

Maximum Nash Welfare and Other Stories About EFX

Georgios Amanatidis^{1,2,3}, Georgios Birmpas⁴, Aris Filos-Ratsikas⁵,
Alexandros Hollender⁴ and Alexandros A. Voudouris^{1,4}

¹University of Essex

²University of Amsterdam

³Sapienza University of Rome

⁴University of Oxford

⁵University of Liverpool

{georgios.amanatidis, alexandros.voudouris}@essex.ac.uk,

{georgios.birmpas, alexandros.hollender}@cs.ox.ac.uk,

aris.filos-ratsikas@liverpool.ac.uk

Abstract

We consider the classic problem of fairly allocating indivisible goods among agents with additive valuation functions and explore the connection between two prominent fairness notions: maximum Nash welfare (MNW) and envy-freeness up to any good (EFX). We establish that an MNW allocation is always EFX as long as there are at most two possible values for the goods, whereas this implication is no longer true for three or more distinct values. As a notable consequence, this proves the existence of EFX allocations for these restricted valuation functions. While the efficient computation of an MNW allocation for two possible values remains an open problem, we present a novel algorithm for directly constructing EFX allocations in this setting. Finally, we study the question of whether an MNW allocation implies any EFX guarantee for general additive valuation functions under a natural new interpretation of approximate EFX allocations.

1 Introduction

Fair division refers to the general problem of allocating a set of resources to a set of agents in a way satisfying a desired fairness criterion. A well-known example of such a criterion is *envy-freeness* [Gamow and Stern, 1958], where each agent perceives the share she receives to be no worse than what any other agent receives. Since the problem was formally introduced by Banach, Knaster and Steinhaus [1948], fair division has attracted the attention of various scientific disciplines, including mathematics, economics, and political science. During the last two decades, the algorithmic aspects of fair division have been the focus of a particularly active line of work within the computer science community, e.g., see [Procaccia, 2016; Bouveret *et al.*, 2016] and references therein.

We consider the classic setting where the resources are indivisible goods that need to be fully allocated and the agents have additive valuation functions. One of the main challenges

in this setting is that classic fairness notions such as equitability, envy-freeness and proportionality—introduced several decades ago having divisible resources in mind—are impossible to satisfy. To see this for envy-freeness, it suffices to consider two agents and one good of value; one agent is going to be envious. This has led to the recent emergence of several weaker fairness notions (see *Related work*). As a result, there is a plethora of open questions about the existence, the computation and the interrelationships of such notions. In this work we focus on two of the most prominent: envy-freeness up to any good (EFX) and maximum Nash welfare (MNW).

EFX, introduced recently by Caragiannis *et al.* [2019b], is an additive relaxation of envy-freeness. Here an agent may envy another agent but only by the value of the least desirable good in the other agent’s bundle. While this added flexibility of EFX takes care of extreme pathological cases like the one mentioned above (2 agents, 1 good), this notion is not well understood yet. Despite the active interest in it, it is unknown whether EFX allocations always exist, even for 4 agents with additive valuation functions.¹ We consider the problem of showing the existence of EFX allocations to be one of the most intriguing currently open questions in fair division.

The *Nash social welfare* (or, simply, Nash welfare) is the geometric mean of the agents’ utilities. By considering maximum Nash welfare (MNW) allocations, i.e., allocations that maximize the product of the utilities, we achieve some kind of balance between the *efficiency* of the maximum utilitarian social welfare—the sum of the utilities—and the *individual fairness* of the maximum egalitarian social welfare—the minimum utility. Although not a fairness concept per se, MNW has strong ties to fairness. In the setting where the goods are divisible, each (possibly fractional) MNW allocation corresponds to a *competitive equilibrium of equal incomes*, a market equilibrium (under the assumption that all agents are endowed with the same budget) that is known to guarantee envy-freeness and Pareto optimality [Varian, 1974]. Even in our setting, Caragiannis *et al.* [2019b] showed that integral

¹Existence of EFX allocations for 3 agents was very recently shown by Chaudhury *et al.* [2020a].

MNW allocations, besides being Pareto optimal, are *envy-free up to one good* (EF1) and satisfy *maximin share fairness* up to a $\Theta(1/\sqrt{n})$ factor, where n is the number of agents. Both these guarantees are significantly weaker than EFX.

In general, MNW does not imply EFX. One of our goals is to identify the cases where it does, in terms of the allowed number of distinct values for the goods. For such cases, we want to obtain efficient algorithms for computing EFX allocations, either through maximizing the Nash welfare or directly. Since, in general, MNW does not even imply a non-trivial approximation of EFX, we further introduce a less stringent, yet natural, new interpretation of approximate EFX and investigate how it is related to MNW.

1.1 Contribution

There are two variants of EFX used in the related literature, depending on whether only the *positively* valued goods are considered or not; for the latter case we adopt the name EFX_0 suggested by Kyropoulou *et al.* [2019]. We start by establishing a strong algorithmic connection between the two variants (Proposition 2.1). Then we explore the relationship between maximizing the Nash welfare and achieving EFX or EFX_0 allocations. In doing so, we also obtain some interesting results for the individual notions. In particular:

- In case there are at most two possible values for the goods (*2-value instances*), we show that any allocation that maximizes the Nash welfare is EFX_0 (Theorem 3.1). This has the following two consequences:
 - (a) For any 2-value instance, there exists an EFX_0 allocation. Note that this is the first such existence result for non-identical valuations that holds for any number of agents and goods;
 - (b) For the special case of binary valuations, by adapting an algorithm of Barman *et al.* [2018c], we can efficiently construct an allocation that is both MNW and EFX_0 .

The implication $\text{MNW} \Rightarrow \text{EFX}_0$ is no longer true for three or more distinct values.

- For 2-value instances, the efficient computation of an MNW allocation remains an open problem. Nevertheless, we propose a polynomial-time algorithm for producing EFX_0 allocations in this case (Theorem 5.1). This algorithm, which we call `MATCH&FREEZE`, is based on repeatedly computing maximum matchings and “freezing” certain agents whenever they acquire too much value compared to their peers. We believe these novel ideas might be a stepping stone for proving the existence of EFX allocations in more general settings.
- For general additive valuations, we show that an MNW allocation does not guarantee any non-trivial approximation of EFX. However, we argue that the current definition of *approximate* EFX allocations is not always meaningful. Instead, we explore a different natural definition based on the idea of (hypothetically) augmenting an agent’s bundle until an EFX-like condition is satisfied. For this new benchmark, which we call *EFX-value*, we show that any MNW allocation is a $1/2$ -approximation of EFX (Theorem 6.2).

1.2 Related Work

As there is a vast literature on fair division, here we focus on the indivisible items setting and on related fairness notions. The concept of *envy-freeness up to one good* (EF1) was implicitly suggested by Lipton *et al.* [2004] and formally defined by Budish [2011]. Budish [2011] also introduced the notion of *maximin share* (MMS), which has been studied extensively [Kurokawa *et al.*, 2018; Amanatidis *et al.*, 2017; Barman and Krishnamurthy, 2017; Garg *et al.*, 2019; Ghodsi *et al.*, 2018; Garg and Taki, 2019] and has yielded several very interesting variants like *pairwise* MMS [Caragiannis *et al.*, 2019b], *groupwise* MMS [Barman *et al.*, 2018a], and MMS for groups of agents [Suksompong, 2018].

As already mentioned, EFX was introduced by Caragiannis *et al.* [2019b]. Plaut and Roughgarden [2018] defined what an α -approximate EFX (or α -EFX) allocation is and studied exact and approximate EFX allocations with both additive and general valuations. Most of their results, including the existence of EFX allocations for identical valuations, hold under the similar but stricter notion of EFX_0 which is implicitly introduced therein. The currently best 0.618-approximation of either EFX or EFX_0 for the additive case is due to Amanatidis *et al.* [2019]. For binary additive valuations, Aleksandrov and Walsh [2019] recently proposed an algorithm that produces EFX—but not necessarily EFX_0 —allocations. Independently and at the same time with our work, Babaioff *et al.* [2020] designed an algorithm that computes an EFX_0 allocation which maximizes the Nash welfare for submodular dichotomous valuations, a class that includes binary, but does not include general 2-value additive valuations.

Besides [Caragiannis *et al.*, 2019b], there are several recent papers which relate allocations that maximize (exactly or approximately) the Nash welfare with other fairness notions. Caragiannis *et al.* [2019a] showed that there exist *incomplete* allocations that are EFX and in which each agent receives at least half of the value they get in a MNW allocation; Chaudhury *et al.* [2020b] achieved the same with only a few unallocated goods. Garg and McGlaughlin [2019] showed how to get an allocation that 2-approximates the Nash welfare of an MNW allocation that is also *proportional up to one good*, satisfies a weak MMS guarantee and is Pareto optimal.

Since computing MNW allocations is an APX-hard problem [Lee, 2017], there is an active interest on special cases like binary valuations [Barman *et al.*, 2018c] or on approximation algorithms for additive [Cole and Gkatzelis, 2018; Cole *et al.*, 2017; Barman *et al.*, 2018b] or even more general valuation functions [Garg *et al.*, 2020].

2 Preliminaries and Notation

We consider fair division instances $I = (N, M, (v_i)_{i \in N})$ in which there is a set N of n agents and a set M of m indivisible goods. Each agent $i \in N$ has a valuation function $v_i : M \rightarrow \mathbb{R}_{\geq 0}$ assigning a non-negative value $v_i(g)$ to each good $g \in M$. We assume v_i is *additive*, i.e., $v_i(A) = \sum_{g \in A} v_i(g)$ for every set (or *bundle*) of goods $A \subseteq M$.

We pay particular attention to the following subclasses of additive valuation functions: (1) *Binary*: $v_i(g) \in \{0, 1\}$ for every $i \in N$ and $g \in M$, and (2) *k-value*: there is a set V

consisting of $|V| = k$ distinct, non-negative real values such that $v_i(g) \in V$ for every $i \in N$ and $g \in M$. Of course, any binary instance is a 2-value instance with $V = \{0, 1\}$ but we distinguish between these cases for technical reasons that will become apparent later on.

A *complete allocation* (or just *allocation*) $\mathbf{A} = (A_i)_{i \in N}$ is a vector listing the bundle A_i of goods that each agent i receives, such that $A_i \cap A_j = \emptyset$ for every $i, j \in N$, and $\cup_{i \in N} A_i = M$. Our goal is to come up with allocations that are considered to be *fair* by all agents. We begin by defining envy-freeness and its additive relaxations.

Definition 2.1. An allocation $\mathbf{A} = (A_i)_{i \in N}$ is

- *envy-free* (EF) if $v_i(A_i) \geq v_i(A_j)$ for every $i, j \in N$;
- *envy-free up to one good* (EF1) if for every $i, j \in N$ with $A_j \neq \emptyset$ there exists a good $g \in A_j$, such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$;
- *envy-free up to any (positively-valued) good* (EFX) if for every $i, j \in N$ and every good $g \in A_j$ for which $v_i(g) > 0$, we have $v_i(A_i) \geq v_i(A_j \setminus \{g\})$;
- *envy-free up to any good* (EFX₀) if for every $i, j \in N$ and every good $g \in A_j$, we have $v_i(A_i) \geq v_i(A_j \setminus \{g\})$.

By definition, we have $\text{EF} \Rightarrow \text{EFX}_0 \Rightarrow \text{EFX} \Rightarrow \text{EF1}$, but no implication works in the opposite direction. For brevity, we say that agent i is \mathcal{E} towards agent j when the criterion of $\mathcal{E} \in \{\text{EF}, \text{EF1}, \text{EFX}, \text{EFX}_0\}$ is true for the ordered pair (i, j) .

Definition 2.2. The *Nash welfare* of an allocation $\mathbf{A} = (A_i)_{i \in N}$ is the product² of the values of the agents for their bundles: $\text{NW}(\mathbf{A}) = \prod_{i \in N} v_i(A_i)$.

We will usually denote by \mathbf{A}^* one of the allocations that maximize the Nash welfare (MNW). Among all such allocations, we will sometimes select \mathbf{A}^* so that some additional properties are satisfied; e.g., see the discussion in Section 3. Caragiannis *et al.* [2019b] showed that $\text{MNW} \Rightarrow \text{EF1}$, but the exact connection between MNW and the variants of EFX is not well-understood.

Before we dive into our main technical results, we show a somewhat surprising connection between EFX and EFX₀. In particular, assuming agents with k -value valuation functions, for any $k \in \mathbb{N}$, the question of finding an EFX₀ allocation reduces to finding an EFX allocation for an instance with only slightly perturbed valuation functions. An immediate corollary is that the existence (resp. the efficient computation) of EFX allocations for additive agents implies the existence (resp. the efficient computation) of EFX₀ allocations; the converse statements are obvious.

Proposition 2.1. *Let $k \in \mathbb{N}$. The problem of computing EFX₀ allocations for k -value instances reduces to the problem of computing EFX allocations for k -value instances. When all values are rational numbers, this reduction requires only polynomial time.*

²As mentioned in the Introduction, the Nash welfare is usually defined as the geometric mean of the values rather than their product. As the allocations (exactly) maximizing the Nash welfare are the same under both definitions, we use the product for simplicity.

3 When Does MNW Imply EFX?

In this section we focus on allocations that maximize the Nash welfare. We first identify the subclasses of valuation functions for which the MNW allocations are always EFX₀, and then consider computational complexity questions. Our main result here is that for all 2-value instances (including binary) any MNW allocation is also EFX₀. Moreover, this result is tight: there exist 3-value instances for which this implication is no longer true.

Throughout this section, we assume that for every good g , there exists at least one agent i with $v_i(g) > 0$, but our results hold without this assumption.

When we talk about MNW allocations, the standard interpretation would be to include *all* complete allocations which achieve the maximum Nash welfare for an instance. When it is possible to achieve positive Nash welfare this is indeed true. However, for the extreme case of instances where all allocations have zero Nash welfare we are going to need a refinement. Following the work of Caragiannis *et al.* [2019b], for instances with maximum Nash welfare equal to zero, we call an allocation an MNW allocation if it (1) maximizes the number of agents with positive value, and then (2) maximizes the product of the values of such agents. This restriction is necessary because when all allocations have zero Nash welfare, the idea of maximizing it clearly fails to distinguish “good” allocations in any sense. To illustrate this, consider the next instance:

	g_1	g_2	g_3
agent 1	1	0	0
agent 2	1	0	0
agent 3	0	1	1

Since the first two agents only like g_1 , the Nash welfare of any allocation is 0. However, not all allocations are EFX₀. The allocation $\{\emptyset, \emptyset, \{g_1, g_2, g_3\}\}$ is clearly not EFX₀ since the first two agents envy agent 3 even after the removal of either g_2 or g_3 . Even an allocation such as $\{\{g_1, g_2\}, \emptyset, \{g_3\}\}$, which maximizes the number of agents with positive value, is not EFX₀ since agent 2 envies agent 1 even after the removal of g_2 . On the other hand, the allocation $\{\{g_1\}, \emptyset, \{g_2, g_3\}\}$, which maximizes the number of agents with positive value *as well as* the product of their values, is indeed EFX₀: the envy of agent 2 towards agent 1 is eliminated by the removal of g_1 .

It should be noted that both (1) and (2) are by default true for MNW allocations in instances with positive Nash welfare. We next show that maximizing the Nash welfare subject to (1) and (2) yields an EFX₀ allocation for any 2-value instance.

Theorem 3.1. *Each MNW allocation is an EFX₀ allocation, for any 2-value instance.*

Proof. Due to space constraints, we prove the statement for two values $\{a, b\}$ such that $a > b > 0$. Then EFX₀ coincides with EFX. The proof for when $b = 0$ (essentially the binary case) is of similar flavor, albeit simpler, despite having to consider instances with zero Nash welfare.

Let i and j be any two agents who are given the sets of goods A_i and A_j in an MNW allocation \mathbf{A}^* . We say that a good is of *type* T_{xy} if i and j have values $v_i(g) = x$ and $v_j(g) = y$ for good g , respectively; so there are four different

types of goods: T_{aa} , T_{ab} , T_{ba} and T_{bb} . If $\min_{g \in A_j} v_i(g) = a$ or $\max_{g \in A_j} v_i(g) = b$, then i is EFX towards j since i is EF1 towards j [Caragiannis *et al.*, 2019b] and the two notions coincide in this case for the pair (i, j) . Therefore, from now on, we will assume that $\min_{g \in A_j} v_i(g) = b$ and $\max_{g \in A_j} v_i(g) = a$, which implies that $|A_j| \geq 2$ and A_j includes at least one good of type T_{ba} or T_{bb} .

Case I: There is at least one good of type T_{bb} in A_j .

Subcase (a): A_j does not include any good of type T_{ab} . Assume, towards a contradiction, that i is not EFX towards j : $v_i(A_i) < v_i(A_j) - b$. Since $v_j(g) \geq v_i(g)$ for all $g \in A_j$, we have that $v_j(A_j) \geq v_i(A_j)$. We now define a new allocation by moving a good $h \in A_j$ of type T_{bb} from j to i . In this new allocation, the product of the values of i and j is

$$\begin{aligned} & (v_i(A_i) + b)(v_j(A_j) - b) \\ &= v_i(A_i)v_j(A_j) + b(v_j(A_j) - v_i(A_i) - b) \\ &\geq v_i(A_i)v_j(A_j) + b(v_i(A_j) - v_i(A_i) - b) \\ &> v_i(A_i)v_j(A_j). \end{aligned}$$

Since the allocation of all other agents has not been changed, the new allocation achieves a strictly larger Nash welfare than \mathbf{A}^* , yielding a contradiction.

Subcase (b): A_j includes at least one good g of type T_{ab} . We will argue about the structure of set A_i . If A_i includes any good x of type T_{aa} , T_{ba} or T_{bb} , then by exchanging g with x , we obtain an allocation with strictly higher Nash welfare, contradicting the choice of \mathbf{A}^* . For example, if x is of type T_{aa} , then in the new allocation (after swapping x and g) agent i has exactly the same value, but agent j 's value has strictly increased by an amount $a - b > 0$. One can verify that the same holds for the other two types. Hence, A_i must include only goods of type T_{ab} , which implies that $v_i(A_i) = |A_i|a$.

Towards a contradiction, assume that i is not EFX towards j . If $|A_j| \leq |A_i| + 1$, since A_j includes some good h for which $v_i(h) = b$, we have that

$$v_i(A_j) \leq (|A_j| - 1)a + b \leq |A_i|a + b = v_i(A_i) + b,$$

i.e., agent i is EFX towards j . So, it must be $|A_j| \geq |A_i| + 2$. We create a new allocation by moving a good $g \in T_{ab}$ from j to i . The product of the values of i and j then becomes

$$\begin{aligned} & (v_i(A_i) + a)(v_j(A_j) - b) \\ &= v_i(A_i)v_j(A_j) + av_j(A_j) - bv_i(A_i) - ab. \end{aligned}$$

Since $v_j(A_j) \geq |A_j|b \geq (|A_i| + 2)b$ and $v_i(A_i) = |A_i|a$, we have that

$$\begin{aligned} & av_j(A_j) - bv_i(A_i) - ab \\ &\geq (|A_i| + 2)ab - |A_i|ab - ab = ab > 0. \end{aligned}$$

Since the bundles of the other agents have not been changed, we have that the new allocation has strictly larger Nash welfare than \mathbf{A}^* , contradicting its choice.

Case II: There are no goods of type T_{bb} in A_j .

Then A_j includes at least one good of type T_{ba} . If A_j includes at least one good of type T_{ab} , then, as we argued in Case I(b) above, in order for \mathbf{A}^* to be an MNW allocation, A_i cannot include any goods of type T_{aa} , T_{ba} or T_{bb} . As a result,

A_i includes only goods of type T_{ab} and by reproducing the analysis used in Case I(b) it follows that \mathbf{A}^* is EFX.

So, we may assume that A_j includes goods of type T_{ba} and T_{aa} only. This implies that $v_j(A_j) = |A_j|a$. Assume towards a contradiction that i is not EFX towards j : $v_i(A_i) < v_i(A_j) - b$. Since A_j contains at least one good that i values as b , we also have that $v_i(A_j) \leq (|A_j| - 1)a + b$. Combining the last two expressions, we obtain that

$$v_i(A_i) + a < |A_j|a = v_j(A_j).$$

Now, consider the allocation that is obtained from \mathbf{A}^* by moving a good of type T_{aa} from j to i . We know that such an item exists since $\max_{g \in A_j} v_i(g) = a$. By using the last inequality, the product of the values of i and j in the new allocation is

$$\begin{aligned} & (v_i(A_i) + a)(v_j(A_j) - a) \\ &= v_i(A_i)v_j(A_j) + a(v_j(A_j) - v_i(A_i) - a) \\ &> v_i(A_i)v_j(A_j), \end{aligned}$$

which combined with the fact that the bundles of the other agents have not been changed, contradicts the choice of \mathbf{A}^* .

In any case, we conclude that \mathbf{A}^* must be EFX. \square

Caragiannis *et al.* [2019b] presented a 3-value instance in which no MNW allocation is EFX. For completeness, we include here a simpler such instance. Let ε be a small positive constant and consider an instance with two agents and three goods with values as shown in the table:

	g_1	g_2	g_3
agent 1	$1 - \varepsilon$	1	$1 + \varepsilon$
agent 2	1	$1 - \varepsilon$	$1 + \varepsilon$

This is a 3-value instance with values $\{1 - \varepsilon, 1, 1 + \varepsilon\}$. It is easy to verify that there are exactly two allocations achieving the maximum Nash welfare of $2 + \varepsilon$: $\mathbf{A}_1 = (\{g_2\}, \{g_1, g_3\})$ and $\mathbf{A}_2 = (\{g_2, g_3\}, \{g_1\})$. The Nash welfare of any other allocation is either $2(1 - \varepsilon)$ or $2 + \varepsilon - \varepsilon^2$. Now, for $\ell \in \{1, 2\}$, observe that in \mathbf{A}_ℓ agent ℓ is not EFX towards the other agent since she envies her even after the removal of g_ℓ .

3.1 Complexity of Maximizing the Nash Welfare

We now turn our attention to the complexity of computing a maximum Nash welfare allocation. This problem is already known to be hard for many domain restrictions, and easy for only a few special cases. Nevertheless, its complexity for k -value instances with $k \in \{2, 3, 4\}$ has been open. Here we make significant progress towards settling these cases.

For binary instances, Barman *et al.* [2018c] gave the greedy algorithm ALG-BINARY that outputs an allocation maximizing the Nash welfare. If the Nash welfare is positive, then the outputted allocation is also EFX₀ by Theorem 3.1. However, if the Nash welfare is zero, ALG-BINARY might *not* output an MNW allocation in our sense, i.e., an allocation that maximizes the number of agents with positive values and their product (in that order). To circumvent this, we define a bipartite graph consisting of nodes corresponding to the agents on the left and nodes corresponding to the goods on the right, while an edge between an agent and a good exists if and only

if the agent has value 1 for the good. By computing a maximum bipartite matching on this graph, it is guaranteed that the number of agents with positive value is maximized. Then, we run ALG-BINARY on the restricted instance where the set of agents includes only the ones that participate in the maximum matching, so that the product of their values (which now is going to be positive) is also maximized. This leads to the following statement.

Theorem 3.2. *For binary instances, computing an MNW allocation (and thus an EFX₀ allocation) can be done in polynomial time.*

For general 2-value instances we were unable to resolve the complexity of computing an MNW allocation, but we show that the problem is NP-hard for 3-value instances. This extends the hardness aspect (but not the inapproximability) of the result of Lee [2017] for 5-value instances.

Theorem 3.3. *Computing an MNW allocation is NP-hard, even for 3-value instances.*

4 EFX Allocations for 2-Value Instances

Even though we showed that any MNW allocation is also EFX₀ for agents with 2-value valuation functions, it remains an open question whether there exists a polynomial-time algorithm for computing such allocations beyond the binary case. In this section, we try to circumvent this and aim to design an efficient algorithm for computing EFX₀ allocations (which might not maximize the Nash welfare) for 2-value instances. We assume values $\{a, b\}$, such that $a > b \geq 0$.

Our algorithm, which we call MATCH&FREEZE, proceeds in rounds and maintains a set of *active* agents L , initially containing everyone. In each round, every active agent is given exactly one of the remaining goods, with the possible exception of the last round in which there might not be enough goods left for all agents. The algorithm terminates when all goods have been allocated.

To determine which good each active agent gets during a round, we create a bipartite graph $G = (L \cup R, E)$ with nodes corresponding to the active agents L on one side and to the remaining goods R on the other. An edge between an active agent i and a good g exists if and only if $v_i(g) = a$. We first compute a maximum matching on this graph. Then each agent gets the good to which she is matched. If there are agents who are not matched to any good and there are still available goods, the unmatched agents receive one arbitrary available good each (subject to availability).

There are two possible reasons why an agent i is not matched to any good in a round: (1) she does not have value a for any good (only b), or (2) the maximum matching is such that all goods for which her value is a are given to other agents. Case (1) does not affect whether the final allocation will be EFX₀, but case (2) is crucial. This is because agent i might now have much smaller value for her own bundle compared to her value for the bundles of some agents that *just received* one good each that i values as a . Let Z be the set of these agents. To make up the distance, agent i should possibly receive multiple goods of value b while all agents in Z must *freeze* for a number of subsequent rounds depending on the ratio a/b .

Algorithm 1 MATCH&FREEZE($N, M, (v_i)_{i \in N}$)

```

1: Input: a 2-value instance using the values  $a, b$  ( $a > b \geq 0$ )
2:  $L \leftarrow N$  ▷ set of active agents
3:  $R \leftarrow M$  ▷ set of unallocated goods
4:  $\ell = (1, 2, \dots, n)$  ▷ ordered list of agents
5: while  $R \neq \emptyset$  do ▷ every iteration is a round
6:   Construct the bipartite graph  $G = (L \cup R, E)$ .
7:   Compute a maximum matching on  $G$ .
8:   for each matched pair  $(i, g)$  do
9:     Allocate good  $g$  to agent  $i$ .
10:    Remove  $g$  from  $R$ .
11:  for each unmatched active agent  $i$  w.r.t.  $\ell$  do
12:    Allocate one arbitrary unallocated good  $g$  to  $i$ .
13:    Remove  $g$  from  $R$ .
14:  Construct the set  $F$  of agents that need to freeze.
15:  Remove agents of  $F$  from  $L$  for the next  $\lfloor a/b - 1 \rfloor$  rounds.
16:  Put agents of  $F$  to the end of  $\ell$ .
17: return the resulting allocation  $\mathbf{A}$ .
```

We define the set F of agents that need to freeze at the end of round r to consist of all those agents who must become inactive because they have obtained too much value from the perspective of other agents (similarly to Case (2) above). Formally, for every active agent i , let g_i be the good she gets in round r . We begin by setting $F = \{i \in L \mid \exists j \in L : v_j(g_i) = a, v_j(g_j) = b\}$. Then, iteratively, as long as there is an agent $i \in L \setminus F$ such that there exists $j \in F$ with $v_j(g_i) = a$, we also add i to F . Each agent in F will remain *frozen* for the next $\lfloor a/b - 1 \rfloor$ rounds. In the case $b = 0$, we use the interpretation $\lfloor a/b - 1 \rfloor = +\infty$ in which case agents in F remain frozen forever. Exploiting the properties of the maximum matchings used to allocate goods we can prove that no agent in F will become envious while frozen, and that F is always a strict subset of L . The latter means that there is at least one (non-frozen) active agent at any time, and thus the algorithm will terminate after at most m rounds.

Theorem 4.1. *MATCH&FREEZE computes an EFX₀ allocation in polynomial time, for any 2-value instance.*

Proof sketch. For any $i \in N$, let r_i be the round in which the last goods of value a for agent i were allocated. Since the goods in every round are allocated by computing a maximum matching, we can show that for each agent $i \in N$:

- She received a good for which she has value a in each of the rounds $1, 2, \dots, r_i - 1$.
- She can freeze only at the end of round r_i , and only if during r_i she received a good for which she has value a .
- She can freeze at most once. After freezing, she has value b for each of the remaining goods.

First, let us see why the algorithm terminates in polynomial time. If no agent ever becomes frozen, then the algorithm terminates after at most $\lceil m/n \rceil$ rounds. Otherwise, let r be the first round at the end of which an agent becomes frozen, i.e., there exists some agent j with $v_j(g_j) = b$ and some agent i with $v_j(g_i) = a$. By the definition of r_j and the above observations, $r = r_j$. Since agent j did not become frozen in round $r = r_j$ (because $v_j(g_j) = b$), this means that she will

never freeze. Since agent j gets a good in every round, the algorithm terminates after at most m rounds.

Now, it remains to show that the final allocation is EFX_0 . This follows by a careful case analysis and the following observation: If at the beginning of some round $r > r_i$, an active agent i is envy-free towards agent j , then agent i will be EFX_0 towards agent j at the end of the algorithm. Due to space constraints, we omit the details. \square

5 MNW and the EFX-Value

As we saw in Section 3, maximizing the Nash welfare does not yield an EFX allocation in general. Here we take a different route and instead of considering exact EFX allocations, we focus on approximation. We start by showcasing that maximizing the Nash welfare does not guarantee *any* meaningful approximation of EFX according to the current definition of approximation used in the literature [Plaut and Roughgarden, 2018; Chan *et al.*, 2019; Amanatidis *et al.*, 2019].

Definition 5.1 (α -EFX allocation). For $\alpha \in (0, 1]$, an allocation \mathbf{A} is α -EFX if for every pair $i, j \in N$ and every good $g \in A_j$ such that $v_i(g) > 0$, it holds that $v_i(A_i) \geq \alpha v_i(A_j \setminus \{g\})$. Let $w > 1$ and $\varepsilon < \frac{1}{2w}$. Consider the following very simple instance with two agents, three goods, and values:

	g_1	g_2	g_3
agent 1	w	0	$1/2$
agent 2	w	1	ε

We first claim that the allocation $\mathbf{A}^* = (A_1 = \{g_1, g_3\}, A_2 = \{g_2\})$ is the only one that achieves the maximum Nash welfare of $w+1/2$. Indeed, the Nash welfare of any allocation that gives g_2 to agent 1 can only increase by moving g_2 to agent 2, while any allocation other than \mathbf{A}^* that gives g_2 to agent 2 has Nash welfare either $(w+1)/2$ or $w+\varepsilon w < w+1/2$. Notice, however, that \mathbf{A}^* is not EFX since $v_2(A_2) = 1 < w = v_2(A_1 \setminus \{g_3\})$. Instead, it is only $1/w$ -EFX, an approximation factor that can be arbitrarily close to zero as w becomes large.

Nevertheless, \mathbf{A}^* is not that far away from being an EFX allocation! To see this, consider the allocation $\mathbf{B} = (B_1 = \{g_1\}, B_2 = \{g_2, g_3\})$ that is obtained from \mathbf{A}^* by only moving g_3 from agent 1 to agent 2. Clearly, agent 2 is EFX towards agent 1, as the latter gets only one good. Moreover, the value agent 2 has now is $v_2(B_2) = 1 + \varepsilon$, which is extremely close to the value $v_2(A_2) = 1$ that agent 2 has in \mathbf{A}^* . So, even though \mathbf{A}^* is $1/w$ -EFX because $v_2(A_2)$ is very low compared to $v_2(A_1 \setminus \{g_3\})$, $v_2(A_2)$ is actually very close to the value she would have in a *nearby* EFX allocation. Consequently, if we accept that agent 2 considers the EFX allocation \mathbf{B} as fair, then she should also consider \mathbf{A}^* as being *almost* fair.

We say that the value $v_2(B_2) = 1 + \varepsilon$ is the EFX-value that agent 2 can achieve by augmenting her bundle with a subset of goods from agent 1 in order to create the closest-to- \mathbf{A}^* (in terms of value) allocation \mathbf{B} which she considers as EFX. Then, since $v_2(A_2) = \frac{1}{1+\varepsilon} v_2(B_2)$, agent 2 achieves an approximation of $\frac{1}{1+\varepsilon}$ of her EFX-value. Let us now formalize these notions for any number of agents.

Definition 5.2 (EFX-value). Let $\mathbf{A} = (A_1, \dots, A_n)$ be an allocation. For every pair of agents $i, j \in N$, let $X_{ij} \subseteq A_j$ be

a set of goods such that $v_i(A_i \cup X_{ij}) \geq v_i(A_j \setminus (X_{ij} \cup \{g\}))$ for every $g \in A_j \setminus X_{ij}$ and $v_i(A_i \cup X_{ij})$ is minimized. Then, the EFX-value of agent i is

$$\chi_i(\mathbf{A}) = \max_{j \in N \setminus \{i\}} v_i(A_i \cup X_{ij}).$$

Definition 5.3 (α -vEFX allocation). For $\alpha \in (0, 1]$, an allocation \mathbf{A} is α -vEFX if $v_i(A_i) \geq \alpha \chi_i(\mathbf{A})$ for every $i \in N$.

We remark that using EFX_0 instead of EFX in the above definitions does not make any difference since adding zeros does not affect the EFX-value. Furthermore, observe that a 1-vEFX allocation is an EFX allocation but not necessarily an EFX_0 allocation.

Our first result in this section illustrates the connection between approximate EFX and vEFX allocations.

Theorem 5.1. *An α -EFX allocation is also an $\frac{\alpha}{1+\alpha}$ -vEFX allocation and this guarantee is tight. On the other hand, an α -vEFX allocation is not guaranteed to be β -EFX, for any $\alpha, \beta \in (0, 1)$.*

Even though maximizing the Nash welfare may not yield a β -EFX for any $\beta \in (0, 1)$ as we showed above, it is guaranteed to produce a $1/2$ -vEFX allocation!

Theorem 5.2. *Any maximum Nash welfare allocation \mathbf{A}^* is a $1/2$ -vEFX allocation and this factor is tight.*

6 Directions for Future Work

We studied the connection between two celebrated notions, that of maximum Nash welfare and envy-freeness up to any good. The first question that our work leaves open is whether it is possible to compute in polynomial time an allocation that maximizes the Nash welfare for 2-value instances. Another direction is to investigate whether the algorithmic ideas presented in Section 5 for 2-value instances can be extended to computing EFX_0 allocations for k -value instances, for any $k \geq 3$. Finally, going beyond exact MNW or EFX allocations, a natural question is whether one can design polynomial-time algorithms with constant approximation guarantees for both the EFX-value and the Nash welfare.

Acknowledgements

We would like to thank Ioannis Caragiannis for fruitful discussions. G. Amanatidis was partially supported by the ERC Advanced Grant AMDROMA (No. 788893), the MIUR PRIN project ALGADIMAR and the NWO Veni project VI.Veni.192.153. G. Birmpas and A. Voudouris were partially supported by the ERC Starting Grant 639945 ACCORD “Algorithms for Complex Collective Decisions on Structured Domains”. A. Hollender was partially supported by an EP-SRC doctoral studentship (Reference 1892947).

References

[Aleksandrov and Walsh, 2019] M. Aleksandrov and T. Walsh. Greedy algorithms for fair division of mixed manna. *arXiv:1911.11005*, 2019.

- [Amanatidis *et al.*, 2017] G. Amanatidis, E. Markakis, A. Nikzad, and A. Saberi. Approximation algorithms for computing maximin share allocations. *ACM Trans. Algorithms*, 13(4):52:1–52:28, 2017.
- [Amanatidis *et al.*, 2019] G. Amanatidis, E. Markakis, and A. Ntokos. Multiple birds with one stone: Beating $1/2$ for EFX and GMMS via envy cycle elimination. *arXiv:1909.07650*, 2019.
- [Babaioff *et al.*, 2020] M. Babaioff, T. Ezra, and U. Feige. Fair and truthful mechanisms for dichotomous valuations. *arXiv:2002.10704*, 2020.
- [Barman and Krishnamurthy, 2017] S. Barman and S. K. Krishnamurthy. Approximation algorithms for maximin fair division. In *Proceedings of ACM EC 2017*, pages 647–664, 2017.
- [Barman *et al.*, 2018a] S. Barman, A. Biswas, S. K. Krishnamurthy, and Y. Narahari. Groupwise maximin fair allocation of indivisible goods. In *Proceedings of AAI 2018*, pages 917–924, 2018.
- [Barman *et al.*, 2018b] S. Barman, S. K. Krishnamurthy, and R. Vaish. Finding fair and efficient allocations. In *Proceedings of ACM EC 2018*, pages 557–574, 2018.
- [Barman *et al.*, 2018c] S. Barman, S. K. Krishnamurthy, and R. Vaish. Greedy algorithms for maximizing Nash social welfare. In *Proceedings of AAMAS 2018*, pages 7–13, 2018.
- [Bouveret *et al.*, 2016] S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair allocation of indivisible goods. In *Handbook of Computational Social Choice*, pages 284–310. Cambridge University Press, 2016.
- [Budish, 2011] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2019a] I. Caragiannis, N. Gravin, and X. Huang. Envy-freeness up to any item with high Nash welfare: The virtue of donating items. In *Proceedings of ACM EC 2019*, pages 527–545, 2019.
- [Caragiannis *et al.*, 2019b] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. *ACM Trans. Economics and Comput.*, 7(3):12:1–12:32, 2019.
- [Chan *et al.*, 2019] H. Chan, J. Chen, B. Li, and X. Wu. Maximin-aware allocations of indivisible goods. In *Proceedings of IJCAI 2019*, pages 137–143, 2019.
- [Chaudhury *et al.*, 2020a] B. R. Chaudhury, J. Garg, and K. Mehlhorn. EFX exists for three agents. *arXiv:2002.05119*, 2020.
- [Chaudhury *et al.*, 2020b] B. R. Chaudhury, T. Kavitha, K. Mehlhorn, and A. Sgouritsa. A little charity guarantees almost envy-freeness. In *Proceedings of ACM-SIAM SODA 2020*, pages 2658–2672, 2020.
- [Cole and Gkatzelis, 2018] R. Cole and V. Gkatzelis. Approximating the Nash social welfare with indivisible items. *SIAM J. Comput.*, 47(3):1211–1236, 2018.
- [Cole *et al.*, 2017] R. Cole, N. R. Devanur, V. Gkatzelis, K. Jain, T. Mai, V. V. Vazirani, and S. Yazdanbod. Convex program duality, Fisher markets, and Nash social welfare. In *Proceedings of ACM EC 2017*, pages 459–460, 2017.
- [Gamow and Stern, 1958] G. Gamow and M. Stern. *Puzzle-Math*. Viking press, 1958.
- [Garg and McGlaughlin, 2019] J. Garg and P. McGlaughlin. Improving Nash social welfare approximations. In *Proceedings of IJCAI 2019*, pages 294–300, 2019.
- [Garg and Taki, 2019] J. Garg and S. Taki. An improved approximation algorithm for maximin shares. *arXiv:1903.00029*, 2019.
- [Garg *et al.*, 2019] J. Garg, P. McGlaughlin, and S. Taki. Approximating maximin share allocations. In *Proceedings of SOSA@SODA 2019*, volume 69 of *OASICS*, pages 20:1–20:11, 2019.
- [Garg *et al.*, 2020] J. Garg, P. Kulkarni, and R. Kulkarni. Approximating Nash social welfare under submodular valuations through (un)matchings. In *Proceedings of ACM-SIAM SODA 2020*, pages 2673–2687, 2020.
- [Ghodsí *et al.*, 2018] M. Ghodsí, M. T. Hajiaghayi, M. Seddighin, S. Seddighin, and H. Yami. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of ACM EC 2018*, pages 539–556, 2018.
- [Kurokawa *et al.*, 2018] D. Kurokawa, A. D. Procaccia, and J. Wang. Fair enough: Guaranteeing approximate maximin shares. *J. ACM*, 65(2):8:1–8:27, 2018.
- [Kyropoulou *et al.*, 2019] M. Kyropoulou, W. Suksompong, and A. A. Voudouris. Almost envy-freeness in group resource allocation. In *Proceedings of IJCAI 2019*, pages 400–406, 2019.
- [Lee, 2017] E. Lee. APX-hardness of maximizing Nash social welfare with indivisible items. *Information Processing Letters*, 122:17–20, 2017.
- [Lipton *et al.*, 2004] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of ACM EC 2004*, pages 125–131, 2004.
- [Plaut and Roughgarden, 2018] B. Plaut and T. Roughgarden. Almost envy-freeness with general valuations. In *Proceedings of ACM-SIAM SODA 2018*, pages 2584–2603, 2018.
- [Procaccia, 2016] A. D. Procaccia. Cake cutting algorithms. In *Handbook of Computational Social Choice*, pages 311–330. Cambridge University Press, 2016.
- [Steinhaus, 1948] H. Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
- [Suksompong, 2018] W. Suksompong. Approximate maximin shares for groups of agents. *Mathematical Social Sciences*, 92:40–47, 2018.
- [Varian, 1974] H. R. Varian. Equity, envy and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.