 2006. Issues in Multiagent Resource Allocation. *Informatica*, 30: 3–31.
- Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair Public Decision Making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 629–646.
- Dror, A.; Feldman, M.; and Segal-Halevi, E. 2023. On fair division under heterogeneous matroid constraints. *Journal of Artificial Intelligence Research*, 76: 567–611.
- Ebadian, S.; Freeman, R.; and Shah, N. 2024. 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)*. Forthcoming.
- Ebadian, S.; Peters, D.; and Shah, N. 2022. How to Fairly Allocate Easy and Difficult Chores. In *Proceedings of the 21st International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 372–380.
- Echenique, F.; Miralles, A.; and Zhang, J. 2021. Constrained Pseudo-Market Equilibrium. *American Economic Review*, 111(11): 3699–3732.
- Fain, B.; Munagala, K.; and Shah, N. 2018. Fair Allocation of Indivisible Public Goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 575–592.
- Gafni, Y.; Huang, X.; Lavi, R.; and Talgam-Cohen, I. 2023. Unified Fair Allocation of Goods and Chores via Copies. *ACM Transactions on Economics and Computation*, 11(3): article 10.
- Gan, J.; Li, B.; and Wu, X. 2023. Approximation Algorithm for Computing Budget-Feasible EF1 Allocations. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 170–178.
- Garg, J.; Murhekar, A.; and Qin, J. 2022. Fair and efficient allocations of chores under bivalued preferences. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 5043–5050.
- Garg, J.; Murhekar, A.; and Qin, J. 2023. New algorithms for the fair and efficient allocation of indivisible chores. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2710–2718.
- Garg, J.; Murhekar, A.; and Qin, J. 2024. Fair Division of Indivisible Chores via Earning Restricted Equilibria. arXiv:2407.03318.
- Garg, N.; Kavitha, T.; Kumar, A.; Mehlhorn, K.; and Mestre, J. 2010. Assigning papers to referees. *Algorithmica*, 58: 119–136.
- Hosseini, H.; Sikdar, S.; Vaish, R.; and Xia, L. 2021. Fair and efficient allocations under lexicographic preferences. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 5472–5480.
- Igarashi, A. 2023. How to cut a discrete cake fairly. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 5681–5688.

- Kawase, Y.; Sumita, H.; and Yokoi, Y. 2023. Random Assignment of Indivisible Goods under Constraints. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2792–2799.
- Lee, E. 2015. APX-Hardness of Maximizing Nash Social Welfare with Indivisible Items. arXiv:1507.01159.
- McGlaughlin, P.; and Garg, J. 2020. Improving Nash social welfare approximations. *Journal of Artificial Intelligence Research*, 68: 225–245.
- Oxley, J. G. 2011. *Matroid Theory*. Oxford University Press, second edition.
- Psomas, A.; and Verma, P. 2022. Fair and Efficient Allocations Without Obvious Manipulations. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 13342–13354.
- Shah, N. 2017. Spliddit: two years of making the world fairer. *XRDS: Crossroads, The ACM Magazine for Students*, 24(1): 24–28.
- Shoshan, H.; Hazon, N.; and Segal-Halevi, E. 2023. Efficient Nearly-Fair Division with Capacity Constraints.
- Suksompong, W. 2021. Constraints in fair division. *ACM SIGecom Exchanges*, 19(2): 46–61.
- Tröbst, T.; and Vazirani, V. V. 2024. Cardinal-Utility Matching Markets: The Quest for Envy-Freeness, Pareto-Optimality, and Efficient Computability. arXiv:2402.08851.
- Wu, X.; Li, B.; and Gan, J. 2021. Budget-feasible Maximum Nash Social Welfare is Almost Envy-free. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 465–471.