

Fair Division Under Cardinality Constraints

Arpita Biswas and Siddharth Barman

Indian Institute of Science

arpitab@iisc.ac.in, barman@iisc.ac.in

Abstract

We consider the problem of *fairly* allocating indivisible goods, among agents, under cardinality constraints and additive valuations. In this setting, we are given a partition of the entire set of goods—i.e., the goods are categorized—and a limit is specified on the number of goods that can be allocated from each category to any agent. The objective here is to find a *fair* allocation in which the subset of goods assigned to any agent satisfies the given cardinality constraints. This problem naturally captures a number of resource-allocation applications, and is a generalization of the well-studied (unconstrained) fair division problem.

The two central notions of fairness, in the context of fair division of indivisible goods, are *envy freeness up to one good* (EF1) and (approximate) *maximin share guarantee* (MMS). We show that the existence and algorithmic guarantees established for these solution concepts in the unconstrained setting can essentially be achieved under cardinality constraints.

Furthermore, focusing on the case wherein all the agents have the same additive valuation, we establish that EF1 allocations exist even under *matroid constraints*.

1 Introduction

A large body of recent work in algorithmic game theory, artificial intelligence, and computational social choice has been directed towards understanding the problem of allocating indivisible goods among agents in “fair” manner; see, e.g., [Brandt *et al.*, 2016] and [Endriss, 2017] for excellent expositions. This recent focus on indivisible goods is motivated, in part, by applications (such as division of inheritance and partitioning computational resources in a cloud computing environment) which inherently entail allocation of resources that cannot be fractionally allocated. In fact, algorithms developed for finding fair allocations of indivisible goods have been implemented in specific settings; for instance, Course Match [Budish *et al.*, 2016] is employed for course allocation at the Wharton School in the University of Pennsylvania

and the website Spliddit (www.spliddit.org) [Goldman and Procaccia, 2014] provides online access to fair division algorithms.

Note that, though the theory of fair division is extensive, classical notions of fairness—such as *envy freeness*¹—typically address allocation of divisible goods and are not representative in the indivisible setting. For instance, while an envy-free allocation of divisible goods is guaranteed to exist [Stromquist, 1980], such an existence result does not hold for indivisible goods.²

Motivated by these considerations, recent results have formulated and studied solution concepts for fairly allocating indivisible goods [Budish, 2011; Procaccia and Wang, 2014; Bouveret and Lemaître, 2014]. Arguably, the two most prominent notions of fairness in this context are (i) *envy freeness up to one good* (EF1) and (ii) the *maximin share guarantee* (MMS). These solution concepts were defined by Budish [Budish, 2011] and they, respectively, provide a cogent analogue of envy-freeness and *proportionality*³ in the context of indivisible goods:

- An allocation is said to be EF1 if every agent values her bundle at least as much as any other agent’s bundle, up to the removal of the most valuable good from the other agent’s bundle. EF1 allocations are guaranteed to exist; by contrast, for indivisible goods envy-free allocations might not exist. Another attractive feature of EF1 is that it is computationally tractable: even under combinatorial valuations, EF1 allocations can be found efficiently [Lipton *et al.*, 2004]. Furthermore, under additive valuations, this notion of fairness is compatible with *Pareto efficiency* [Caragiannis *et al.*, 2016].
- An allocation is said to satisfy MMS if each agent receives a bundle of value at least as much as her *maximin share*. These shares are defined as the maximum value that an agent can guarantee for herself if she were to

¹An allocation is said to be envy-free if each agent values her bundle at least as much as she values any other agent’s bundle [Foley, 1967; Varian, 1974].

²If we have a single indivisible good and two agents, then in any allocation, the losing agent is bound to be envious.

³An allocation is said to be proportionally fair among n agents, if every agent gets a bundle of value at least $1/n$ times her value for the entire set of goods.

partition the set of goods into n bundles and then, from those bundles, receive the minimum valued one; here, n is the total number of agents. This can be interpreted via an application of the standard *cut-and-choose* protocol over indivisible goods: if agent i is (hypothetically) asked to partition the set of goods into n bundles and the remaining $(n - 1)$ agents were to select their bundles before i , then a risk-averse agent would find a partition which maximizes the least valued bundle. This value, that the agent i can guarantee for herself, is called the maximin share of i . Even though MMS allocations are not guaranteed to exist [Procaccia and Wang, 2014; Kurokawa *et al.*, 2016], it admits efficient approximation guarantees, in particular, assuring that each agent receives a bundle of value at least $2/3$ times her maximin share [Procaccia and Wang, 2014; Amanatidis *et al.*, 2015; Barman and Krishnamurthy, 2017].⁴ Note that these results establish an absolute guarantee, i.e., they show that an approximately maximin fair allocation always exists.

The existence and computational results developed for EF1 and MMS provide a sound understanding of fair division of indivisible goods. However, it is relevant to note that the vast majority of work in this thread of research is solely focussed on the unconstrained version of the problem.⁵ To address this limitation, and motivated by the fact that in many real-world settings the allocations are required to satisfy certain criteria, we study a relevant, constrained version of the fair division problem.

In particular, we consider a setting wherein the indivisible goods are categorized and a limit is specified on the number of goods that can be allocated from each category to any agent. Here, the objective is to find a fair allocation in which the subset of goods assigned to any agent satisfies the given cardinality constraints. We shall see that this corresponds to a fair allocation problem under a *partition matroid constraint*.

The following stylized example—adapted from [Gourvès *et al.*, 2014]—demonstrates the applicability of such constraints: A museum decides to open new branches, and thereby needs to transfer some of the exhibits from the main museum to the newly opened ones. The exhibits are categorized into, say, statues, paintings, and pottery. In addition, there is an upper limit on the number of exhibits that every newly opened branch can accommodate from each category. The question now is to find a feasible division of the exhibits which is fair to the curators of each of the new branches.

Our Contributions: We establish the following results under additive valuations and cardinality (partition matroid) constraints:

1. EF1 allocations are guaranteed to exist. In particular, we develop a combinatorial algorithm which, for a given

⁴For additive valuations, Ghodsi *et al.* [Ghodsi *et al.*, 2017] provide an improved approximation guarantee of $3/4$.

⁵The work of Bouveret *et al.* [Bouveret *et al.*, 2017] along with [Gourvès and Monnot, 2017] and [Gourvès *et al.*, 2014] are notable exceptions. These results are discussed in Section 1.1.

fair division instance with additive valuations and cardinality constraints, finds an EF1 allocation in polynomial time (Theorem 1).

2. In this constrained setting, a constant-factor approximate MMS allocation always exists and can be computed in polynomial time (Theorem 2). Note that, in this setting, the value of the maximin share of each of the n agents is obtained by considering only feasible allocations (see Equation (1) in Section 2). That is, here, the maximin share of an agent is defined to be maximum value that she can guarantee for herself if she were to partition the set of goods into n bundles, each of which must satisfy the cardinality constraints, and from them receive the minimum valued one.
3. We also consider fair division subject to matroid constraints (Section 6) and show that if the agents have identical, additive valuations, then again an EF1 allocation is guaranteed to exist (Theorem 3).

1.1 Related Work

The fairness notions EF1 and MMS were defined by Budish [Budish, 2011] (see also [Moulin, 1990]), and have been extensively studied since then. These two solution concepts are incomparable, i.e., one does not imply the other [Cariannidis *et al.*, 2016]. The existence of EF1 allocations can be established via the *cycle-elimination* algorithm [Lipton *et al.*, 2004].

Fair division instances wherein the agents’ valuations are binary and additive always admit MMS allocations [Bouveret and Lemaître, 2014]. However, there exists intricate counterexamples that refutes the universal existence of MMS allocations, even under additive valuations [Procaccia and Wang, 2014; Kurokawa *et al.*, 2016]. This lead to a study of approximate maximin share allocations, α -MMS, where each agent receives a bundle whose value to her is at least $\alpha \in (0, 1)$ times her maximin share. In particular, efficient algorithms are provided to compute a $2/3$ -MMS allocation [Procaccia and Wang, 2014; Amanatidis *et al.*, 2015].

More recently, constant-factor approximation guarantees for MMS in settings wherein the valuations are not necessarily additive have been established in [Barman and Krishnamurthy, 2017] and [Ghodsi *et al.*, 2017]. Specifically, for *submodular* valuations Ghodsi *et al.* [Ghodsi *et al.*, 2017] have developed an efficient algorithm for computing $1/3$ -MMS allocations.

As mentioned previously, the work on fair division of indivisible goods is primarily confined to the unconstrained setting. Exceptions include the setting considered in [Bouveret *et al.*, 2017] and [Ferraioli *et al.*, 2014]. In [Bouveret *et al.*, 2017], they consider fair division of goods which correspond to vertices of a given graph, and the problem is to fairly allocate a connected subgraph to each agent. They also show that MMS allocations might not exist for general graphs; however, it can be efficiently computed when the underlying graph is a tree.

Ferraioli *et al.* [Ferraioli *et al.*, 2014] consider fair division problems where each agent must receive exactly k goods, for a given integer k , and provide an algorithm to efficiently com-

pute a $1/k$ -approximate MMS allocation. A different kind of matroid constraint is considered in [Gourvès and Monnot, 2017] and [Gourvès *et al.*, 2014]. In particular, their problem formulation requires that the *union* of all the allocated goods is an independent set of a given matroid. This setup is incomparable to the one considered in this paper. Specifically, while we ensure that in the computed fair allocation *each* agent’s bundle satisfies the partition matroid constraint and all the goods are allocated, these requirements are not imposed in [Gourvès and Monnot, 2017] and [Gourvès *et al.*, 2014].

2 Notation and Preliminaries

An instance of the fair division problem comprises of a tuple $\langle [m], [n], (v_i)_{i \in [n]} \rangle$, where $[m] = \{1, 2, \dots, m\}$ denotes the set of indivisible goods, $[n] = \{1, 2, \dots, n\}$ denotes the set of agents, and v_i s specify the valuation (preferences) of the agents, $i \in [n]$, over the set of goods. Throughout, we will assume that, for each agent i , the valuation $v_i : 2^{[m]} \mapsto \mathbb{R}_+$ is additive, i.e., for all agents $i \in [n]$ and subsets $S \subseteq [m]$, $v_i(S) = \sum_{g \in S} v_i(\{g\})$. Also, we have $v_i(\{g\}) \geq 0$ for all $i \in [n]$ and $g \in [m]$. For ease of notation, we will use $v_i(g)$, instead of $v_i(\{g\})$, to denote the valuation of agent i for a good g .

Write $\Pi_t(S)$ to denote the set of all t -partitions of a subset of goods $S \subseteq [m]$. An *allocation*, \mathcal{A} , refers to an n -partition $\mathcal{A} = (A_1, A_2, \dots, A_n) \in \Pi_n([m])$, where A_i is the subset of goods (bundle) allocated to agent i . In this work we focus on finding fair allocations which satisfy given cardinality constraints. Specifically, we are given a partition of the set of goods consisting of ℓ different “categories” $\{C_1, C_2, \dots, C_\ell\}$, and associated with each category $h \in \{1, 2, \dots, \ell\}$ we have a (cardinality) threshold k_h . In this setup, an allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ is said to be *feasible* iff, for every bundle A_i and category h , the cardinality constraint holds: $|A_i \cap C_h| \leq k_h$. Throughout, we will use \mathcal{F} to denote the set of feasible allocations, $\mathcal{F} := \{\mathcal{A} \in \Pi_n([m]) \mid |A_i \cap C_h| \leq k_h \text{ for all } i \in [n] \text{ and } h \in [\ell]\}$. To ensure that \mathcal{F} is nonempty, we require that the threshold $k_h \geq \frac{|C_h|}{n}$ for all categories $h \in [\ell]$.

We will use $\langle [m], [n], (v_i)_{i \in [n]}, \mathcal{F} \rangle$ to denote an instance of the fair division problem subject to cardinality constraints. Overall, our goal is to find fair allocations contained in \mathcal{F} .⁶ In this work, we provide existential and algorithmic results for the following fairness notions:

- **Envy-free up to one good (EF1):** In a fair division instance $\langle [m], [n], (v_i)_{i \in [n]}, \mathcal{F} \rangle$, an allocation $\mathcal{A} = (A_1, A_2, \dots, A_n) \in \mathcal{F}$ is said to be EF1 iff for every pair of agents $i, j \in [n]$ there exists a good $g \in A_j$ such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$.
- **Maximin Share Guarantee (MMS):** Given an instance $\langle [m], [n], (v_i)_{i \in [n]}, \mathcal{F} \rangle$, the (constrained) maximin share

⁶Note that the set \mathcal{F} might be exponential in size, but it is specified in an efficient manner via the partition $\{C_1, C_2, \dots, C_\ell\}$ and thresholds k_h . Setting $\ell=1$, $C_1=[m]$, $k_1=m$, we get $\mathcal{F}=\Pi_n([m])$. Hence, this formulation is a strict generalization of the unconstrained fair division problem.

of agent i is defined as

$$\text{CMMS}_i := \max_{(P_1, \dots, P_n) \in \mathcal{F}} \min_{j \in [n]} v_i(P_j). \quad (1)$$

An allocation $\mathcal{A} = (A_1, \dots, A_n) \in \mathcal{F}$ is said satisfy MMS iff for $i \in [n]$, we have $v_i(A_i) \geq \text{CMMS}_i$. Since MMS allocations are not guaranteed to exist, the objective is to find feasible allocations $(A_1, \dots, A_n) \in \mathcal{F}$ wherein each agent gets a bundle of value at least $\alpha \in (0, 1]$ times CMMS_i ; with factor α being as large as possible. We call such allocations α -MMS.

3 Main Results

The key results established in this paper are:

Theorem 1. *Given any fair division instance $\langle [m], [n], (v_i)_{i \in [n]}, \mathcal{F} \rangle$ with additive valuations and cardinality constraints ($\mathcal{F} \neq \emptyset$), there exists a polynomial time algorithm for finding a feasible EF1 allocation⁷.*

This theorem is established in Section 4 and it implies that as long as the set of feasible allocations \mathcal{F} is nonempty it admits an EF1 (i.e., a fair) allocation.

Theorem 2. *Given any fair division instance $\langle [m], [n], (v_i)_{i \in [n]}, \mathcal{F} \rangle$ with cardinality constraints ($\mathcal{F} \neq \emptyset$) and additive valuations, a $1/3$ -MMS allocation can be computed in polynomial time.*

Analogous to the EF1 case, this theorem provides an absolute, existence guarantee for approximate maximin fair allocations under cardinality constraints. A proof of this result appears in Section 5.

We also establish that when the valuations are identical, then EF1 allocations exist even under matroid constraints; see Theorem 3 in Section 6.

4 EF1 Allocations Under Cardinality Constraints: Proof of Theorem 1

In the unconstrained setting, there exist efficient algorithms for finding EF1 allocations; see, e.g., the *cycle-elimination* algorithm [Lipton *et al.*, 2004] and the round-robin method [Caragiannis *et al.*, 2016]. However, the allocations obtained by these algorithms are not guaranteed to satisfy the given cardinality constraints. We bypass this issue by developing a polynomial-time algorithm, ALG 1, for finding an allocation \mathcal{A} which is not only EF1, but also feasible, i.e., $\mathcal{A} \in \mathcal{F}$.

ALG 1 is based on an interesting modification of the round-robin algorithm: initially, ALG 1 selects an arbitrary order (permutation) over the agents $\sigma := (\sigma(1), \dots, \sigma(n))$. It then picks an unallocated category h and executes the *Greedy-Round-Robin* algorithm (ALG 2) with the n agents, $|C_h|$ goods (from category C_h), and the selected order σ .

ALG 2 follows the ordering σ in a round-robin fashion (i.e., it selects agents, one after the other, from $\sigma(1)$ to $\sigma(n)$), and

⁷One can construct examples to show that stronger fairness notions than EF1—in particular, *envy-free up to the least valued good* (EFX) [Caragiannis *et al.*, 2016] and *envy-free up to one least-preferred good* (EFL) [Barman *et al.*, 2018]—are not guaranteed to exist under cardinality constraints.

Algorithm 1 ALG 1

Input : A fair division instance $\langle [m], [n], (v_i)_i, \mathcal{F} \rangle$ with additive valuations and cardinality constraints.

Output: A feasible EF1 allocation

- 1: Initialize allocation $\mathcal{A}^0 = (A_1^0, \dots, A_n^0)$ with $A_i^0 \leftarrow \emptyset$ for each agent $i \in [n]$.
 - 2: Fix an (arbitrary) ordering of the agents $\sigma = (\sigma(1), \sigma(2), \dots, \sigma(n))$
 - 3: **for** $h = 1$ to ℓ **do**
 - 4: $\mathcal{B}^h \leftarrow$ Greedy-Round-Robin($C_h, [n], (v_i)_i, \sigma$).
 - 5: Set $A_i^h \leftarrow A_i^{h-1} \cup B_i^h$ for all $i \in [n]$.
 - 6: Using Lemma 1, update $\mathcal{A}^h = (A_1^h, \dots, A_n^h)$ to obtain an acyclic envy graph $\mathcal{G}(\mathcal{A}^h)$.
 - 7: Update σ to be a topological ordering of $\mathcal{G}(\mathcal{A}^h)$.
 - 8: **Return** \mathcal{A}^ℓ .
-

iteratively assigns to the selected agent an unallocated good from C_h that she desires the most. Finally, it returns an allocation $\mathcal{B}^h \in \Pi_n(C_h)$.

Algorithm 2 Greedy-Round-Robin (ALG 2)

Input : An instance $\langle C, [n], (v_i)_i \rangle$ with additive valuations and an ordering σ of $[n]$.

Output: An allocation of the given $|C|$ goods among n agents

- 1: Initialize bundle $B_i \leftarrow \emptyset$, for each agent $i \in [n]$, the set of unallocated goods $M \leftarrow C$, and $t \leftarrow 0$.
 - 2: **while** $M \neq \emptyset$ **do**
 - 3: $t \leftarrow t + 1$.
 - 4: **for** $i = 1$ to n **do**
 - 5: Set $g_{\sigma(i)}^t \in \arg \max_{g \in M} v_{\sigma(i)}(g)$
 - 6: Update $B_{\sigma(i)} \leftarrow B_{\sigma(i)} \cup \{g_{\sigma(i)}^t\}$
 - 7: Update $M \leftarrow M \setminus \{g_{\sigma(i)}^t\}$.
 - 8: **if** $M == \emptyset$ **then**
 - 9: **break**;
 - 10: **Return** $\mathcal{B} = (B_1, \dots, B_n)$.
-

After allocating all the goods of a category h , Step 6 of ALG 1 creates an *envy graph*⁸ $\mathcal{G}(\mathcal{A}^h)$. It was established in [Lipton *et al.*, 2004] that one can always efficiently update a given partial allocation such that the resulting envy graph is acyclic:

Lemma 1. [Lipton *et al.*, 2004] *Given a partial allocation $(A_1, \dots, A_n) \in \Pi_n(S)$ of a subset of goods S , we can find another partial allocation $\mathcal{B} = (B_1, \dots, B_n) \in \Pi_n(S)$ of S in polynomial time such that*

- (i) *The valuations of the agents for their bundles do not decrease: $v_i(B_i) \geq v_i(A_i)$ for all $i \in [n]$.*
- (ii) *The envy graph $\mathcal{G}(\mathcal{B})$ is acyclic.*

Finally, σ is updated to be a topological ordering of the

⁸An envy graph, for an allocation \mathcal{A} , is a directed graph that captures the envy between agents in \mathcal{A} . Specifically, the nodes in the envy graph represent the agents and it contains a directed edge from i to j iff i envies j , i.e., iff $v_i(A_i) < v_i(A_j)$.

acyclic directed graph $\mathcal{G}(\mathcal{A}^h)$. This new ordering is then used for the next category of goods.

The feasibility of the computed allocation, \mathcal{A}^ℓ , directly follows from the fact that (for each $h \in [\ell]$) ALG 2 distributes the $|C_h|$ goods evenly among n agents. In particular, the round-robin nature of ALG 2 ensures that, B_i^h , the set of goods allocated to agent i from category h satisfies:⁹

$$|B_i^h| \leq \left\lceil \frac{|C_h|}{n} \right\rceil \leq k_h.$$

In Lemma 2 we show that, for each $h \in [\ell]$, the partial allocation obtained after the allocating the first h categories, \mathcal{A}^h , is EF1. Since Algorithm 1 allocates all the m goods, the final allocation, \mathcal{A}^ℓ , is EF1 as well. This, along with the observation that ALG 1 runs in polynomial time completes the proof of Theorem 1.

Next, we establish a proposition which will be used in the proof of Lemma 2.

Proposition 1. *Given any fair division instance $\langle C, [n], (v_i)_i \rangle$ with additive valuations and an ordering of agents σ , the allocation $\mathcal{B} = (B_1, \dots, B_n)$ obtained by ALG 2 satisfies the following properties:*

- (1) *For any two indices $i < j$, the agent $\sigma(i)$ does not envy agent $\sigma(j)$, i.e., $v_{\sigma(i)}(B_{\sigma(i)}) \geq v_{\sigma(i)}(B_{\sigma(j)})$.*
- (2) *\mathcal{B} is EF1.*

Proof. Write $T := \lceil \frac{|C|}{n} \rceil$ to denote the total number of rounds of ALG 2. In each round, for $i < j$, agent $\sigma(i)$ gets to choose her most desired good among the unallocated goods before agent $\sigma(j)$. Hence, if $g_{\sigma(i)}^t$ and $g_{\sigma(j)}^t$ denote the good assigned to agent $\sigma(i)$ and $\sigma(j)$, respectively, in the t th round, then $v_{\sigma(i)}(g_{\sigma(i)}^t) \geq v_{\sigma(i)}(g_{\sigma(j)}^t)$ for all $t \in \{1, \dots, T\}$. Since the valuations are additive, the stated property holds: $v_{\sigma(i)}(B_{\sigma(i)}) \geq v_{\sigma(i)}(B_{\sigma(j)})$.

It is known that the round-robin algorithm results in an EF1 allocation [Caragiannis *et al.*, 2016], we repeat the argument for completeness: If index i is less than j , then agent $\sigma(i)$ does not envy agent $\sigma(j)$. On the other hand, even if $i > j$ the good allocated to agent $\sigma(i)$ in the t th round is of value (under $v_{\sigma(i)}$) no less than the good allocated to agent $\sigma(j)$ in the $(t + 1)$ th round: $v_{\sigma(i)}(g_{\sigma(i)}^t) \geq v_{\sigma(i)}(g_{\sigma(j)}^{t+1})$ for all $t \in \{1, \dots, T - 1\}$. Summing we get, $\sum_{t=1}^{T-1} v_{\sigma(i)}(g_{\sigma(i)}^t) \geq v_{\sigma(i)}(B_{\sigma(j)}) - v_{\sigma(i)}(g_{\sigma(j)}^1)$. Thus, the allocation \mathcal{B} is EF1. \square

Lemma 2. ALG 1 returns an EF1 allocation.

Proof. We will show inductively that, for each $h \in [\ell]$, the partial allocation obtained after allocating the first h categories, \mathcal{A}^h , is EF1. Hence, the returned allocation, \mathcal{A}^ℓ is EF1 as well.

The base case ($h = 1$) follows from Proposition 1; since $\mathcal{A}^1 = \mathcal{B}^1$, the proposition ensures that \mathcal{B}^1 is EF1.

By the induction hypothesis we have that \mathcal{A}^h is EF1 and, by construction, the corresponding envy graph $\mathcal{G}(\mathcal{A}^h)$ is

⁹Recall that the underlying feasible set of allocations \mathcal{F} is nonempty if and only if the integer limit $k_h \geq \frac{|C_h|}{n}$.

acyclic. Next, we will show that this continues to hold for the next category $h + 1$. Note that the ordering of agents for (executing ALG 2 over) the category C_{h+1} is obtained by topologically sorting $\mathcal{G}(\mathcal{A}^h)$. Let π be that topological ordering and write $\pi^{-1}(a)$ as the index of agent a according to the ordering π .

Now, if agent a envies b in \mathcal{A}^h , then there is a directed edge from a to b in $\mathcal{G}(\mathcal{A}^h)$ and, hence, $\pi^{-1}(a) < \pi^{-1}(b)$. In this case, Proposition 1 ensures that a does not envy b in \mathcal{B}^{h+1} , i.e., we have $v_a(B_a^{h+1}) \geq v_a(B_b^{h+1})$. The fact that \mathcal{A}^h is EF1 gives us $v_a(A_a^h) \geq v_a(A_b^h) - v_a(g)$, for some $g \in A_b^h$. Summing the last two inequalities and noting that $A_a^{h+1} = A_a^h \cup B_a^{h+1}$ and $A_b^{h+1} = A_b^h \cup B_b^{h+1}$ we get that \mathcal{A}^{h+1} is EF1 with respect to agent a and b .

The complementary case wherein a does not envy b in \mathcal{A}^h is analogous, since \mathcal{B}^{h+1} is guaranteed to be EF1 in itself. Overall, this establishes the stated claim, that the final allocation \mathcal{A}^ℓ is EF1. \square

Lemma 2 along with the feasibility argument mentioned above (and the direct observation that ALG 1 runs in polynomial time) proves Theorem 1.

5 MMS Under Cardinality Constraints: Proof of Theorem 2

To obtain an α -approximate maximin share allocation (α -MMS) under cardinality constraints we define a nonnegative, monotone, submodular function $F_i(\cdot)$, for each agent $i \in [n]$ which helps in reducing a fair division problem under additive valuations and cardinality constraints to an unconstrained fair division problem under monotone, submodular valuations F_i . Recall that a function $F_i(\cdot)$ is said to be monotone and submodular iff for all subsets $A \subseteq B \subseteq [m]$ and $g \in [m] \setminus B$, we have $F(A) \leq F(B)$ and $F(A \cup \{g\}) - F(A) \geq F(B \cup \{g\}) - F(B)$.

We define the function $F_i(\cdot) : 2^{[m]} \rightarrow \mathbb{R}$ as follows $F_i(S) := \sum_{h \in [\ell]} f_i^h(S)$ where,

$$f_i^h(S) := \begin{cases} \sum_{g \in S \cap C_h} v_i(g), & \text{if } |S \cap C_h| \leq k_h \\ \sum_{g \in \text{Top}_i^{k_h}(S \cap C_h)} v_i(g), & \text{otherwise} \end{cases}$$

Here, $\text{Top}_i^k(T)$ denotes the set of the k most valued (by agent i) goods contained in T . The following lemma asserts that F_i s are submodular; its proof is deferred to a full version of the paper.

Lemma 3. *For each agent $i \in [n]$ the function $F_i(\cdot)$ (defined above) is monotone, nonnegative, and submodular.*

Recall that the set of feasible allocations \mathcal{F} is nonempty iff the cardinality thresholds $k_h \geq \frac{|C_h|}{n}$ for all categories $h \in [\ell]$. Assuming nonempty set of feasible allocations, we establish the following lemma.

Lemma 4. *If for all the categories $h \in [\ell]$, the cardinality threshold satisfies $k_h \geq \frac{|C_h|}{n}$, then for each agent*

$i \in [n]$, the value of $\max_{(P_1, \dots, P_n) \in \mathcal{F}} \min_{j \in [n]} v_i(P_j)$ equals $\max_{(P_1, \dots, P_n) \in \Pi_n([m])} \min_{j \in [n]} F_i(P_j)$.

Proof. We fix an agent $i \in [n]$ and first show that the left-hand side of the stated equality is upper bounded by the right-hand side. Write $(A_1^*, \dots, A_n^*) \in \arg \max_{(P_1, \dots, P_n) \in \mathcal{F}} \min_{j \in [n]} v_i(P_j)$. Since $|A_j^* \cap C_h| \leq k_h$ for all h and j , we have $F_i(A_j^*) = v_i(A_j^*)$ for all j . Now, the inequality $\min_{j \in [n]} F_i(A_j^*) \leq \max_{(P_1, \dots, P_n) \in \Pi_n([m])} \min_{j \in [n]} F_i(P_j)$ establishes the upper bound.

We complete the proof by showing that an inequality holds in the other direction as well. Write $(B_1^*, \dots, B_n^*) \in \arg \max_{(P_1, \dots, P_n) \in \Pi_n([m])} \min_{j \in [n]} F_i(P_j)$. Say $(B_1^*, \dots, B_n^*) \notin \mathcal{F}$, then there exists an index $b \in [n]$ and $h \in [\ell]$ such that $|B_b^* \cap C_h| > k_h$. Since $k_h \geq \frac{|C_h|}{n}$, an averaging argument implies that there exists another index $a \in [n]$ for which $|B_a^* \cap C_h| < k_h$. Now, consider the lowest valued (by agent i) good $g \in B_b^* \cap C_h$. Note that $F_i(B_b^*) = F_i(B_b^* \setminus \{g\})$ and $F_i(B_a^* \cup \{g\}) \geq F_i(B_a^*)$. Hence, we can iteratively perform such swaps till all the cardinality constraints are satisfied. That is, we can obtain an allocation, say $\mathcal{B}' = (B'_1, \dots, B'_n) \in \mathcal{F}$, which is feasible and satisfies $F_i(B'_j) \geq F_i(B_j^*)$ for all $j \in [n]$. The feasibility of \mathcal{B}' ensures that the following equality holds for all j : $F_i(B'_j) = v_i(B'_j)$. Therefore, $\min_{j \in [n]} F_i(B'_j) = \min_{j \in [n]} v_i(B'_j) \leq \max_{(P_1, \dots, P_n) \in \mathcal{F}} \min_{j \in [n]} v_i(P_j)$. Hence, the left-hand side of the equality stated in the lemma is at least as much as the right-hand side. \square

Under submodular valuations a 1/3-MMS allocation can be computed efficiently in the unconstrained setting [Ghodsi *et al.*, 2017]. Therefore, for an unconstrained fair division instance over the m goods and with submodular valuations of the n agents as $(F_i)_i$, we can efficiently find an allocation $\mathcal{A} = (A_1, \dots, A_n) \in \Pi_n([m])$ which is 1/3-MMS. Employing a swap argument—similar to the one used in the proof of Lemma 4—we can efficiently convert \mathcal{A} into a feasible allocation $\mathcal{B} \in \mathcal{F}$ which satisfies $F_i(B_i) \geq F_i(A_i)$, for all $i \in [n]$. That is, for the unconstrained instance, \mathcal{B} is a 1/3-MMS allocation as well. In addition, the feasibility of \mathcal{B} implies that $F_i(B_i) = v_i(B_i)$, for all i . Overall, via Lemma 4 (i.e., the fact that the maximin shares of each agent i in the constructed unconstrained instance is equal to the underlying CMMS _{i} value), we get that \mathcal{B} is a feasible, 1/3-MMS allocation for the constrained instance. This completes the proof of Theorem 2.

6 Identical Valuations and Matroid Constraint

This section shows that if the additive valuations of the agents are identical (i.e., $v_i = v_j$ for all $i, j \in [n]$), then an EF1 allocation is guaranteed to exist even under a matroid constraint.

Matroids have been studied extensively in mathematics and computer science; see, e.g. [Oxley, 1992]. These structures provide an encompassing framework for representing combinatorial constraints; in particular, the cardinality constraints

considered in the previous sections correspond to a particular matroid, called the partition matroid. Formally, a *matroid* is defined as a pair $([m], \mathcal{I})$ where $[m]$ is the ground set of elements and \mathcal{I} —referred to as *independent sets*—is a nonempty collection of subsets of $[m]$ that satisfies: (i) Hereditary property: If $B \in \mathcal{I}$ and $A \subset B$, then $A \in \mathcal{I}$, and (ii) Independent Set Exchange: If $A, B \in \mathcal{I}$ and $|A| < |B|$, then there exist an element $x \in B \setminus A$ such that $A \cup \{x\} \in \mathcal{I}$.

We consider a fair division instance $\langle [m], [n], (v_i)_{i \in [n]}, \mathcal{M} \rangle$ where \mathcal{M} denotes the set of all allocations which satisfy the underlying matroid constraint, $\mathcal{M} := \{\mathcal{A} = (A_1, \dots, A_n) \in \Pi_n([m]) \mid A_i \in \mathcal{I} \text{ for all } i \in [n]\}$. Here, the main result is as follows

Theorem 3. *Every fair division instance $\langle [m], [n], (v_i)_{i \in [n]}, \mathcal{M} \rangle$ under additive, identical valuations and matroid constraint, $\mathcal{M} \neq \emptyset$, admits an EF1 allocation.*

Proof. We establish the result by showing that an allocation which maximizes the *Nash Social Welfare* over the set \mathcal{M} is necessarily EF1. For an allocation $\mathcal{A} = (A_1, \dots, A_n)$, the Nash Social Welfare (NSW) is defined as the geometric mean of the agents' valuations, $\text{NSW}(\mathcal{A}) := (\prod_{i \in [n]} v_i(A_i))^{1/n}$.

Let $v(\cdot)$ denote identical, additive valuations of the agents. We assume $v(g) > 0$ for all $g \in [m]$ and consider an optimal allocation $\mathcal{A} \in \arg \max_{\mathcal{B} \in \mathcal{M}} \text{NSW}(\mathcal{B})$, which satisfies

$\text{NSW}(\mathcal{A}) > 0$.¹⁰ We will prove that \mathcal{A} is EF1. Since $\mathcal{A} \in \mathcal{M}$, the stated claim follows.

Say, for contradiction, that \mathcal{A} is not an EF1 allocation, then we will show that there exists another allocation $\mathcal{A}' \in \mathcal{M}$ along with agents i and j , such that $A'_h = A_h$, for all other agents $h \in [n] \setminus \{i, j\}$, and $\min\{v(A'_i), v(A'_j)\} > \min\{v(A_i), v(A_j)\}$. The last inequality and the fact that $v(A'_i) + v(A'_j) = v(A_i) + v(A_j)$ imply $v(A'_i) \cdot v(A'_j) > v(A_i) \cdot v(A_j)$. Therefore, we get $\text{NSW}(\mathcal{A}') > \text{NSW}(\mathcal{A})$, which contradicts the optimality of \mathcal{A} .

Note that, if \mathcal{A} is not EF1 then there exists a pair of agents i, j such that

$$(E) : v(A_i) < v(A_j) - v(g) \quad \text{for all } g \in A_j.$$

Note that if there exists a good g in A_j such that $A_i \cup \{g\}$ is independent (i.e., $A_i \cup \{g\} \in \mathcal{I}$), then swapping g from A_j to A_i gives us the desired allocation \mathcal{A}' (with a strictly higher NSW than \mathcal{A}). Hence, we analyze the case in which no such good exists. In particular, we have $A_i \cup g \notin \mathcal{I}$ for all $g \in A_j$. This condition implies that $|A_i| \geq |A_j|$; otherwise, the Independent Set Exchange property of matroids would ensure the existence of a good $g \in A_j$ such that $A_i \cup \{g\} \in \mathcal{I}$.

Write $t := |A_j|$. Let \hat{A}_i denote a subset of A_i of cardinality t and hence, \hat{A}_i is also independent by the Hereditary property of independent sets of matroids. Since \hat{A}_i and A_j are independent and $|\hat{A}_i| = |A_j| = t$, there exist t component-wise distinct pairs of goods $\{(g_z^i, g_z^j) \in (\hat{A}_i, A_j) \mid z \in [t]\}$

¹⁰If the optimal NSW over \mathcal{M} is zero, then it must be the case that we have less than n goods. For such an instance, an allocation wherein each agent gets at most one good is both feasible and EF1.

such that $(\hat{A}_i \setminus \{g_z^i\}) \cup \{g_z^j\}$ and $(A_j \setminus \{g_z^j\}) \cup \{g_z^i\}$ are independent for all $z \in [t]$ [Goemans, 2009]. Since the pairs are distinct, $A_j = \{g_z^j \mid z \in [t]\}$ and $\hat{A}_i = \{g_z^i \mid z \in [t]\}$. In addition, the envy between agent i and j (i.e., $v(A_j) > v(A_i) \geq v(\hat{A}_i)$) implies that there exists index $y \in [t]$ for which $v(g_y^j) > v(g_y^i)$.

Consider the allocation \mathcal{A}' wherein $A'_i = (A_i \setminus \{g_y^i\}) \cup \{g_y^j\}$, $A'_j = (A_j \setminus \{g_y^j\}) \cup \{g_y^i\}$, and $A'_h = A_h$ for all other agents. Note that $v(A'_i) = v(A_i) + v(g_y^j) - v(g_y^i) > v(A_i)$. Furthermore, $v(A'_j) = v(A_j) - v(g_y^j) + v(g_y^i) > v(A_j)$; the last inequality follows from (E). Hence, $\min\{v(A'_i), v(A'_j)\} > v(A_i) = \min\{v(A_i), v(A_j)\}$, which implies that the NSW of \mathcal{A}' is higher than that of \mathcal{A} . This, by contradiction, completes the proof. \square

7 Conclusion and Future Work

This paper extends the active line of work on fair division of indivisible goods and shows that fairness guarantees are not lost by imposing cardinality constraints. In particular, we show that EF1 allocations are guaranteed to exist even under cardinality constraints. Note that, though the round-robin method [Caragiannis *et al.*, 2016] and cycle-elimination algorithm [Lipton *et al.*, 2004] efficiently find EF1 allocations in the unconstrained setting, these algorithms can lead to an allocation which does not satisfy the cardinality constraints. In this paper we bypass this issue by combining the round-robin method with envy graphs in an interesting manner.

The universality of EF1 is further strengthened by our result which shows that if the agents' valuations are identical (and additive), then fair (EF1) allocations exist even under matroid constraints. Establishing such a guarantee—along with computational results—for heterogeneous valuations (additive and beyond) remains an interesting direction of future work.¹¹

We also show that, even under cardinality constraints, approximate maximin fair allocations always exist and can be computed efficiently. This result is obtained by reducing the constrained version of the problem (with additive valuations) to an unconstrained one (with submodular valuations). This reduction might be useful for addressing other constrained, fair division problems.

Going forward it would be quite relevant to study fair division, both unconstrained and constrained, among strategic agents.

Acknowledgments

Siddharth Barman was supported by a Ramanujan Fellowship (SERB - SB/S2/RJN-128/2015). Arpita Biswas gratefully acknowledges the support of a Google PhD Fellowship Award.

¹¹In and of itself, the approach used in Section 6 does not establish such a guarantee: one can construct instances where the agents have additive (but different) valuations and the allocation that maximizes the Nash social welfare, subject to cardinality constraints, is not EF1.

References

- [Amanatidis *et al.*, 2015] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. Approximation Algorithms for Computing Maximin Share Allocations. In *42nd International Colloquium on Automata, Languages, and Programming, ICALP*, pages 39–51, 2015.
- [Barman and Krishnamurthy, 2017] Siddharth Barman and Sanath Kumar Krishnamurthy. Approximation Algorithms for Maximin Fair Division. In *ACM Conference on Economics and Computation, EC*, pages 647–664, 2017.
- [Barman *et al.*, 2018] Siddharth Barman, Arpita Biswas, Sanath Kumar Krishnamurthy, and Y. Narahari. Groupwise Maximin Fair Allocation of Indivisible Goods. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, 2018.
- [Bouveret and Lemaître, 2014] Sylvain Bouveret and Michel Lemaître. Characterizing Conflicts in Fair Division of Indivisible Goods using a Scale of Criteria. In *International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1321–1328, 2014.
- [Bouveret *et al.*, 2017] Sylvain Bouveret, Katarína Cechlárová, Edith Elkind, Ayumi Igarashi, and Dominik Peters. Fair Division of a Graph. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI*, pages 135–141, 2017.
- [Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [Budish *et al.*, 2016] 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, 2016.
- [Budish, 2011] Eric Budish. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2016] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The Unreasonable Fairness of Maximum Nash Welfare. In *ACM Conference on Economics and Computation, EC*, pages 305–322, 2016.
- [Endriss, 2017] Ulle Endriss. *Trends in Computational Social Choice*. Elsevier, 2017.
- [Ferraioli *et al.*, 2014] Diodato Ferraioli, Laurent Gourvès, and Jérôme Monnot. On Regular and Approximately Fair Allocations of Indivisible Goods. In *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS*, pages 997–1004, 2014.
- [Foley, 1967] DC Foley. Resource allocation in the public sector. *Yale Economic Essays*, 7:73–76, 1967.
- [Ghodsi *et al.*, 2017] Mohammad Ghodsi, MohammadTaghi HajiAghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods: Improvement and generalization. *arXiv preprint arXiv:1704.00222*, 2017.
- [Goemans, 2009] M Goemans. Lecture notes on Matroid Intersection (Lecture 11). *Massachusetts Institute of Technology, Combinatorial Optimization*, 2009.
- [Goldman and Procaccia, 2014] Jonathan R. Goldman and Ariel D. Procaccia. Spliddit: unleashing fair division algorithms. *SIGecom Exchanges*, 13(2):41–46, 2014.
- [Gourvès and Monnot, 2017] Laurent Gourvès and Jérôme Monnot. Approximate Maximin Share Allocations in Matroids. In *10th International Conference on Algorithms and Complexity, CIAC*, pages 310–321, 2017.
- [Gourvès *et al.*, 2014] Laurent Gourvès, Jérôme Monnot, and Lydia Tlili. Near fairness in matroids. In *Proceedings of the 21st European Conference on Artificial Intelligence, ECAI*, pages 393–398, 2014.
- [Kurokawa *et al.*, 2016] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. When Can the Maximin Share Guarantee Be Guaranteed? In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 523–529, 2016.
- [Lipton *et al.*, 2004] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On Approximately Fair Allocations of Indivisible Goods. In *ACM conference on Electronic Commerce, EC*, pages 125–131, 2004.
- [Moulin, 1990] Hervé Moulin. Uniform Externalities: Two axioms for fair allocation. *Journal of Public Economics*, 43(3):305–326, 1990.
- [Oxley, 1992] James G Oxley. *Matroid Theory*, 1992.
- [Procaccia and Wang, 2014] Ariel D Procaccia and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. In *ACM Conference on Economics and Computation, EC*, pages 675–692, 2014.
- [Stromquist, 1980] Walter Stromquist. How to Cut a Cake Fairly. *The American Mathematical Monthly*, 87(8):640–644, 1980.
- [Varian, 1974] Hal R Varian. Equity, Envy, and Efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.