

Fair and Efficient Allocations under Subadditive Valuations

Bhaskar Ray Chaudhury ¹, Jugal Garg ², Ruta Mehta ²

¹ Max Planck Institute for Informatics, Saarland Informatics Campus

² University of Illinois at Urbana-Champaign

braycha@mpi-inf.mpg.de, {jugal, rutameht}@illinois.edu

Abstract

We study the problem of allocating a set of indivisible goods among agents with subadditive valuations in a fair and efficient manner. Envy-Freeness up to any good (EFX) is the most compelling notion of fairness in the context of indivisible goods. Although the existence of EFX is not known beyond the simple case of two agents with subadditive valuations, some good approximations of EFX are known to exist, namely $\frac{1}{2}$ -EFX allocation and EFX allocations with bounded charity.

Nash welfare (the geometric mean of agents' valuations) is one of the most commonly used measures of efficiency. In case of additive valuations, an allocation that maximizes Nash welfare also satisfies fairness properties like Envy-Free up to one good (EF1). Although there is substantial work on approximating Nash welfare when agents have additive valuations, very little is known when agents have subadditive valuations. In this paper, we design a polynomial-time algorithm that outputs an allocation that satisfies either of the two approximations of EFX as well as achieves an $\mathcal{O}(n)$ approximation to the Nash welfare. Our result also improves the current best-known approximation of $\mathcal{O}(n \log n)$ and $\mathcal{O}(m)$ to Nash welfare when agents have submodular and subadditive valuations, respectively.

Furthermore, our technique also gives an $\mathcal{O}(n)$ approximation to a family of welfare measures, p -mean of valuations for $p \in (-\infty, 1]$, thereby also matching asymptotically the current best approximation ratio for special cases like $p = -\infty$ while also retaining the remarkable fairness properties.

1 Introduction

Discrete fair division of resources is a fundamental problem in various multi-agent settings, where the goal is to partition a set M of m *indivisible* goods among n agents in a *fair* and *efficient* manner. Each agent i has a valuation function $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$ that quantifies the amount of utility i derives from every subset of goods. We assume that v_i 's are *monotone*, i.e., $v_i(A) \leq v_i(A \cup \{g\})$ for all $g \in M$, *normalized* i.e., $v_i(\emptyset) = 0$ and *subadditive*, i.e., $v_i(A \cup B) \leq v_i(A) + v_i(B)$, for all $A, B \subseteq M$. Subadditive functions naturally arise in practice because they capture the notion of *complement-freeness* (Lehmann, Lehmann, and

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Nisan 2006). Furthermore, they strictly contain *submodular* functions¹, which capture the notion of diminishing marginal returns.

Among various choices, *envy-freeness* is the most natural fairness concept, where no agent i envies another agent j 's bundle, i.e., partition of goods into n bundles X_1, X_2, \dots, X_n so that for all agents i and j , we have $v_i(X_i) \geq v_i(X_j)$. However, envy-free allocation do not always exist, e.g., consider allocating a single valuable good among two agents. Its mild relaxation envy-freeness up to *any* good (EFX) (Caragiannis et al. 2016) is arguably the most compelling notion of fairness in discrete setting, where no agent envies other's allocation after the removal of *any* good, i.e., for all agents i and j , we have $v_i(X_i) \geq v_i(X_j \setminus \{g\})$ for all $g \in X_j$. While it is not known whether an EFX allocation always exists or not beyond the simple case of two agents under subadditive valuations, the following relaxations exist:

- $\frac{1}{2}$ -EFX allocation $X = \langle X_1, X_2, \dots, X_n \rangle$ where $v_i(X_i) \geq \frac{1}{2} \cdot v_i(X_j \setminus \{g\})$, for all $g \in X_j$ (Plaut and Roughgarden 2018). In this paper we will be referring to a relaxed version of $\frac{1}{2}$ -EFX namely, $(\frac{1}{2} - \varepsilon)$ -EFX allocation where $v_i(X_i) \geq (\frac{1}{2} - \varepsilon) \cdot v_i(X_j \setminus \{g\})$ for all $g \in X_j$. A $(\frac{1}{2} - \varepsilon)$ -EFX allocation can also be computed in polynomial time when agents have subadditive valuations.
- EFX allocation with *bounded charity* $X = \langle X_1, X_2, \dots, X_n \rangle$ where we do not allocate a set P of goods (set P is donated to charity) where $|P| < n$ and $v_i(X_i) \geq v_i(P)$, for all $i \in [n]$ and the partial allocation X is EFX (Chaudhury et al. 2020). There is also a polynomial time algorithm to find an $(1 - \varepsilon)$ -EFX allocation with bounded charity for general valuations for any $\varepsilon > 0$ (Chaudhury et al. 2019)².

Another popular (and stronger) relaxation is envy-freeness up to *one* good (EF1) (Budish 2011), where no agent envies other's allocation after the removal of *some* good from the other's bundle, i.e., $v_i(X_i) \geq v_i(X_j \setminus \{g\})$, for *some* $g \in X_j$. Clearly, EFX implies EF1. Although

¹A function $v(\cdot)$ is submodular if $v(A) + v(B) \geq v(A \cup B) + v(A \cap B)$, $\forall A, B \subseteq M$.

²This is an updated version of the paper which goes beyond the preliminary version published in SODA 2020

the existence of EFX allocations still remains a major open question, an EF1 allocation always exists for general valuations and can be obtained in polynomial time (Lipton et al. 2004).

We note that none of the above algorithms provides, as such, any efficiency guarantees. For efficiency, among many choices, maximum Nash welfare, defined as the geometric mean of agents' valuations, serves as a focal point. In contrast to other popular welfare measures such as social welfare and max-min welfare, Nash welfare is scale invariant, i.e., scaling one agent's valuation by any positive constant does not change the outcome. In case of additive valuations³, an allocation that maximizes Nash welfare is both EF1 and Pareto optimal⁴ (Caragiannis et al. 2016). However, such an allocation does not provide the EF1 property beyond additive (e.g., subadditive valuations (Caragiannis et al. 2016)), and further, no meaningful guarantee in terms of EFX even in the case of additive valuations (Amanatidis et al. 2020). Furthermore, maximizing the Nash welfare is a hard problem, and the best known approximation guarantees are $\mathcal{O}(n \log n)$ and $\mathcal{O}(m)$ for submodular (Garg, Kulkarni, and Kulkarni 2020) and subadditive (Nguyen and Rothe 2014) valuations, respectively. As is the case with the algorithms providing fairness guarantees, these Nash welfare approximation algorithms do not provide any fairness guarantees. Therefore, a natural question is:

Does there exist a polynomial-time algorithm that provides the best known fairness guarantees as well as the best known efficiency guarantees simultaneously?

In this paper, we answer this question affirmatively. We design a simple algorithm that outputs an allocation that provides (i) either of the best-known EFX approximations mentioned above, (ii) EF1 guarantee, and (iii) $\mathcal{O}(n)$ approximation to the maximum Nash welfare. The latter also improves the best-known approximation factor. Further, we show that our algorithm can be easily adapted to obtain the same guarantees for the entire family of p -mean welfare measures $M_p(X)$, defined as,

$$M_p(X) = \left(\sum_i \frac{1}{n} (v_i(X_i))^p \right)^{1/p} \quad \text{for } p \in (-\infty, 1].$$

The $p = -\infty, 0$, and 1 correspond to the well-studied cases of max-min welfare, Nash welfare, and social welfare, respectively. We note that this also matches the current best approximation ratio for the max-min welfare (Khot and Ponnuswami 2007) while also retaining the above mentioned fairness guarantees.

One crucial difference between Nash welfare and p -mean welfare when $p \neq 0$ is that p -mean is no longer scale invariant. Therefore, it is not intuitive that the allocation that maximizes welfare will be fair when our fairness measures are

³A valuation function $v(\cdot)$ is additive if $v_i(S) = \sum_{j \in S} v_i(\{j\}), \forall S$.

⁴An allocation $X' = (X'_1, \dots, X'_n)$ Pareto dominates another allocation $X = (X_1, \dots, X_n)$ if $v_i(X'_i) \geq v_i(X_i), \forall i$ and $v_k(X'_k) > v_k(X_k)$ for some k . An allocation X is Pareto optimal if no allocation X' dominates X .

relaxations of envy-freeness.⁵ However, we manage to give a polynomial time algorithm that achieves a good approximation (independent of the number of goods in the instance) to the p -mean welfare while still retaining all the fairness properties.

1.1 Technical Overview

In this section, we briefly sketch our main result and overall approach. We first state the main result of our paper

Theorem 1. *Given a discrete fair division instance with a set $[n]$ of n agents, a set M of m indivisible goods, where each agent i has a subadditive valuation function $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$, for any $\varepsilon > 0$, we can find in polynomial time*

- a partition $\langle X_1, X_2, \dots, X_n \rangle$ of M such that X is $(\frac{1}{2} - \varepsilon)$ -EFX and $M_p(X) \geq \frac{1-2\varepsilon}{8(n+1)} \cdot M_p(X^*)$, and
- a partition $\langle X_1, X_2, \dots, X_n, P \rangle$ of M such that X is $(1 - \varepsilon)$ -EFX with bounded charity and $M_p(X) \geq \frac{1-\varepsilon}{4(n+1)} \cdot M_p(X^*)$,

where X^* is the allocation with maximum p -mean value.

We now briefly sketch our main techniques: Let us consider the scenario that a given instance admits an envy-free allocation, i.e., a partition of the goods into n bundles X_1, X_2, \dots, X_n such that for all pairs of agents i and j we have $v_i(X_i) \geq v_i(X_j)$. In that case for each agent i we have

$$\begin{aligned} n \cdot v_i(X_i) &\geq \sum_{j \in [n]} v_i(X_j) \\ &\geq v_i(\cup_{j \in [n]} X_j) \quad (\text{by subadditivity}) \\ &= v_i(M) \end{aligned}$$

This implies that $v_i(X_i) \geq \frac{1}{n} \cdot v_i(M)$. Since in any optimal allocation no agent can get a valuation more than $v_i(M)$, we can conclude that each agent has a bundle worth $\frac{1}{n}$ times his bundle at optimum. This would immediately give us an n approximation for generalized p -mean welfare. However, most instances may not admit an envy-free allocation. Naturally, we then look into the closest relaxation of envy-freeness that is known to exist in the context of indivisible goods⁶. $(\frac{1}{2})$ -EFX (Plaut and Roughgarden 2018) and EFX with bounded charity⁷ (Chaudhury et al. 2020)). So let us consider the $\frac{1}{2}$ -EFX allocation: Here we can partition the given instance into n bundles X_1, X_2, \dots, X_n such that for all pairs of agents i and j we have $v_i(X_i) \geq \frac{1}{2} v_i(X_j \setminus \{g\})$ for all $g \in X_j$. Let us first look into all the bundles

⁵Consider a special case when $p = -\infty$. Here the p -mean welfare is equal to the valuation of the agent with smallest valuation. In particular, consider the scenario with two agents and n goods where agent 1 has a valuation of 1 for each good and agent 2 has a valuation of $\varepsilon < \frac{1}{n}$ for each good. The allocation that maximizes the p -mean welfare here will give exactly one good to agent 1 and $n-1$ goods to agent 2, which is very far from satisfying any relaxation of envy-freeness.

⁶In our algorithm we consider relaxed variants of these notions like $(\frac{1}{2} - \varepsilon)$ -EFX and $(1 - \varepsilon)$ -EFX with bounded charity, but for clarity in this section we keep the original notions

⁷defined in Section 2

X_j that are not singleton, i.e., $|X_j| \geq 2$: We have that $v_i(X_i) \geq \frac{1}{2} \cdot v_i(X_j \setminus \{g\})$ for all $g \in X_j$, implying that $v_i(X_i) \geq \frac{1}{2} \cdot \max(v_i(X_j \setminus \{g\}), v_i(\{g\}))$ (as $|X_j| \geq 2$). Thus,

$$\begin{aligned} n \cdot v_i(X_i) &\geq \sum_{|X_j| \geq 2} \frac{1}{2} \cdot \frac{1}{2} \cdot (v_i(X_j \setminus \{g\}) + v_i(\{g\})) \\ &\geq \frac{1}{4} \cdot \sum_{|X_j| \geq 2} v_i(X_j) \\ &\quad \text{(by subadditivity)} \\ &\geq \frac{1}{4} \cdot v_i(\bigcup_{|X_j| \geq 2} X_j) \\ &\quad \text{(by subadditivity)} \end{aligned} \tag{1}$$

Let S be the set of all the goods in singleton bundles in X , i.e., $S = \{g \mid \text{there is a } X_j = \{g\}\}$. Then from (1) we have the guarantee that for every agent $v_i(X_i) \geq \frac{1}{4n} \cdot v_i(M \setminus S)$. Therefore, in any $\frac{1}{2}$ -EFX allocation every agent has an $\frac{1}{4n}$ fraction of his valuation on the goods he receives from $M \setminus S$ in the optimal allocation, i.e., $v_i(X_i^* \cap (M \setminus S))$ where $X^* = (X_1^*, \dots, X_n^*)$ is the allocation that has the highest generalized p -mean welfare. The only problem is how to allocate the goods in the set S appropriately.

The only scenario where an incorrect allocation of the goods in S causes a significant decrease in the p -mean welfare is when there are agents who have a substantially high valuation for some goods in S . However, we could very well be in a scenario where there are only a few goods in S (say less than $\frac{n}{3}$) which are very valuable to many agents and then we may not be able to give every agent a bundle that he values $\frac{1}{n}$ times the whole set S .⁸ Therefore we need to compare our allocation with the allocation that maximizes the p -mean welfare.

We briefly sketch how we overcome this barrier. The good aspect of the situation is that the number of goods in S is small, i.e., $|S| \leq n$. Let H_i denote the set of n goods that are valued by agent i the most, i.e., all goods in H_i are more valuable than any good outside H_i . Now we find a *single good allocation* (where each agent gets exactly one good) of the *high valued goods*, namely the set $\mathbf{H} = \bigcup_{i \in [n]} H_i$, optimally to the agents assuming that we can give each agent at least $\frac{1}{n}$ times their valuation for the *low valued goods*, namely the set $M \setminus H_i$. That is, we find a single good allocation, where every agent i gets exactly one high valued good $h_i \in H_i$, that maximizes $\sum_{i \in [n]} (v_i(\{h_i\}) + \frac{1}{n} v_i(M \setminus H_i))^p$ (such allocations can be found efficiently with matching algorithms). Let us call the current single good allocation Y . Note that Y is trivially EFX as every agent has exactly one good (therefore Y is also $\frac{1}{2}$ -EFX). We then run the $\frac{1}{2}$ -EFX

⁸A very simple scenario is to divide n goods among n agents with identical additive valuations, where all agents have a valuation of 1 for a single good and $\varepsilon \ll \frac{1}{n}$ for the rest of the goods. In any division there will be $n - 1$ agents who do not get $\frac{1}{n}$ of their valuation on the set of n goods

algorithm starting with Y as the initial partial $\frac{1}{2}$ -EFX allocation. The intuition being that the low valued goods appear in non-singleton bundles and the high valued goods occur in singleton bundles in the final $\frac{1}{2}$ -EFX allocation, but we have allocated the high valued goods correctly (up to a factor of $\frac{1}{n}$ as we computed a single good allocation, while the optimum need not necessarily give every agent exactly one high valued good) as we started out with an optimal allocation of the high valued goods. Since the low valued goods occur in non-singleton bundles we are indeed able to give every agent $\frac{1}{n}$ times their valuation for the low valued goods.

1.2 Related Work

Fair division has been extensively studied for more than seventy years since the seminal work of Steinhaus (Steinhaus 1948). A complete survey of all different settings and the fairness and efficiency notions used is well beyond the scope of this paper. We limit our attention to fair division of indivisible goods (as mentioned in Section 1). There have been extensive studies on relaxations of envy-freeness like EF1 (Budish et al. 2017; Barman et al. 2018; Lipton et al. 2004; Caragiannis et al. 2016) and EFX (Chaudhury et al. 2020; Caragiannis, Gravin, and Huang 2019; Plaut and Roughgarden 2018) and relaxations of proportionality like maximin shares (MMS) (Budish 2011; Bouveret and Lemaître 2016; Amanatidis et al. 2017; Barman and Krishnamurthy 2017; Kurokawa, Procaccia, and Wang 2018; Ghodsi et al. 2018; Garg, McGlaughlin, and Taki 2019; Garg and Taki 2019) and proportionality up to one good (PROP1) (Conitzer, Freeman, and Shah 2017; Barman and Krishnamurthy 2019; Garg and McGlaughlin 2019).

While there is a significant interest in finding fair allocations, there has also been a lot of interest in guaranteeing *efficient* fair allocations. A common measure of economic efficiency is *Pareto-Optimality*⁹. (Caragiannis et al. 2016) showed that when agents have additive valuations, any allocation that maximizes the Nash welfare is guaranteed to be Pareto-optimal (efficient) and EF1 (fair). Hence the Nash welfare is also a good measure of efficiency of fair allocations. Unfortunately, finding an allocation with the maximum Nash welfare is APX-hard (Lee 2017). However, approximation algorithms for Nash welfare under different types of valuations have received significant attention recently, e.g., (Cole and Gkatzelis 2018; Cole et al. 2017; Anari et al. 2017; Garg, Hoefer, and Mehlhorn 2018; Anari et al. 2018; Barman, Krishnamurthy, and Vaish 2018; Chaudhury et al. 2018; Garg, Kulkarni, and Kulkarni 2020; Babaioff, Ezra, and Feige 2020; Benabbou et al. 2020).

1.3 Independent Work

Independently of our work, (Barman et al. 2020) also find an $\mathcal{O}(n)$ -approximation algorithm for the generalized p -mean welfare when agents have subadditive valuations. On a high level, both algorithms, first carefully allocate a single *highly valuable* good to each agent and then carefully allocate the remaining goods. However, the procedures used to

⁹Defined in Section 1

determine the initial (the single highly valuable good allocation) and the final allocations are significantly different. Also, contrary to the allocation determined by the algorithm in (Barman et al. 2020), we are able to obtain guarantees on the fairness of our allocation, namely the properties of EF1 and the two relaxations of EFX.

In the same paper, (Barman et al. 2020) show that it requires an exponential number of value queries to provide any sublinear approximation for the generalized p -mean welfare under subadditive valuations. Therefore, in polynomial time, our algorithm yields an allocation that satisfies the best relaxations of EFX known for subadditive valuations, and also achieves the best approximation for the generalized p -mean welfare possible in polynomial time.

2 Preliminaries

Any discrete fair division instance I is a tuple $\langle [n], M, \mathcal{V} \rangle$ comprising of a set of n agents ($[n]$), a set of m goods (M) and a set of valuation functions $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$. The valuation function $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$ tries to capture agent i 's utility for each subset of goods. Throughout this paper we will be dealing with the case where all the valuation functions are

- *normalized*: $v_i(\emptyset) = 0$ for all $i \in [n]$,
- *monotone*: $v_i(A \cup \{g\}) \geq v_i(A)$ for all $i \in [n]$ and $A \subset M$, and
- *subadditive*: for any sets $A, B \subseteq M$ we have $v_i(A) + v_i(B) \geq v_i(A \cup B)$ for all $i \in [n]$.

For ease of notation we use $v_i(g)$ instead of $v_i(\{g\})$ and $v_i(A \setminus g)$ instead of $v_i(A \setminus \{g\})$

Generalized p -mean welfare: Given an allocation X the p -mean welfare of X (parametrized by p) is defined by

$$M_p(X) = \left(\frac{1}{n} \sum_{i \in [n]} v_i(X_i)^p \right)^{\frac{1}{p}}$$

This captures a wide range of fairness and efficiency measures that have been used frequently in the literature: Nash welfare for $p = 0$, max-min welfare (also known as the egalitarian welfare) for $p = -\infty$ and social welfare for $p = 1$. Barman and Sundaram (Barman and Sundaram 2020) also mention that,

“generalized means with $p \in (-\infty, 1]$ exactly constitute the family of welfare functions that satisfy the Pigou-Dalton transfer principle and a few other key axioms.”

In the same paper they show that when agents have identical valuations, there is an algorithm that provides an $\mathcal{O}(1)$ factor approximation to the p -mean welfare.

EFX Allocations and Relaxations: EFX is arguably the strongest notion of fairness in the context of indivisible goods. Formally,

Definition 2. An allocation $X = \langle X_1, X_2, \dots, X_n \rangle$ is said to be an EFX allocation if for all pairs of agents i and j , we have $v_i(X_i) \geq v_i(X_j \setminus g)$ for all $g \in X_j$.

Although the existence of complete EFX allocations is not known yet, there have been results pertaining to the existence of certain relaxations of EFX. We state two major relaxations here. Plaut and Roughgarden (Plaut and Roughgarden 2018) introduced the notion of *approximate EFX* or equivalently α -EFX:

Definition 3. An allocation X is α -EFX with $\alpha \in (0, 1)$ if and only if for all pairs of agents i and j we have $v_i(X_i) \geq \alpha \cdot v_i(X_j \setminus g)$ for all $g \in X_j$.

(Plaut and Roughgarden 2018) showed that $\frac{1}{2}$ -EFX allocations exist and can be computed in pseudo-polynomial time. With a very minor change in their algorithm¹⁰ we can obtain an $(\frac{1}{2} - \varepsilon)$ -EFX allocation in polynomial time.

Another relaxation introduced by (Chaudhury et al. 2020) is *EFX with bounded charity*:

Definition 4. A partial allocation X is an EFX allocation with bounded charity with the set of unallocated goods P such that

- X is EFX,
- $|P| < n$, and
- $v_i(X_i) \geq v_i(P)$ for all $i \in [n]$.

The updated version of the paper (Chaudhury et al. 2019) gives a polynomial time algorithm to determine $(1 - \varepsilon)$ -EFX allocation with bounded charity.¹¹

Throughout the paper we will outline algorithms that find allocations with high welfare and are flexible with the fairness that the allocations satisfy, i.e., we can get $(\frac{1}{2} - \varepsilon)$ -EFX allocations with high welfare or $(1 - \varepsilon)$ -EFX allocations with bounded charity and high welfare. Therefore we now introduce some common notation for ease of referring to both these relaxations of EFX at the same time:

Definition 5. An allocation X is an (α, c) -EFX allocation with $\alpha \in (0, 1)$ and $c \in \{0, 1\}$ if and only if

- X is α -EFX and EF1,
- $|P| < n$, and
- $v_i(X_i) \geq v_i(P)$ for all $i \in [n]$.
- $P = \emptyset$ if $c = 1$.¹²

Therefore an $(\alpha, 1)$ -EFX allocation would refer to an α -EFX (which is also EF1) allocation and a $(\alpha, 0)$ -EFX allocation would refer to an α -EFX allocation with bounded charity. In particular we would only be interested in $((\frac{1}{2} - \varepsilon), 1)$ -EFX allocation and $(1 - \varepsilon, 0)$ -EFX allocation.

Similarly, we also introduce the notion of an (α, c) -EFX algorithm:

¹⁰Just run the same algorithm replacing the violation condition from $\frac{1}{2}$ -EFX to $(\frac{1}{2} - \varepsilon)$ -EFX

¹¹Just relax the first condition in Definition 4 to “ X is $(1 - \varepsilon)$ -EFX”

¹² c serves as an indicator to whether the allocation is complete or not.

Definition 6. An (α, c) -EFX algorithm takes as input any partial α -EFX allocation X and outputs an (α, c) -EFX allocation Y as the final allocation such that

- the valuation of every agent in the final allocation is at least as high as his valuation in the initial allocation, i.e., $v_i(Y_i) \geq v_i(X_i)$, and
- if there exists an agent i such that $|Y_i| = 1$ and $Y_i \neq X_i$, then $v_i(Y_i) > v_i(X_i)$.

In particular an (α, c) -EFX algorithm outputs a final (α, c) -EFX allocation that preserves (if not improves) all the welfare guarantees of the initial α -EFX allocation. Both the existing algorithms for determining an $(\frac{1}{2} - \varepsilon, 1)$ -EFX allocation (a trivial modification of the algorithm in (Plaut and Roughgarden 2018)) and $(1 - \varepsilon, 0)$ -EFX (Chaudhury et al. 2019) allocation are an $(\frac{1}{2} - \varepsilon, 1)$ -EFX algorithm and an $(1 - \varepsilon, 0)$ -EFX algorithm respectively.

3 Algorithm

In this section, we show that we can determine an (α, c) -EFX allocation with an $\mathcal{O}(n)$ approximation on the p -mean welfare. The algorithm is very simple and it has just two phases: In the first phase we try to determine a reasonable allocation of high valued goods (we call this allocation Y) and then in the second phase we just run an (α, c) -EFX algorithm with the remaining set of goods (we call our final allocation Z).

Allocating the high valued goods Y : We first formally define the notion of high valued goods for an agent: For each agent i let us order the goods in M as $\{g_1^i, g_2^i, \dots, g_m^i\}$ such that $v_i(g_1^i) \geq v_i(g_2^i) \geq \dots \geq v_i(g_m^i)$. Let $H_i = \{g_1^i, g_2^i, \dots, g_n^i\}$. We refer to H_i to be the set of high valued goods for agent i . Also for each good g_k^i and an agent i we define $rank_i(g_k^i) = k$. Notice that if for any agent i $rank_i(g) < rank_i(g')$, then $v_i(g) \geq v_i(g')$.

We now outline how we compute the initial allocation Y . We construct the complete bipartite graph $G = ([n] \cup M, [n] \times M)$ with the weight of the edge from agent i to good g , w_{ig} being

- $n \cdot v_i(g) + v_i(M \setminus H_i)$ if $p = -\infty$,
- $\log(n \cdot v_i(g) + v_i(M \setminus H_i))$ if $p = 0$ and
- $(n \cdot v_i(g) + v_i(M \setminus H_i))^p$ otherwise.

Depending on the value of p we choose an appropriate matching mechanism to determine Y : Y is determined such that $\cup_{i \in [n]}(i, Y_i)$ is

- a maximum weight matching in G if $p \geq 0$,
- a minimum weight perfect matching in G if $p < 0$ and $p \neq -\infty$,
- a max-min matching¹³ in G if $p = -\infty$.

¹³This is a matching that maximizes the weight of the smallest edge in the matching.

Let Y be the allocation outputted by the corresponding matching subroutine. Also let $\mathbf{Y} = \cup_{i \in [n]} Y_i$. We modify the allocation Y slightly such that $\cup_{i \in [n]}(i, Y_i)$ still remains the optimum matching, but no agent will prefer a good outside \mathbf{Y} to the good allocated to him in Y (Y_i), i.e., we wish to determine an allocation Y such that for all agent $i \in [n]$ and all $g \notin \mathbf{Y}$ we have that $rank_i(Y_i) < rank_i(g)$. To achieve this, as long as there is an agent $i \in [n]$ and a good $g \notin \mathbf{Y}$ such that $rank_i(g) < rank_i(Y_i)$ we set $Y_i \leftarrow \{g\}$. Note that such an operation does not decrease the optimum value of the matching: $v_i(g) \geq v_i(Y_i)$ (as $rank_i(g) < rank_i(Y_i)$) and hence $w_{ig} \geq w_{iY_i}$ for $p = -\infty$ and $p \in [0, 1]$, while $w_{ig} \leq w_{iY_i}$ for $p < 0$ and $p \neq -\infty$. This implies that the objective value of the matching does not decrease when $p \in [0, 1]$ and $p = -\infty$ and the objective value of the matching does not increase when $p < 0$ and $p \neq -\infty$. Therefore, $\cup_{i \in [n]}(i, Y_i)$ still stays an optimum matching, but $\sum_{i \in [n]} rank_i(Y_i)$ strictly decreases. Since $n \leq \sum_{i \in [n]} rank_i(Y_i) \leq nm$, after $\mathcal{O}(nm)$ iterations we will have an allocation Y such that $\cup_{i \in [n]}(i, Y_i)$ is still an optimum matching, but for all agents $i \in [n]$ and for all goods $g \notin \mathbf{Y}$ we have $rank_i(Y_i) < rank_i(g)$.

The complete algorithm is outlined in Algorithm 1 (Selection of the allocation Y is captured in steps 1 to 5).

Lemma 7. For all $i \in [n]$ we have $Y_i \subset H_i$. Furthermore, for all

- $g \notin H_i$, and

- $g \notin \mathbf{Y}$,

we have $v_i(Y_i) \geq v_i(g)$.

Proof. We first show that $Y_i \subset H_i$. We prove the same by contradiction. Assume otherwise, i.e., $Y_i \not\subset H_i$. In that case note that there is always a good $g \in H_i \setminus \mathbf{Y}$ (as $|H_i| = |\mathbf{Y}| = n$ and there is a good in \mathbf{Y} (namely Y_i) which is not in H_i). By the definition of H_i it is clear that $rank_i(g) < rank_i(g')$ for all $g' \notin H_i$. Thus we have $rank_i(g) < rank_i(Y_i)$ when $g \notin \mathbf{Y}$, which is a contradiction. Therefore $Y_i \subset H_i$. This also immediately shows that for all $g \notin H_i$ we have $v_i(g) \leq v_i(Y_i)$ (as $Y_i \subset H_i$ and any good in H_i is at least as valuable as any good outside H_i to agent i).

The proof of the last statement of the lemma is immediate. We have that $rank_i(Y_i) < rank_i(g)$ for all $g \notin \mathbf{Y}$, immediately implying that $v_i(Y_i) \geq v_i(g)$. \square

Run (α, c) -EFX algorithm on the remaining goods: Once we determined the initial allocation Y , we run an (α, c) -EFX algorithm on the remaining goods starting with Y as the initial allocation (Y is a feasible initial allocation as it is trivially an α -EFX allocation as every agent has exactly a single good). Let Z be the final (α, c) -EFX allocation. As mentioned earlier in Section 1.1 the singleton sets allocated to the agents are the barriers to proving our desired approximation for any (α, c) -EFX allocation. However since we started with a good initial allocation (namely Y), we first show that we have some nice properties about the singleton sets in the final allocation Z .

Observation 8. If $|Z_\ell| = 1$ for any $\ell \in [n]$, then we have $Z_\ell \subseteq \mathbf{Y}$.

Proof. Since Z is obtained by running an (α, c) -EFX algorithm starting with Y as the initial allocation, we have for every agent i that $v_i(Z_i) \geq v_i(Y_i)$ (from the definition of (α, c) -EFX algorithm). Note that if for any agent i we have $Z_i \neq Y_i$, and $|Z_i| = 1$, then $v_i(Z_i) > v_i(Y_i)$ (from the definition of (α, c) -EFX algorithm). Now consider the agent ℓ such that $|Z_\ell| = 1$. If $Z_\ell = Y_\ell$, then we immediately have $Z_\ell \subseteq \mathbf{Y}$. So now consider the case when $Z_\ell \neq Y_\ell$. Then we have $v_\ell(Z_\ell) > v_\ell(Y_\ell)$. By Lemma 7 we know that no good outside \mathbf{Y} can be more valuable to agent ℓ than Y_ℓ . Therefore $Z_\ell \subseteq \mathbf{Y}$. \square

Now we show a lower bound on the final valuation of an agent in terms of the low valued goods.

Observation 9. We have $v_i(Z_i) \geq \frac{\alpha v_i(M \setminus \mathbf{Y})}{2(n+1)}$ for all $i \in [n]$.

Proof. Fix an agent i . Now consider any agent j such that Z_j is not a singleton. Since Z is an α -EFX allocation, we have that $v_i(Z_i) \geq \alpha \cdot v_i(Z_j \setminus g)$ for all $g \in Z_j$. Since $|Z_j| \geq 2$ we can say that $v_i(Z_i) \geq \alpha \cdot \max(v_i(Z_j \setminus g), v_i(g))$. Therefore we have,

$$\begin{aligned} v_i(Z_i) &\geq \frac{\alpha \cdot (v_i(Z_j \setminus g) + v_i(g))}{2} \\ &\geq \frac{\alpha \cdot v_i(Z_j)}{2} \quad (\text{by subadditivity}) \end{aligned}$$

Let $S = \bigcup_{|Z_\ell|=1} Z_\ell$. By Observation 8 we know that $S \subseteq \mathbf{Y}$. Note that even if Z is a partial allocation (if $c = 0$ in the (α, c) -EFX allocation Z) and there exists a set of goods P unallocated, we have $v_i(Z_i) \geq v_i(P)$ (since Z is an (α, c) -EFX allocation). Therefore we have,

$$\begin{aligned} (n+1-|S|)v_i(Z_i) &\geq \frac{\alpha}{2} \sum_{|Z_j| \geq 2} v_i(Z_j) + v_i(P) \\ &\geq \frac{\alpha}{2} v_i \left(\bigcup_{|Z_j| \geq 2} Z_j \cup P \right) \\ &\quad (\text{by subadditivity}) \\ &= \frac{\alpha}{2} v_i(M \setminus S) \\ &\geq \frac{\alpha}{2} v_i(M \setminus Y) \\ &\quad (\text{since } S \subseteq \mathbf{Y}) \end{aligned}$$

Therefore we have $v_i(Z_i) \geq \frac{\alpha}{2(n+1-|S|)} v_i(M \setminus \mathbf{Y}) \geq \frac{\alpha}{2(n+1)} v_i(M \setminus \mathbf{Y})$. \square

Now we show a lower bound on $v_i(Z_i)$ in terms of the initial allocation Y and the set of low-valued goods for agent i , i.e., $M \setminus H_i$.

Lemma 10. For all $i \in [n]$ we have $v_i(Z_i) \geq \frac{\alpha}{4(n+1)} \cdot \left(n \cdot v_i(Y_i) + v_i(M \setminus H_i) \right)$.

Algorithm 1 Determining an (α, c) -EFX allocation with $\mathcal{O}(n)$ approximation on optimum p -mean.

1: **Construct** $G = ([n] \cup M, [n] \times M)$ with

$$w_{ig} = \begin{cases} n \cdot v_i(g) + v_i(M \setminus H_i) & \text{if } p = -\infty \\ \log(n \cdot v_i(g) + v_i(M \setminus H_i)) & \text{if } p = 0 \\ (n \cdot v_i(g) + v_i(M \setminus H_i))^p & \text{otherwise} \end{cases}$$

2: **Set** Y such that

$$\cup_{i \in [n]} (i, Y_i) = \begin{cases} \mathbf{Max-Min-Matching}(G) & \text{if } p = -\infty \\ \mathbf{Min-Weight-Perfect-Matching}(G) & \text{if } p < 0 \text{ and } p \neq -\infty \\ \mathbf{Max-Weight-Matching}(G) & \text{otherwise} \end{cases}$$

- 3: **while** $\exists i \in [n]$ and $\exists g \notin \mathbf{Y}$ such that $\text{rank}_i(g) < \text{rank}_i(Y_i)$ **do**
- 4: $Y_i \leftarrow \{g\}$.
- 5: **Set** $Z \leftarrow (\alpha, c)$ -**EFX** $\left(\langle Y_1, \dots, Y_n \rangle, (M \setminus \cup_{i \in [n]} Y_i) \right)$

Proof. Proof can be found in the full version of the paper (Lemma 10 in (Chaudhury, Garg, and Mehta 2020)). \square

The final allocation is the set Z which is obtained by running an (α, c) -EFX algorithm starting with Y as the initial allocation. Therefore our final allocation is an (α, c) -EFX allocation. We will now show the approximation guarantees that the algorithm achieves. The section that follows, proves that the allocation Z has a Nash-welfare that is $\frac{\alpha}{4(n+1)}$ times Nash-welfare achieved by the optimal allocation.

We now discuss the case where $p = 0$ and $M_p(X) = \left(\prod_{i \in [n]} v_i(X_i) \right)^{\frac{1}{n}}$. Let X^* be the allocation with the highest p -mean value and let g_i^* be agent i 's most valuable good in X_i^* . Like in the earlier sections we will show in this section that $M_p(Z) \geq \frac{\alpha}{4(n+1)} \cdot M_p(X^*)$. First observe that by Lemma 10, we have that for all $i \in [n]$, $v_i(Z_i) \geq \frac{\alpha}{4(n+1)} \cdot \left(n \cdot v_i(Y_i) + v_i(M \setminus H_i) \right)$. Therefore, $\left(\prod_{i \in [n]} v_i(Z_i) \right)^{\frac{1}{n}}$ is greater than or equal to,

$$\begin{aligned} &\left(\prod_{i \in [n]} \frac{\alpha}{4(n+1)} \cdot \left(n \cdot v_i(Y_i) + v_i(M \setminus H_i) \right) \right)^{\frac{1}{n}} \\ &= \frac{\alpha}{4(n+1)} \cdot \left(\prod_{i \in [n]} \left(n \cdot v_i(Y_i) + v_i(M \setminus H_i) \right) \right)^{\frac{1}{n}} \end{aligned}$$

Recall that Y was chosen such that (i, Y_i) is a maximum weight matching in the bipartite graph $G = ([n] \cup M, [n] \times M)$ where the weight of an edge from agent i to good g , $w_{ig} = \log(n \cdot v_i(g) + v_i(M \setminus H_i))$. Note that

$\cup_{i \in [n]} (i, g_i^*)$ is a feasible matching in G . Thus we have $\sum_{i \in [n]} \log(n \cdot v_i(Y_i) + v_i(M \setminus H_i)) \geq \sum_{i \in [n]} \log(n \cdot v_i(g_i^*) + v_i(M \setminus H_i))$. This implies that $\prod_{i \in [n]} (n \cdot v_i(Y_i) + v_i(M \setminus H_i)) \geq \prod_{i \in [n]} (n \cdot v_i(g_i^*) + v_i(M \setminus H_i))$. Therefore we have, $\left(\prod_{i \in [n]} v_i(Z_i) \right)^{\frac{1}{n}}$ to be greater than or equal to

$$\begin{aligned}
& \frac{\alpha}{4(n+1)} \cdot \left(\prod_{i \in [n]} (n \cdot v_i(Y_i) + v_i(M \setminus H_i)) \right)^{\frac{1}{n}} \\
& \geq \frac{\alpha}{4(n+1)} \cdot \left(\prod_{i \in [n]} (n \cdot v_i(g_i^*) + v_i(M \setminus H_i)) \right)^{\frac{1}{n}} \\
& \geq \frac{\alpha}{4(n+1)} \cdot \left(\prod_{i \in [n]} (n \cdot v_i(g_i^*) + \right. \\
& \quad \left. v_i(X_i^* \cap (M \setminus H_i))) \right)^{\frac{1}{n}} \\
& \geq \frac{\alpha}{4(n+1)} \cdot \left(\prod_{i \in [n]} (v_i(X_i^* \cap H_i) + \right. \\
& \quad \left. v_i(X_i^* \cap (M \setminus H_i))) \right)^{\frac{1}{n}} \\
& \quad (\text{as } |H_i| = n) \\
& \geq \frac{\alpha}{4(n+1)} \cdot \left(\prod_{i \in [n]} v_i(X_i^*) \right)^{\frac{1}{n}} \\
& \quad (\text{by subadditivity})
\end{aligned}$$

This shows that $M_p(Z) \geq \frac{\alpha}{4(n+1)} \cdot M_p(X^*)$ when $p = 0$. In the full version of the paper (Chaudhury, Garg, and Mehta 2020), we show that with a very similar proof, we can prove $M_p(Z) \geq \frac{\alpha}{4(n+1)} \cdot M_p(X^*)$ for all $p \in [-\infty, 1]$. The following theorem summarizes the main result of our paper.

Theorem 11. *Given any instance $\langle [n], M, \mathcal{V} \rangle$, in polynomial time we can determine an allocation Z such that*

- Z is either $(1 - \varepsilon, 0)$ -EFX allocation or $(\frac{1}{2} - \varepsilon, 1)$ -EFX allocation for any positive ε and
- $M_p(Z) \geq \frac{\alpha}{4(n+1)} M_p(X^*)$.

where X^* is the allocation with maximum p -mean welfare.

Proof. The proof can be found in the full version of the paper (Theorem 12 in (Chaudhury, Garg, and Mehta 2020)). \square

Remark: Theorem 11 also suggests that we can find a $\frac{4(n+1)}{1-\varepsilon}$ approximation to the p -mean welfare in polynomial time. Also, it can be verified that a minor variant of our approach (changing the weights of the edges of the complete bipartite graph $G([n] \cup B, [n] \times B)$ appropriately - step 1

of Algorithm 1) gives a $\mathcal{O}(n)$ approximation on weighted generalized p -mean, defined as $WM_p(X) = \left(\sum_{i \in [n]} \eta_i \cdot v_i(X_i)^p \right)^{\frac{1}{p}}$.

Acknowledgements

Jugal Garg is supported by NSF CAREER Grant CCF-1942321. Ruta Mehta is supported by NSF CAREER Grant CCF 1750436.

References

Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2020. Maximum Nash Welfare and Other Stories About EFX. *CoRR* abs/2001.09838.

Amanatidis, G.; Markakis, E.; Nikzad, A.; and Saberi, A. 2017. Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms* 13(4): 52:1–52:28.

Anari, N.; Gharan, S. O.; Saberi, A.; and Singh, M. 2017. Nash Social Welfare, Matrix Permanent, and Stable Polynomials. In *8th Innovations in Theoretical Computer Science Conference (ITCS)*, 1–12.

Anari, N.; Mai, T.; Gharan, S. O.; and Vazirani, V. V. 2018. Nash Social Welfare for Indivisible Items under Separable, Piecewise-Linear Concave Utilities. In *Proc. 29th Symp. Discrete Algorithms (SODA)*, 2274–2290.

Babaioff, M.; Ezra, T.; and Feige, U. 2020. Fair and Truthful Mechanisms for Dichotomous Valuations. *CoRR* abs/2002.10704.

Barman, S.; Bhaskar, U.; Krishna, A.; and Sundaram, R. G. 2020. Tight Approximation Algorithms for p-Mean Welfare Under Subadditive Valuations. *CoRR* abs/2005.07370.

Barman, S.; Biswas, A.; Murthy, S. K. K.; and Narahari, Y. 2018. Groupwise Maximin Fair Allocation of Indivisible Goods. In *AAAI*, 917–924. AAAI Press.

Barman, S.; and Krishnamurthy, S. K. 2017. Approximation algorithms for maximin fair division. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 647–664.

Barman, S.; and Krishnamurthy, S. K. 2019. On the Proximity of Markets with Integral Equilibria. In *Proc. 33rd Conf. Artif. Intell. (AAAI)*.

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 Sundaram, R. G. 2020. Uniform Welfare Guarantees Under Identical Subadditive Valuations. *CoRR* abs/2005.00504.

Benabbou, N.; Chakraborty, M.; Igarashi, A.; and Zick, Y. 2020. Finding Fair and Efficient Allocations When Valuations Don’t Add Up. In *SAGT*, volume 12283 of *Lecture Notes in Computer Science*, 32–46. Springer.

Bouveret, S.; and Lemaître, M. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *Autonomous Agents and Multi-Agent Systems (AAMAS) 30*, 2, 259–290.

Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6): 1061–1103.

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.

Caragiannis, I.; Gravin, N.; and Huang, X. 2019. Envy-Freeness Up to Any Item with High Nash Welfare: The Virtue of Donating Items. In *EC*, 527–545. ACM.

Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, 305–322.

Chaudhury, B. R.; Cheung, Y. K.; Garg, J.; Garg, N.; Hoefer, M.; and Mehlhorn, K. 2018. On Fair Division for Indivisible Items. In *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS*, 25:1–25:17.

Chaudhury, B. R.; Garg, J.; and Mehta, R. 2020. Fair and Efficient Allocations under Subadditive Valuations. *CoRR* abs/2005.06511.

Chaudhury, B. R.; Kavitha, T.; Mehlhorn, K.; and Sgouritsa, A. 2019. A Little Charity Guarantees Almost Envy-Freeness. *CoRR* abs/1907.04596.

Chaudhury, B. R.; Kavitha, T.; Mehlhorn, K.; and Sgouritsa, A. 2020. A Little Charity Guarantees Almost Envy-Freeness. In *Proceedings of the 31st Symposium on Discrete Algorithms (SODA)*, 2658–2672.

Cole, R.; Devanur, N.; Gkatzelis, V.; Jain, K.; Mai, T.; Vazirani, V.; and Yazdanbod, S. 2017. Convex Program Duality, Fisher Markets, and Nash Social Welfare. In *Proc. 18th Conf. Economics and Computation (EC)*.

Cole, R.; and Gkatzelis, V. 2018. Approximating the Nash Social Welfare with Indivisible Items. *SIAM J. Comput.* 47(3): 1211–1236.

Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair Public Decision Making. In *Proc. 18th Conf. Economics and Computation (EC)*, 629–646.

Garg, J.; Hoefer, M.; and Mehlhorn, K. 2018. Approximating the Nash Social Welfare with Budget-Additive Valuations. In *Proc. 29th Symp. Discrete Algorithms (SODA)*.

Garg, J.; Kulkarni, P.; and Kulkarni, R. 2020. Approximating Nash Social Welfare under Submodular Valuations through (Un)Matchings. In *SODA*.

Garg, J.; and McGlaughlin, P. 2019. Improving Nash Social Welfare Approximations. In *IJCAI*, 294–300. ijcai.org.

Garg, J.; McGlaughlin, P.; and Taki, S. 2019. Approximating Maximin Share Allocations. In *Proceedings of the 2nd Symposium on Simplicity in Algorithms (SOSA)*, volume 69, 20:1–20:11. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.

Garg, J.; and Taki, S. 2019. An Improved Approximation Algorithm for Maximin Shares. *CoRR* abs/1903.00029.

Ghodsi, M.; Hajiaghayi, M. T.; Seddighin, M.; Seddighin, S.; and Yami, H. 2018. Fair Allocation of Indivisible Goods: Improvements and Generalizations. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, 539–556.

Khot, S.; and Ponnuswami, A. K. 2007. Approximation Algorithms for the Max-Min Allocation Problem. In *APPROX-RANDOM*, volume 4627 of *Lecture Notes in Computer Science*, 204–217. Springer.

Kurokawa, D.; Procaccia, A. D.; and Wang, J. 2018. Fair enough: Guaranteeing approximate maximin shares. *Journal of ACM* 65(2): 8:1–27.

Lee, E. 2017. APX-hardness of Maximizing Nash Social Welfare with Indivisible Items. *Inf. Process. Lett.* 122: 17–20.

Lehmann, B.; Lehmann, D.; and Nisan, N. 2006. Combinatorial auctions with decreasing marginal utilities. *Games Econ. Behav.* 55(2): 270–296.

Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, 125–131.

Nguyen, T. T.; and Rothe, J. 2014. Minimizing envy and maximizing average Nash social welfare in the allocation of indivisible goods. *Discrete Applied Mathematics* 179: 54–68.

Plaut, B.; and Roughgarden, T. 2018. Almost Envy-Freeness with General Valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2584–2603.

Steinhaus, H. 1948. The Problem of Fair Division. *Econometrica* 16(1): 101–104.