On Optimal Tradeoffs between EFX and Nash Welfare*

Michal Feldman[1,2], **Simon Mauras**[1], **Tomasz Ponitka**[1],

[1]Tel Aviv University,
[2]Microsoft ILDC
michal.feldman@cs.tau.ac.il, smauras@tauex.tau.ac.il, tomaszp@mail.tau.ac.il

Abstract

A major problem in fair division is how to allocate a set of indivisible resources among agents fairly and efficiently. The goal of this work is to characterize the tradeoffs between two well-studied measures of fairness and efficiency — *envy freeness up to any item* (EFX) for fairness, and *Nash welfare* for efficiency — by saying, for given constants α and β, whether there exists an α-EFX allocation that guarantees a β-fraction of the maximum Nash welfare (β-MNW). For additive valuations, we show that for any $\alpha \in [0, 1]$, there exists a partial allocation that is α-EFX and $\frac{1}{\alpha+1}$-MNW. This tradeoff turns out to be tight (for every α) as demonstrated by an impossibility result that we give. We also show that for $\alpha \in [0, \varphi-1 \approx 0.618]$ these partial allocations can be turned into complete allocations where all items are assigned. Furthermore, for any $\alpha \in [0, 1/2]$, we show that the tight tradeoff of α-EFX and $\frac{1}{\alpha+1}$-MNW with complete allocations holds for the more general setting of *subadditive* valuations. Our results improve upon the current state of the art, for both additive and subadditive valuations, and match the best-known approximations of EFX under complete allocations, regardless of Nash welfare guarantees. Notably, our constructions for additive valuations also provide EF1 and constant approximations for maximin share guarantees.

Introduction

A common resource allocation setting has m indivisible goods and n agents with (possibly different) preferences over bundles of goods. One of the biggest questions in such scenarios, dating back to Steinhaus (1948), is how to allocate goods *fairly* among agents, given their preferences. Another central goal is to divide the items *efficiently* so that the collective welfare of the agents, represented by some efficiency measure, is maximized. There is an extensive line of work on providing allocations that achieve both criteria simultaneously, e.g., (Caragiannis et al. 2009; Bertsimas, Farias, and Trichakis 2011; Bei et al. 2019).

The aim of this paper is to analyze the tradeoffs between fairness and efficiency for two standard notions, namely *EFX* for fairness and *Nash welfare* for efficiency. Previous work has discovered partial results for this problem (Caragiannis, Gravin, and Huang 2019; Garg et al. 2023). In this work, we provide a complete characterization of this tradeoff when agents have additive valuations and a partial characterization for the more general case of subadditive valuations.

Resource allocation problems. A resource allocation problem is given by a set of m indivisible goods, a set of n agents, and a valuation function $v_i : 2^{[m]} \to \mathbb{R}^{\geq 0}$ for every agent i, which assigns a value $v_i(S)$ to every bundle of goods $S \subseteq [m]$. The most common type of valuation functions are *additive* valuations, where there exist values v_{ig} such that $v_i(S) = \sum_{g \in S} v_{ig}$ for every bundle S and agent i. We also consider the widely studied class of *subadditive* valuations, where $v_i(S \cup T) \leq v_i(S) + v_i(T)$ for every $S, T \subseteq [m]$. Subadditive valuations constitute the frontier of complement-free valuations (Lehmann, Lehmann, and Nisan 2006).

An allocation of m goods amongst n agents is given by a collection $X = (X_1, \ldots, X_n)$ where the sets X_i and X_j are disjoint for every pair of distinct agents i and j. The set X_i represents the set of goods allocated to agent i. An allocation is said to be *complete* if $\bigcup_i X_i = [m]$, and *partial* if some items might be unallocated. In resource allocation problems, one is typically interested in complete allocations, but in some contexts, including in the context of this paper, it also makes sense to consider partial allocations.

Fairness notions. The problem of finding allocations that are fair is called fair division. Many different definitions of fairness have been proposed in the fair division domain, see (Amanatidis et al. 2022) for a survey. One of the most compelling of those notions is that of *envy-freeness*, introduced by Foley (1966). An allocation $X = (X_1, \ldots, X_n)$ is envy-free (EF) if for every pair of agents i and j, it holds that $v_i(X_i) \geq v_i(X_j)$, namely, every agent (weakly) prefers her own allocation over that of any other agent. In the context of envy-freeness, it is usually required that the allocations are complete. Indeed, without this constraint, every instance vacuously admits an EF allocation, namely one where no item is allocated.

A major weakness of EF is that even the simplest settings may not admit complete EF allocations. For example, if there is a single good and two agents with value 1 for this good, then in any complete allocation, only one of the agents gets the single good and the other agent is envious.

Consequently, various relaxations of EF have been introduced. One such relaxation is *envy-freeness up to one item*

*See (Feldman, Mauras, and Ponitka 2023) for the full version.
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

(EF1), defined by Budish (2011). An allocation X is EF1 if for every pair of agents i and j, there exists an item $g \in X_j$, such that $v_i(X_i) \geq v_i(X_j - g)$, namely, agent i does not envy agent j after removing some item from j's bundle. Using a well-known technique of eliminating envy cycles, it is not too difficult to show that a complete EF1 allocation always exists (Lipton et al. 2004). For concreteness, consider the following example.

Example 1. Consider a setting with three items, $\{a, b, c\}$, and two agents with identical additive valuations, where $v(a) = v(b) = 1$ and $v(c) = 2$.

In Example 1, the allocation $X_1 = \{a\}$ and $X_2 = \{b, c\}$ is EF1. Indeed, after removing c from X_2, agent 1 has no envy towards agent 2.

A stronger notion than EF1 is envy freeness up to *any* item (EFX), introduced by Caragiannis et al. (2016). An allocation X is EFX if for every pair of agents i and j, and *every* item $g \in X_j$, it holds that $v_i(X_i) \geq v_i(X_j - g)$. This is a stronger condition than EF1 since the requirement for no envy applies to the removal of *any* item from j's bundle. For instance, in Example 1, the allocation $X_1 = \{a\}$ and $X_2 = \{b, c\}$ is not EFX, since after removing b from X_2, agent 1 still envies agent 2.

Unfortunately, complete EFX allocations are not known to exist even for additive valuations, except in several special cases such as identical valuations (Plaut and Roughgarden 2018), identical items (Lipton et al. 2004), or the case of three agents (Chaudhury, Garg, and Mehlhorn 2020). Arguably, the existence of EFX allocations is the most enigmatic problem in fair division (Procaccia 2020).

A natural step to approach the existence of EFX allocations is to consider the approximate notion of α-EFX, defined by Plaut and Roughgarden (2018). An allocation X is α-EFX, for some $\alpha \in [0, 1]$, if for every pair of agents i and j, and every item $g \in X_j$, it holds that $v_i(X_i) \geq \alpha \cdot v_i(X_j - g)$. The existence of α-EFX allocations has been studied for several classes of valuations. In particular, previous work has established (i) the existence of $(\varphi - 1)$-EFX allocations for additive valuations, where $\varphi \approx 1.618$ is the golden ratio (Amanatidis, Markakis, and Ntokos 2020), and (ii) the existence of $1/2$-EFX allocations for subadditive valuations (Plaut and Roughgarden 2018).

Efficiency notions. Common measures that capture the efficiency of an allocation are: (i) social welfare — the sum of agent values, $\mathrm{SW}(X) = \sum_{i \in [n]} v_i(X_i)$, and (ii) Nash welfare — the geometric mean of agent values, $\mathrm{NW}(X) = \prod_{i \in [n]} v_i(X_i)^{1/n}$. In this paper, as in many previous studies, we consider the Nash welfare notion.

The most important property of the Nash welfare is that it encourages more balanced allocations relative to social welfare. For example, consider a setting with two agents and two items, where every agent values every item at 1. The unique maximum Nash welfare (MNW) allocation is the one that allocates one item per agent. In terms of social welfare, however, every allocation is equally good, including one that gives two items to one of the agents and none to the other, as all complete allocations have social welfare of 2.

Remarkably, under additive valuations, every allocation that maximizes Nash welfare is EF1 (Caragiannis et al. 2016). In fact, maximizing Nash welfare is the only welfarist rule satisfying EF1 (Yuen and Suksompong 2023). Similar results have been shown for other valuation classes: every allocation that maximizes Nash welfare is $1/4$-EF1 when the valuations are subadditive (Wu, Li, and Gan 2021), and it is EFX when the valuations are additive and bi-valued (Amanatidis et al. 2020), or when they are submodular and dichotomous (Babaioff, Ezra, and Feige 2021).

An allocation X is said to be β-max Nash welfare (β-MNW) if the Nash welfare of X is at least a β fraction of the maximum Nash welfare.

Fairness vs efficiency trade-off. A natural question is whether fairness and efficiency can be achieved simultaneously. Bei et al. (2019) studied this question with respect to the efficiency measure of the social welfare for instances with additive valuations. Specifically, they provided bounds for the "price of fairness" with respect to several fairness notions; namely, the fraction of the maximum social welfare that can be achieved, when constrained by the corresponding fairness condition.

They showed that the price of fairness of EFX is $\Omega(\sqrt{n})$, i.e., there are instances with additive valuations in which no EFX allocation can achieve more than a $O(1/\sqrt{n})$ fraction of the optimal social welfare. On the other hand, it is possible to find (partial) EFX allocations that obtain a constant fraction of the maximum Nash welfare (Caragiannis, Gravin, and Huang 2019). This motivates the study of optimal tradeoffs between α-EFX and β-MNW, which is the focus of this paper.

Notably, the original motivation for considering α-EFX has been the embarrassment around the EFX existence problem, and the original motivation for considering β-MNW has been the NP-hardness of maximizing Nash welfare (Ramezani and Endriss 2009; Nguyen, Roos, and Rothe 2012). As it turns out, an equally important motivation for studying approximate notions of EFX and MNW is the fact that one may come at the expense of another. Thus, understanding the tradeoffs between these fairness and efficiency measures is crucial when designing a resource allocation scheme.

Prior to our work, the following tradeoffs (demonstrated in Figure 1) have been known: (i) instances with additive valuations admit a partial EFX allocation that is $1/2$-MNW (Caragiannis, Gravin, and Huang 2019), (ii) instances with additive valuations admit a complete allocation that is $(\varphi - 1)$-EFX (with no Nash welfare guarantees) (Amanatidis, Markakis, and Ntokos 2020), and (iii) instances with subadditive valuations admit a complete allocation that is $1/2$-EFX and $1/2$-MNW (Garg et al. 2023).

As described in the next section, we extend these results to give a more complete picture of the optimal tradeoffs between α-EFX and β-MNW.

Our Results

In this paper we provide results on the optimal trade-offs between approximate EFX and approximate maximum Nash

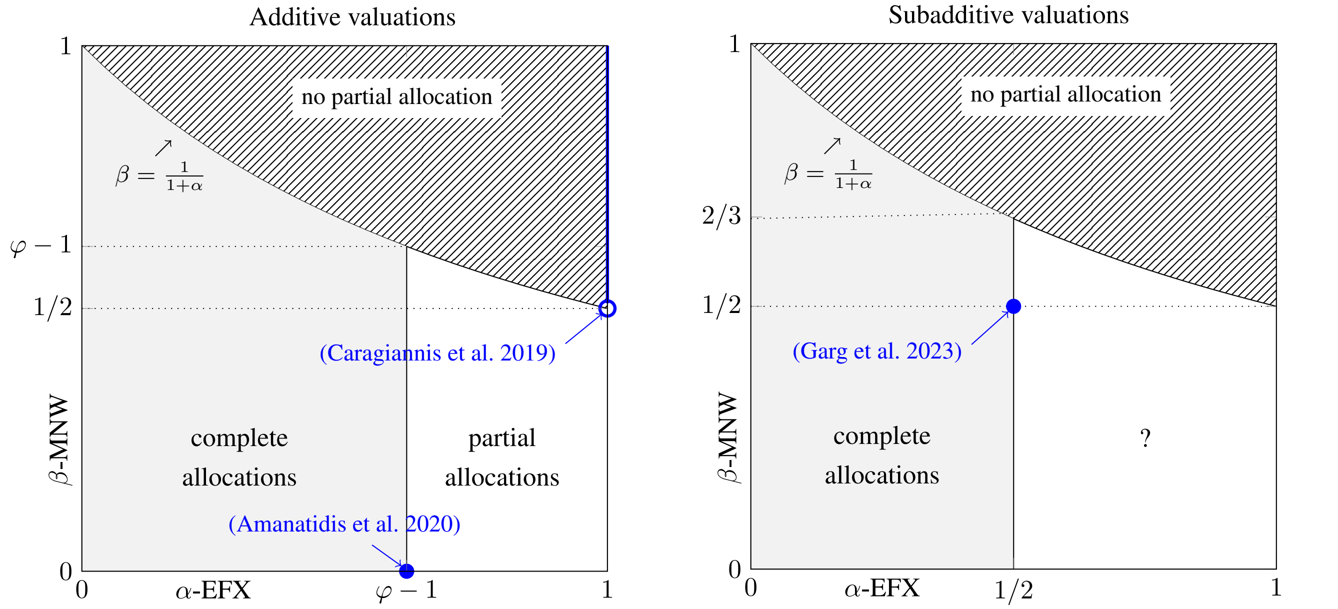

Figure 1: Trade-off between the existence of α-EFX and β-MNW allocations, for additive (left) and subadditive (right) valuations. Previous positive results are represented by dots (complete dots for complete allocations, and hollow dots for partial allocations). Dots that lie on the curves belong to the regions of positive results. All positive results for additive valuations also guarantee EF1.

welfare (MNW), for both additive valuations and subadditive valuations. Our results are demonstrated in Figure 1, where the left and right figures correspond to additive and subadditive valuations, respectively.

Our first result gives existence guarantees on partial allocations with approximate EFX and approximate MNW (and EF1), for additive valuations.

Theorem 1. *Every instance with additive valuations admits a partial allocation that is α-EFX, EF1, and $\frac{1}{\alpha+1}$-MNW, for every $0 \leq \alpha \leq 1$.*

For $0 \leq \alpha \leq \varphi - 1 \approx 0.618$, we show that our partial allocations can be turned into complete ones without any loss. This is cast in the following theorem.

Theorem 2. *Every instance with additive valuations admits a complete allocation that is α-EFX, EF1, and $\frac{1}{\alpha+1}$-MNW, for every $0 \leq \alpha \leq \varphi - 1 \approx 0.618$.*

In particular, Theorem 2 extends the existence of $(\varphi - 1)$-EFX complete allocation (Amanatidis, Markakis, and Ntokos 2020) to $(\varphi - 1)$-EFX complete allocation that is also $(\varphi - 1)$-MNW. Note that, by Theorem 4 below, $(\varphi - 1)$ is the highest possible MNW approximation of a (complete or partial) $(\varphi - 1)$-EFX allocation.

Our final positive result shows that for any $0 \leq \alpha \leq 1/2$, the same trade-off between EFX and MNW approximation extends to subadditive valuations.

Theorem 3. *Every instance with subadditive valuations admits a complete allocation that is α-EFX and $\frac{1}{\alpha+1}$-MNW, for every $0 \leq \alpha \leq 1/2$.*

In particular, Theorem 3 extends the existence of a $1/2$-EFX and $1/2$-MNW complete allocation by (Garg et al. 2023) to the existence of a $1/2$-EFX complete allocation that is also $2/3$-MNW. Note that, by Theorem 4 below, $2/3$ is the highest possible MNW approximation of a (complete or partial) $1/2$-EFX allocation, even for additive valuations.

More generally, our tradeoffs are tight, as the following theorem shows.

Theorem 4 (Impossibility results)**.** *For every $0 < \alpha \leq 1$ and $\beta > \frac{1}{\alpha+1}$, there exists an instance with additive (and hence also subadditive) valuations that admits no allocation (even partial) that is α-EFX and β-MNW. Moreover, for every $\alpha, \beta > 0$, there exists an instance with monotone valuations that admits no allocation that is α-EFX and β-MNW.*

Computational remarks. While our positive results (Theorems 1, 2, 3) are stated as existence results, to prove the existence of allocations with the stated guarantees, we construct polynomial-time algorithms that find such allocations, given a max Nash welfare (MNW) allocation as input.

Moreover, the algorithms used to prove Theorems 1 and 3 apply also when given an *arbitrary* allocation as input. In particular, these are poly-time algorithms which, given an arbitrary allocation X as input, produce an α-EFX allocation that gives at least a $1/(\alpha + 1)$ fraction of the Nash welfare of X, under the corresponding conditions.

Thus, when given black-box access to an algorithm that computes a β-MNW allocation, they provide an α-EFX allocation that is also $\beta/(\alpha + 1)$-MNW. This extension is important in light of the fact that finding a MNW allocation is NP-hard, even for additive valuations (Ramezani and Endriss 2009), while constant approximation algorithms exist, even for subadditive valuations.

In particular, using this extension, combined with the

$(e^{-1/e} - \varepsilon)$-MNW approximation for additive valuations (Barman, Krishnamurthy, and Vaish (2018)) and the $(1/4 - \varepsilon)$-MNW approximation for submodular valuations (Garg et al. (2023)) and the constant-factor-MNW approximation for subadditive valuations (Dobzinski et al. (2023)), our algorithms find in polynomial time a partial α-EFX allocation with constant-factor-MNW, for any $0 \leq \alpha \leq 1/2$ when valuations are subadditive, and any $0 \leq \alpha \leq 1$ when valuations are additive. See section A of the full version for more details.

Maximin share guarantees. Even though the allocations that we construct in the proof of Theorem 2 are designed with EFX and MNW in mind, they also satisfy other desirable fairness notions related to the maximin share guarantee. More specifically, on top of α-EFX, EF1, and $\frac{1}{\alpha+1}$-MNW, these allocations are also $\frac{\alpha}{\alpha^2+1}$-GMMS (and hence $\frac{\alpha}{\alpha^2+1}$-MMS) and $(\varphi-1)$-PMMS. See section C of the full version for definitions and more details.

It is worth noting that while the result of Amanatidis, Markakis, and Ntokos (2020) guarantees for every instance with additive valuations the existence of a complete allocation that is $(\varphi-1)$-EFX, EF1, ($\frac{2}{2+\varphi} \approx 0.55$)-GMMS, (2/3)-PMMS with no efficiency guarantees, our result gives a complete allocation that is $(\varphi-1)$-EFX, EF1, ($\frac{\varphi-1}{\varphi^2-2\varphi} \approx 0.44$)-GMMS, $(\varphi-1)$-PMMS, and $(\varphi-1)$-MNW.

Model and Preliminaries

Our Model

We consider settings with a set $[n] = \{1, \ldots, n\}$ of n agents, and a set $[m] = \{1, \ldots, m\}$ of m items. Every agent $i \in [n]$ has a valuation function denoted by $v_i : 2^{[m]} \to \mathbb{R}^{\geq 0}$, which assigns a real value $v_i(S)$ to every set of items $S \subseteq [m]$. We consider the following valuation classes:

- additive: $v_i(S \cup T) = v_i(S) + v_i(T)$ for any disjoint $S, T \subseteq [m]$.
- subadditive: $v_i(S \cup T) \leq v_i(S) + v_i(T)$ for any (not necessarily disjoint) $S, T \subseteq [m]$.

An instance of a resource allocation problem is given by a collection of valuation functions $v_1, \ldots, v_n$ over the set of m items. Throughout this paper, we use the standard notation $v_i(g) = v_i(\{g\})$ for $g \in [m]$ and $Z - g = Z \setminus \{g\}$ and $Z + g = Z \cup \{g\}$ for $\{g\}, Z \subseteq [m]$.

An allocation $X = (X_1, \ldots, X_n)$ is a collection of disjoint subsets of items, i.e., $X_i \cap X_j = \emptyset$ for $i \neq j$ and $X_i \subseteq [m]$ for every $i \in [n]$. We say that X is complete if $\bigcup_{i \in [n]} X_i = [m]$, and that it is partial otherwise.

The Nash welfare of an allocation X is denoted by $\mathrm{NW}(X) = \prod_{i \in [n]} v_i(X_i)^{1/n}$. For every instance, a maximum Nash welfare (MNW) allocation X is any allocation that maximizes $\mathrm{NW}(X)$ among all possible allocations. We say that an allocation Z is β-MNW for some $\beta \in [0, 1]$ if it holds that $\mathrm{NW}(Z) \geq \beta \cdot \mathrm{NW}(X)$.

An allocation X is α-EFX if for every $i, j \in [n]$ and $g \in X_j$, it holds that $v_i(X_i) \geq \alpha \cdot v_i(X_j - g)$. Whenever this condition is violated, i.e., $v_i(X_i) < \alpha \cdot v_i(X_j - g)$ for some $g \in Z_j$, then we say that the agent i envies agent j in the α-EFX sense. We say that an allocation is EFX if it is 1-EFX.

From Partial Allocations to Complete Allocations

The main tool that we use to turn partial allocations into complete ones with the same fairness and Nash welfare guarantees is the envy-cycles procedure (Lipton et al. 2004). In this procedure, as long as the allocation is not complete, we take one of the unallocated items and give it to an agent that no other agent envies, or if there is no such agent, then we find a cycle of agents with the property that each agent prefers the bundle of the following agent, and then we improve the allocation by moving the bundles along the cycle.

For our purposes, the crucial observation is that if the value of any agent for any of the unallocated items is bounded, then this procedure preserves the initial EFX guarantees. The same observation was used by Garg et al. (2023) to show that there exists a complete $1/2$-EFX and $1/2$-MNW allocation for subadditive valuations, by Amanatidis, Markakis, and Ntokos (2020) to show that there exists a complete $(\varphi - 1)$-EFX allocation (with no guarantees on Nash welfare) for additive valuations, and by Farhadi et al. (2021) to show that there exists a complete 0.73-EFR (envy-free up to a random good) allocation for additive valuations.

More formally, consider the following definition.

Definition 1 (γ-separation)**.** Let $Z = (Z_1, \ldots, Z_n)$ be a partial allocation, and let U be the set of unallocated items in Z. We say that Z satisfies γ-*separation*, for some $\gamma \in [0, 1]$, if for every agent i, it holds that $\gamma \cdot v_i(Z_i) \geq v_i(x)$ for all $x \in U$, i.e., agent i prefers Z_i significantly more (by a factor of $1/\gamma$) than any single unallocated item.

The following key lemma is based on the fact that for any partial allocation that is α-EFX and γ-separated, the envy cycles procedure produces a complete $\min(\alpha, 1/(1+\gamma))$-EFX allocation with weakly higher Nash welfare.

Lemma 1. *Let $Z = (Z_1, \ldots, Z_n)$ be a partial allocation that is α-EFX and satisfies γ-separation. Then, there exists a complete allocation Y that is $\min(\alpha, 1/(1+\gamma))$-EFX such that $\mathrm{NW}(Y) \geq \mathrm{NW}(Z)$. Moreover, if Z is EF1, then Y is also EF1.*

With this lemma in hand, in order to establish the existence of complete allocations with good EFX and MNW guarantees, it suffices to produce partial allocations with good Nash welfare and separation guarantees.

Additive Valuations

In this section, we study instances with additive valuations. We first prove Theorem 1 by constructing a partial allocation with the desired EFX and MNW guarantees, and then we prove Theorem 2 by using Lemma 1 to turn the partial allocation into a complete allocation.

Partial Allocations

Our main result in this section is the following:

Theorem 1. *Every instance with additive valuations admits a partial allocation that is α-EFX, EF1, and $\frac{1}{\alpha+1}$-MNW, for every $0 \leq \alpha \leq 1$.*

Algorithm 1: Additive valuations.

Input $(X_1, \ldots, X_n)$ is a complete MNW allocation.
Output $(M_1, \ldots, M_n)$ is a partial α-EFX and $\frac{1}{\alpha+1}$-MNW allocation.

1: match i to J means $M_i \leftarrow J$
2: unmatch Z_i means $M_u \leftarrow \perp$ if there is u matched to Z_i
3: **procedure** ALG
4: $\quad Z \leftarrow (X_1, \ldots, X_n)$
5: $\quad M \leftarrow (\perp, \ldots, \perp)$
6: $\quad$ **while** there is an agent i with $M_i = \perp$ **do**
7: $\quad\quad i^\star \leftarrow$ any agent with $M_i = \perp$
8: $\quad\quad$ **if** $(v_{i^\star}(Z_{i^\star}) \geq \alpha \cdot v_{i^\star}(Z_j - g)$ for all j and $g \in Z_j$
9: $\quad\quad\quad$ **and** $Z_{i^\star} = X_{i^\star})$ **or**
10: $\quad\quad\quad (v_{i^\star}(Z_{i^\star}) \geq v_{i^\star}(Z_j - g)$ for all j and $g \in Z_j$
11: $\quad\quad\quad$ **and** $Z_{i^\star} \neq X_{i^\star})$ **then**
12: $\quad\quad\quad$ unmatch $Z_{i^\star}$
13: $\quad\quad\quad$ match $i^\star$ to $Z_{i^\star}$
14: $\quad\quad$ **else**
15: $\quad\quad\quad j^\star, g \leftarrow$ any $j^\star$ and $g \in Z_{j^\star}$
16: $\quad\quad\quad\quad$ that maximize $v_{i^\star}(Z_{j^\star} - g)$
17: $\quad\quad\quad$ unmatch $Z_{j^\star}$
18: $\quad\quad\quad$ change $Z_{j^\star}$ to $Z_{j^\star} - g$
19: $\quad\quad\quad$ match $i^\star$ to $Z_{j^\star}$
20: $\quad\quad$ **end if**
21: $\quad$ **end while**
22: $\quad$ **return** $(M_1, \ldots, M_n)$
23: **end procedure**

To prove Theorem 1, we use Algorithm 1, which generalizes the algorithm used in Caragiannis, Gravin, and Huang (2019) to produce a partial allocation that is (fully) EFX and $1/2$-MNW. For the special case of $\alpha = 1$, our algorithm is essentially the same as the original one. A very similar modification of the original algorithm was also used by Garg et al. (2023) for the special case of $\alpha = 1/2$; however, their modification differs from Algorithm 1 in that it does not guarantee EF1 for additive valuation. Both of the algorithms used in Caragiannis, Gravin, and Huang (2019) and in Garg et al. (2023) were described using a certain notion of EFX feasibility graphs. Here, we present Algorithm 1 in a different way without referring to that notion, which also gives a more direct description of the previous algorithms.

In Algorithm 1, we start with a maximum Nash welfare allocation $X = (X_1, \ldots, X_n)$, and we iteratively drop items that cause envy, possibly reordering the bundles between agents at the same time. We continue to do so until we reach an α-EFX and EF1 allocation $M = (M_1, \ldots, M_n)$. For example, in the instance described in Example 1, with allocation $X_1 = \{a\}, X_2 = \{b, c\}$, we first remove item b from X_2; this eliminates the envy of agent 1 for agent 2.

We refer to M as a matching between agents and bundles of items. We say that an agent i is *matched* to M_i if $M_i \neq \perp$, and that i is *unmatched* otherwise. The algorithm maintains a set of bundles $(Z_1, \ldots, Z_n)$ which are initially set to $Z_i = X_i$. Throughout the algorithm, every matched agent i is matched to some bundle Z_j. We say that Z_j is *matched* to i if $M_i = Z_j$. The matching M never assigns the same bundle to two agents, and at the end of the algorithm, every bundle Z_j is matched to some agent i. This is formally stated in the following claim.

Claim 1. *At the end of each iteration of the algorithm, for any agent i, it holds that $Z_i \subseteq X_i$ and $M_i \in \{\perp\} \cup \{Z_j \mid j \in [n]\}$. Moreover, for any distinct agents i_1 and i_2 with $M_{i_1}, M_{i_2} \neq \perp$, it holds that $M_{i_1} \neq M_{i_2}$.*

The high-level idea of the algorithm is as follows: The algorithm starts off by setting $Z_i = X_i$ for every agent i; recall that X is an MNW allocation. At first, no agent is matched. Then, the algorithm proceeds by shrinking the bundles Z_j's and matching them to agents in a way that eliminates envy. More precisely, as long as there is an unmatched agent $i^\star$, one of the following two operations takes place. (i) If $Z_{i^\star}$ is "good enough" for $i^\star$, then $i^\star$ is matched to $Z_{i^\star}$. The condition for $Z_{i^\star}$ to be good enough for $i^\star$ in the case where $Z_{i^\star} = X_{i^\star}$ is that if $i^\star$ gets $Z_{i^\star}$, then she does not envy any other bundle Z_j for any agent j, in the α-EFX sense. If $Z_{i^\star} \subsetneq X_{i^\star}$, then the condition is that if $i^\star$ gets $Z_{i^\star}$, then she does not envy any other bundle Z_j for any agent j, in the EFX sense, which is a stronger requirement than in the $Z_{i^\star} = X_{i^\star}$ case. (ii) Otherwise, if $Z_{i^\star}$ is not good enough for $i^\star$, the algorithm picks the most valuable (from the perspective of $i^\star$) *strict* subset of $Z_{j^\star}$ for some $j^\star$, shrinks $Z_{j^\star}$ to the chosen strict subset, and matches $i^\star$ to the new $Z_{j^\star}$ (leaving the agent previously matched to $Z_{j^\star}$, if any, unmatched). It can be shown that in this case, $i^\star$ does not envy any other bundle Z_j in the stronger sense of EFX (rather than α-EFX).

Let us first make a few simple observations that are crucial to the analysis of the algorithm. First, the matching of the algorithm ensures that $M_{i^\star}$ is "good enough" for $i^\star$ whenever $i^\star$ is matched. This in fact implies that the final allocation M is α-EFX and EF1. Second, whenever an agent i with an untouched bundle is matched by the second operation (lines 15-19), she is matched to a bundle that she prefers significantly more (by a factor of $1/\alpha$) to Z_i. Third, any touched bundle, i.e., one from which the algorithm removed an element, is matched. These observations are formally stated in the following claim.

Claim 2. *Consider the state of the algorithm at the end of any iteration. Let i be any matched agent. It holds that if $Z_i = X_i$, then $v_i(M_i) \geq \alpha \cdot v_i(Z_j - g)$ for all j and all $g \in Z_j$, and if $Z_i \neq X_i$, then $v_i(M_i) \geq v_i(Z_j - g)$ for all j and all $g \in Z_j$. It also holds that if $M_i \neq Z_i$, then $v_i(M_i) > v_i(Z_i)$, and if $M_i \neq Z_i$ and $Z_i = X_i$, then $v_i(M_i) > (1/\alpha) \cdot v_i(Z_i)$. Finally, for any agent i with $Z_i \subsetneq X_i$, there is some agent u with $M_u = Z_i$.*

The crucial part of the analysis is to show that the items are removed conservatively so that for each agent her final bundle is worth at least a $1/(\alpha + 1)$ fraction of the bundle she started with, which yields a $1/(\alpha + 1)$ approximation to the maximum Nash welfare. This is the key lemma in our analysis.

Lemma 2. *At the end of the run of the algorithm, we have $v_i(Z_i) \geq \frac{1}{\alpha+1} \cdot v_i(X_i)$.*

To prove this, we assume that this condition is violated at some point, and we use this assumption to construct an allocation with higher Nash welfare than the initial one, which contradicts the assumption that the algorithm is given the Nash welfare maximizing allocation as input.

With this lemma in hand, we can prove Theorem 1.

Proof of Theorem 1. The allocation $(M_1, \dots, M_n)$ is α-EFX by Claims 1 and 2. Moreover, M is EF1 since either $Z_i \subsetneq X_i$ and the property follows from Claim 2, or $Z_i = X_i$ and then, since X is EF1 by the result of Caragiannis et al. (2016), for all agents j, it holds that $v_i(M_i) \geq v_i(Z_i) = v_i(X_i) \geq v_i(X_j - g) \geq v_i(Z_j - g)$ for some $g \in X_j$.

Fix any agent i. Let j be the unique agent matched to Z_i. By Claim 2, it holds that $v_i(Z_j) \geq v_i(Z_i)$. Therefore, by Lemma 2, $\prod_{i\in[n]} v_i(M_i)^{1/n} \geq \prod_{i\in[n]} v_i(Z_i)^{1/n} \geq \prod_{i\in[n]}(\frac{1}{\alpha+1} \cdot v_i(X_i))^{1/n} = \frac{1}{\alpha+1} \cdot \prod_{i\in[n]} v_i(X_i)^{1/n}$ and so the result follows. ☐

Complete Allocations

In this section, we provide the following result.

Theorem 2. *Every instance with additive valuations admits a complete allocation that is α-EFX, EF1, and $\frac{1}{\alpha+1}$-MNW, for every $0 \leq \alpha \leq \varphi - 1 \approx 0.618$.*

Here, the crucial part of the analysis is to provide the appropriate bounds on the value of the unallocated items so that we can use Lemma 1. Let $(Z_1, \dots, Z_n)$ be the bundles at the end of the run of Algorithm 1. The following lemma is the key component in the proof of Theorem 2. It offers additional analysis of Algorithm 1, showing that the allocation returned by this algorithm is α-separated.

Lemma 3. *The partial allocation $(Z_1, \dots, Z_n)$ satisfies α-separation.*

The remainder of this section is dedicated to proving this lemma. We prove the α-separation property by constructing an allocation $\widehat{X}$ which is built from X and using the optimality of X to infer that the Nash welfare of $\widehat{X}$ is at most the Nash welfare of X. The construction of $\widehat{X}$ is based on a non-trivial redistribution of the items among agents, which requires additional definitions, as follows.

Definition 2 (Touching)**.** For any agent i with $Z_i \subsetneq X_i$, we say that agent k was the *last one to touch* i if in the last iteration in Algorithm 1 with $j^\star = i$, it was the case that $i^\star = k$. If, on the other hand, $Z_i = X_i$, then we say that agent i is *untouched*.

Definition 3 (Touching sequence)**.** Given any agent i, we define a *touching sequence* $k_1, \dots, k_\ell$ in the following way. Let $k_1 = i$. For every $s \geq 1$ define k_{s+1} to be the last agent to touch k_s, until either (i) k_s is untouched, or (ii) the last agent to touch k_s is already in the sequence $k_1, \dots, k_{s-1}$. Let $\ell = s$ for s where the process ends.

We also use the following technical lemma.

Lemma 4. *Let i and j be any distinct agents. Let $\widehat{X}_i = (X_i \setminus Z_i) \cup (X_j \setminus Z_j)$. Suppose that $v_i(\widehat{X}_i) \leq \alpha \cdot v_i(X_i)$. Then, $v_i(x) \leq \alpha \cdot v_i(Z_i)$ for all $x \in X_j \setminus Z_j$.*

We now present a simplified proof of Lemma 3.

Simplified proof of Lemma 3. Fix any agents i and j with $Z_j \neq X_j$ and any item $x \in X_j \setminus Z_j$. The goal is to show that $v_i(x) \leq \alpha \cdot v_i(Z_i)$. Let $k_1, \dots, k_\ell$ be the touching sequence starting with $k_1 = i$ (Definition 3). In the simplified version of the proof, we assume that $i \neq j$, that k ends with condition (i), and that $j = k_t$ for some $2 \leq t \leq \ell$.

Consider the following allocation $(\widehat{X}_1, \dots, \widehat{X}_n)$.

$$\begin{aligned} \widehat{X}_j &= Z_{k_{t-1}} \\ \widehat{X}_i &= (X_i \setminus Z_i) \cup (X_j \setminus Z_j) \\ \widehat{X}_{k_\ell} &= X_{k_\ell} \cup Z_{k_{\ell-1}} \\ \widehat{X}_{k_s} &= (X_{k_s} \setminus Z_{k_s}) \cup Z_{k_{s-1}} \qquad \text{for } s \notin \{1, t, \ell\} \end{aligned}$$

where $\widehat{X}_r = X_r$ for any $r \notin \{k_1, \dots, k_\ell\}$. By Lemma 2 and Claim 2, it holds that $v_j(\widehat{X}_j) = v_j(Z_{k_{t-1}}) \geq v_j(Z_j) \geq (1/(\alpha+1)) \cdot v_j(X_j)$. Similarly, for $s \notin \{1, t, \ell\}$, it holds that $v_{k_s}(\widehat{X}_{k_s}) = v_{k_s}((X_{k_s} \setminus Z_{k_s}) \cup Z_{k_{s-1}}) \geq v_{k_s}(X_{k_s})$. Finally, since $X_{k_\ell} = Z_{k_\ell}$, Claim 2 gives $v_{k_\ell}(\widehat{X}_{k_\ell}) = v_{k_\ell}(X_{k_\ell} \cup Z_{k_{\ell-1}}) \geq v_{k_\ell}(X_{k_\ell}) + (1/\alpha) \cdot v_{k_\ell}(X_{k_\ell}) = ((\alpha+1)/\alpha) \cdot v_{k_\ell}(X_{k_\ell})$. Combining all the inequalities above gives, by optimality of X, that

$$\begin{aligned} 1 &\geq \frac{v_j(\widehat{X}_j)}{v_j(X_j)} \cdot \frac{v_i(\widehat{X}_i)}{v_i(X_i)} \cdot \frac{v_{k_3}(\widehat{X}_{k_3})}{v_{k_3}(X_{k_3})} \cdots \frac{v_{k_\ell}(\widehat{X}_{k_\ell})}{v_{k_\ell}(X_{k_\ell})} \\ &\geq (1/(\alpha+1)) \cdot \frac{v_i(\widehat{X}_i)}{v_i(X_i)} \cdot 1 \cdots 1 \cdot ((\alpha+1)/\alpha). \end{aligned}$$

After rearranging we can apply Lemma 4 which proves α-separation. ☐

Finally, we prove the main theorem of this section.

Proof of Theorem 2. Since $v_i(M_i) \geq v_i(Z_i)$ by Claim 2, it follows from Lemma 3 that the allocation $(M_1, \dots, M_n)$ returned by the algorithm is α-separated. Hence, by Lemma 1, there exists a complete allocation that is $\min(\alpha, \frac{1}{\alpha+1})$-EFX, EF1, and $\frac{1}{\alpha+1}$-MNW. This proves the theorem since for $\alpha \leq \varphi - 1$ it holds that $\frac{1}{\alpha+1} \geq \alpha$ and so $\min(\alpha, \frac{1}{\alpha+1}) = \alpha$. ☐

Subadditive Valuations

In this section, we study settings with subadditive valuations. We first prove Theorem 5 by constructing a partial allocation with the desired EFX and MNW guarantees, and then we prove Theorem 3 by using Lemma 1 to turn the partial allocation into a complete allocation.

Partial Allocations

In this section, we establish the following theorem.

Theorem 5. *Every instance with subadditive valuations admits a partial allocation that is α-EFX and $\frac{1}{\alpha+1}$-MNW, for every $0 \leq \alpha \leq 1/2$.*

The proof of Theorem 5 is based on a non-trivial modification of Algorithm 1 which, on a high level, proceeds as follows. As in Algorithm 1, we start with a maximum Nash welfare allocation and we keep removing items to eliminate envy. In the modified algorithm, however, we also allow some additional operations, e.g., we might match an agent to the bundle $X_j \setminus Z_j$ of items that have been removed from Z_j.

Due to space limitations, the complete description of the algorithm for subadditive valuations, along with its analysis, is deferred to Section 4 of the full version. Here, we discuss some of the challenges that arise in settings with subadditive valuations and some of the techniques we use to address them. For concreteness, we focus on the special case of $\alpha = 1/2$. Let us consider the following example.

Example 2. Consider an instance with n agents and a set of items $\{a_1, \ldots, a_{n-1}, b_1, \ldots, b_{n-1}, c\}$. For all $1 \le i \le n-1$, v_i is additive and given by $v_i(a_i) = 1/2+\varepsilon$, $v_i(b_i) = 1/2-\varepsilon$, and $v_i(r) = 0$ for all the remaining items r. Moreover, v_n is subadditive and given by $v_n(a_i) = 2 + \varepsilon$ and $v_n(b_i) = 0$ for all i, $v_n(c) = 1$, and $v_n(S) = \max_{g \in S} v_n(g)$.

Note that the unique maximum Nash welfare allocation in Example 2 is $X_i = \{a_i, b_i\}$ and $X_n = \{c\}$. Executing Algorithm 1 on this instance results in the following. Initially, $Z_i = X_i$ and $M_i = \perp$. Suppose that the algorithm first selects $i^\star = n$ as the unmatched agent and $j^\star = 1$ as the envied agent with $g = b_1$. Then, b_1 is removed from Z_1, i.e., $Z_1 = \{a_1\}$, and agent n is matched to Z_1, i.e., $M_n = Z_1$. Now, suppose that the algorithm selects agent $i^\star = 1$ as the unmatched agents. Then, Z_1 is unmatched from agent n and matched to agent 1, i.e., $M_1 = Z_1$ and $M_n = \perp$. Now, suppose that the same happens with all the remaining agents $2, \ldots, n-1$ so that it holds that $M_i = Z_i = \{a_i\}$ for all $1 \le i \le n-1$. Then, the algorithm finally matches agent n to Z_n which results in a 1/2-EFX allocation with Nash welfare of $(1/2 + \varepsilon)^{(n-1)/n}$.

As clearly demonstrated by this example, the technique from Algorithm 1 does not yield the 2/3-MNW guarantee given by Theorem 5. In fact, it cannot give any guarantee better than 1/2-MNW which matches the result of Garg et al. (2023) based on a variant of Algorithm 1.

Let us now discuss how to overcome the problem that arises in the instance described above. When the algorithm selects the unmatched agent $i^\star = n$ and the envied agent $j^\star = 1$ with $g = b_1$, the issue is that b_1 is too large of a part of $X_1 = \{a_1, b_1\}$ to be removed; more precisely, $v_1(X_1 - b_1) = v_1(a_1) = (1/2+\varepsilon) \cdot v_1(X_1) < (2/3) \cdot v_1(X_1)$ and so removing b_1 from Z_1 violates the condition given in Lemma 2 which was the key part in proving the desired Nash welfare guarantees.

What the modified algorithm does in this case is to set $Z_1, X_1 \leftarrow \{b_1\}$ and $M_n \leftarrow \{a_1\}$. Note that this violates two invariants that were crucial in the analysis of Algorithm 1. First, it no longer necessarily holds that $M_i \in \{\perp\} \cup \{Z_j : j \in [n]\}$. We call the bundles for which this condition does *not* hold the *blue* bundles. Second, it no longer holds that X_i remains constant throughout the execution of the algorithm. However, violating the two invariants allows the modified algorithm to satisfy an analogous guarantee to Lemma 2.

Lemma 5. *At the end of the execution of the modified algorithm, we have* $v_i(Z_i) \ge \frac{1}{\alpha+1} \cdot v_i(X_i)$.

Note that this lemma no longer immediately implies that the final allocation is $\frac{1}{\alpha+1}$-MNW, as was the case for Algorithm 1. This is because at the end of the run of the algorithm, the bundles X_i are not the original Nash welfare maximizing bundles. However, we make two crucial observations. (i) We can lower bound the Nash welfare of the final allocation using the number of the blue bundles. (ii) Every time that the modified algorithm shrinks a bundle X_i, it also adds an additional blue bundle to the current matching which compensates for that change. This is cast in the following lemma.

Lemma 6. *At the end of the execution of the modified algorithm, we have*

$$(i) \quad \prod_{i\in[n]} v_i(M_i) \ge \left(\frac{1}{\alpha+1}\right)^{n-\ell} \left(\frac{1}{\alpha}\right)^{\ell} \prod_{i\in[n]} v_i(X_i).$$

$$(ii) \quad \left(\frac{\alpha+1}{\alpha}\right)^{\ell} \prod_{i\in[n]} v_i(X_i) \ge \prod_{i\in[n]} v_i(X_i^0)$$

where X^0 is the initial MNW allocation and ℓ denotes the number of blue bundles in the final matching.

Let us also remark that since some of the crucial invariants of the analysis of Algorithm 1 are violated by the operations performed by the modified algorithm, it becomes much more challenging to ensure that the final allocation satisfies α-EFX. In particular, note that during the execution of Algorithm 1, once an agent i is matched to some bundle Z_j, this bundle remains good enough for i after any number of further iterations. The same might not hold for the modified algorithm. Indeed, note that when we remove a good from the bundle Z_r for some agent r, the bundle $X_r \setminus Z_r$ increases by this one good, and it might happen that the bundle $X_r \setminus Z_r$ is then matched to some other agent which potentially violates the α-EFX condition for agent i who is still matched to Z_j. This issue requires a careful treatment both in the design and the analysis of the modified algorithm; the details can be found in the full version.

Finally, we combine the statements of Lemma 6 to obtain the Nash welfare guarantees of the allocation returned by the modified algorithm.

Proof of Theorem 5. By Lemma 6, we get $\prod_{i\in[n]} v_i(M_i) \ge (\frac{1}{\alpha+1})^{n-\ell}(\frac{1}{\alpha})^{\ell} \prod_{i\in[n]} v_i(X_i) = (\frac{1}{\alpha+1})^{n}(\frac{\alpha+1}{\alpha})^{\ell} \prod_{i\in[n]} v_i(X_i) \ge (\frac{1}{\alpha+1})^{n} \prod_{i\in[n]} v_i(X_i^0)$ where X^0 is the initial MNW allocation and ℓ denotes the number of blue bundles in the final matching, and hence, M is $\frac{1}{\alpha+1}$-MNW. □

Complete Allocations

The main result of this section is the following theorem.

Theorem 3. *Every instance with subadditive valuations admits a complete allocation that is α-EFX and $\frac{1}{\alpha+1}$-MNW, for every $0 \le \alpha \le 1/2$.*

Proof of Theorem 3. Using the swapping with singletons technique from (Garg et al. 2023, Section 4.1), we can show that the partial allocation given in Theorem 5 can be turned into a 1-separated allocation while preserving the α-EFX and $\frac{1}{\alpha+1}$-MNW guarantees. The result then follows directly from Lemma 1. □

Acknowledgments

We thank Amos Fiat for invaluable discussions, and Vishnu V. Narayan for his feedback and helpful comments.

This work was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 866132), by the Israel Science Foundation (grant number 317/17), by an Amazon Research Award, by the NSF-BSF (grant number 2020788), and by a grant from the Tel Aviv University Center for AI and Data Science (TAD).

References

Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2020. Maximum Nash Welfare and Other Stories About EFX. In *IJCAI*, 24–30. ijcai.org.

Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; and Voudouris, A. A. 2022. Fair Division of Indivisible Goods: A Survey. In *IJCAI*, 5385–5393. ijcai.org.

Amanatidis, G.; Markakis, E.; and Ntokos, A. 2020. Multiple Birds with One Stone: Beating 1/2 for EFX and GMMS via Envy Cycle Elimination. In *AAAI*, 1790–1797. AAAI Press.

Babaioff, M.; Ezra, T.; and Feige, U. 2021. Fair and Truthful Mechanisms for Dichotomous Valuations. In *AAAI*, 5119–5126. AAAI Press.

Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Finding Fair and Efficient Allocations. In *EC*, 557–574. ACM.

Bei, X.; Lu, X.; Manurangsi, P.; and Suksompong, W. 2019. The Price of Fairness for Indivisible Goods. In *IJCAI*, 81–87. ijcai.org.

Bertsimas, D.; Farias, V. F.; and Trichakis, N. 2011. The Price of Fairness. *Oper. Res.*, 59(1): 17–31.

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

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.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2009. The Efficiency of Fair Division. In *WINE*, volume 5929 of *Lecture Notes in Computer Science*, 475–482. Springer.

Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *EC*, 305–322. ACM.

Chaudhury, B. R.; Garg, J.; and Mehlhorn, K. 2020. EFX Exists for Three Agents. In *EC*, 1–19. ACM.

Dobzinski, S.; Li, W.; Rubinstein, A.; and Vondrák, J. 2023. A constant factor approximation for Nash social welfare with subadditive valuations. *CoRR*, abs/2309.04656.

Farhadi, A.; Hajiaghayi, M. T.; Latifian, M.; Seddighin, M.; and Yami, H. 2021. Almost Envy-freeness, Envy-rank, and Nash Social Welfare Matchings. In *AAAI*, 5355–5362. AAAI Press.

Feldman, M.; Mauras, S.; and Ponitka, T. 2023. On Optimal Tradeoffs between EFX and Nash Welfare. *CoRR*, abs/2302.09633.

Foley, D. K. 1966. *Resource allocation and the public sector*. Yale University.

Garg, J.; Husic, E.; Li, W.; Végh, L. A.; and Vondrák, J. 2023. Approximating Nash Social Welfare by Matching and Local Search. In *STOC*, 1298–1310. ACM.

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 *EC*, 125–131. ACM.

Nguyen, T. T.; Roos, M.; and Rothe, J. 2012. A Survey of Approximability and Inapproximability Results for Social Welfare Optimization in Multiagent Resource Allocation. In *ISAIM*.

Plaut, B.; and Roughgarden, T. 2018. Almost Envy-Freeness with General Valuations. In *SODA*, 2584–2603. SIAM.

Procaccia, A. D. 2020. An answer to fair division's most enigmatic question: technical perspective. *Commun. ACM*, 63(4): 118.

Ramezani, S.; and Endriss, U. 2009. Nash Social Welfare in Multiagent Resource Allocation. In *AMEC/TADA*, volume 59 of *Lecture Notes in Business Information Processing*, 117–131. Springer.

Steinhaus, H. 1948. The problem of fair division. *Econometrica*, 16: 101–104.

Wu, X.; Li, B.; and Gan, J. 2021. Budget-feasible Maximum Nash Social Welfare is Almost Envy-free. In *IJCAI*, 465–471. ijcai.org.

Yuen, S. M.; and Suksompong, W. 2023. Extending the Characterization of Maximum Nash Welfare. *CoRR*, abs/2301.03798.