

Almost Envy-freeness, Envy-rank, and Nash Social Welfare Matchings

Alireza Farhadi,¹ MohammadTaghi Hajiaghayi,¹ Mohamad Latifian,² Masoud Seddighin,³ Hadi Yami⁴

¹ University of Maryland

² University of Toronto

³ Institute for Research in Fundamental Sciences (IPM)

⁴ Microsoft

farhadi@cs.umd.edu, Hajiagha@cs.umd.edu, latifian@cs.toronto.edu, seddighin.masood@gmail.com, hayami@microsoft.com

Abstract

Envy-freeness up to one good (EF1) and envy-freeness up to any good (EFX) are two well-known extensions of envy-freeness for the case of indivisible items. It is shown that EF1 can always be guaranteed for agents with subadditive valuations (Lipton et al. 2004). In sharp contrast, it is unknown whether or not an EFX allocation always exists, even for four agents and additive valuations. In addition, the best approximation guarantee for EFX is $(\phi - 1) \simeq 0.61$ by Amanatidis et al. (Amanatidis, Markakis, and Ntokos 2020).

In order to find a middle ground to bridge this gap, in this paper we suggest another fairness criterion, namely *envy-freeness up to a random good* or EFR, which is weaker than EFX, yet stronger than EF1. For this notion, we provide a polynomial-time 0.73-approximation allocation algorithm. For our algorithm we use Nash Social Welfare Matching which makes a new connection between Nash Social Welfare and envy freeness.

Introduction

Fair division is a fundamental and interdisciplinary problem that has been extensively studied in economics, mathematics, political science, and computer science. Generally, the goal is to find an allocation of a resource to n agents, which is agreeable to all the agents according to their preferences. The first formal treatment of this problem was in 1948 by Steinhaus (Steinhaus 1948). Following his work, a vast literature has been developed and several notions for measuring fairness have been suggested (Steinhaus 1948; Foley 1967; Budish 2011; Lipton et al. 2004; Caragiannis et al. 2019). One of the most prominent and well-established fairness notions, introduced by Foley (Foley 1967), is envy-freeness, which requires that each agent prefers his share over that of any other agent.

Traditionally, envy-freeness has been studied for both divisible and indivisible resources. When the resource is a single heterogeneous divisible item (i.e. can be fractionally allocated), envy-freeness admits strong theoretical guarantees. However, beyond divisibility, when dealing with a set of indivisible goods, envy-freeness is too strong to be attained; for example, for two agents and a single indivisible good,

the agent that receives no good envies the other party. Therefore, several relaxations of envy-freeness are introduced for the case of indivisible items (Lipton et al. 2004; Budish 2011; Caragiannis et al. 2019). One of these relaxations, suggested by Budish (Budish 2011), is *envy-freeness up to one good* (EF1)¹. An allocation of indivisible goods is EF1 if any possible envy of an agent for the share of another can be resolved by removing some good from the envied share. In contrast to envy-freeness, EF1 allocation always exists. Indeed, a simple round-robin algorithm always guarantees EF1 for additive valuations, and a standard envy-graph based allocation guarantees EF1 for more general (sub-additive) valuations. Furthermore, it is shown that any Nash welfare maximizing allocation (allocation that maximizes the product of the agents' utilities) is both Pareto efficient and EF1.

Recently, Caragiannis *et al.* (Caragiannis et al. 2019) suggested another intriguing relaxation of envy-freeness, namely *envy-free up to any good* (EFX), which attracted great deal of attention. An allocation is said to be EFX, if no agent envies another agent after the removal of any item from the other agent's bundle. Theoretically, this notion is strictly stronger than EF1 and is strictly weaker than EF. In contrast to EF1, questions related to the EFX notion are relatively unexplored. As an example, despite significant effort (Caragiannis et al. 2019), the existence of such allocations is still unknown. The most impressive breakthrough in this area is the recent work of Chaudhury, Garg, and Mehlhorn (Chaudhury, Garg, and Mehlhorn 2020), which shows that for the case of 3 agents with additive valuations an EFX allocation always exists. Furthermore, unlike EF1, Nash social welfare maximizing allocations are not necessarily EFX (Caragiannis et al. 2019).

Given this impenetrability of EFX, a growing strand of research started considering its relaxations. For example, Plaut and Roughgarden (Plaut and Roughgarden 2020), consider an approximate version of EFX² and provide a $1/2$

¹It is worth mentioning that before the work of Budish (Budish 2011) EF1 was implicitly addressed by Lipton et al (Lipton et al. 2004).

²An allocation is α -approximate EFX, if for every pair of agents i and j , agent i believes that the share allocated to him is worth at least α fraction of the share allocated to agent j , after removal of agent j 's least valued item (according to agent i 's preference).

approximation solution for agents with sub-additive valuation functions. For additive valuations, this factor has been recently improved to 0.618 by Amanatidis *et al.* (Amanatidis, Markakis, and Ntokos 2020). Another interesting relaxation is EFX-with-charity. Such allocations donate a bundle of items to charity and guarantee EFX for the rest of the items. The less valuable the donated items are, the more desirable the allocation is. Caragiannis *et al.* (Caragiannis, Gravin, and Huang 2019) show that there always exists an EFX-with-charity allocation where every agent receives half the value of his bundle in the optimal Nash social welfare allocation. Recently, Chaudhury *et al.* (Chaudhury *et al.* 2020) have proposed an EFX-with-charity allocation such that no agent values the donated items more than his bundle and the number of donated items is less than the number of agents.

Considering the huge discrepancy between EFX and EF1, in this paper we wish to find a middle ground to bridge this gap. We therefore suggest another fairness criterion, namely *envy-freeness up to a random item* or EFR, which is weaker than EFX, yet stronger than EF1. For this notion, we provide a polynomial time 0.73-approximation algorithm, i.e., an algorithm that constructs 0.73-EFR allocations in polynomial time. Our allocation method is based on a special type of matching, namely Nash Social Welfare Matching. In Section , we briefly discuss our techniques to obtain these results.

Our Results and Techniques

Envy-freeness up to a random item. We suggest a new fairness notion, namely *envy-free up to a random good* (EFR). Roughly speaking, in an EFR allocation, no agent i envies another agent j (in expectation), if we remove a random good from the bundle of agent j . In other words, the expected value of agent i for the bundle allocated to agent j , after removing a random item from it is at most as much as the value of his own bundle. Obviously, EFR is a weaker notion than EFX, yet stronger than EF1.

The intuition behind EFR is to use randomness to reduce the *severe impact of small items*. To see what we mean by this term, consider the following scenario: suppose that the value of agent i for his share is 1000. In addition, assume that the bundle allocated to an agent j contains two items, each with value 600 to agent i . Even though the allocation is currently EFX with respect to agent i , allocating even a very small item (say, with value close to 0 to agent i) to agent j violates EFX condition for agent i . This is counter-intuitive in the sense that the last item allocated to agent j was totally worthless to agent i . On the other hand, allocating any item with value less than 300 to agent j preserves EFR condition for agent i . This property makes EFR more flexible, especially when the number of items is not very large. On the other hand, as the number of items allocated to an agent grows larger, we expect EFX and EFR to be more and more aligned.

Similar to EFX, we provide a counter example which shows that a Nash Social Welfare allocation is not necessarily EFR (see Example 4). This separates EFR from EF1 given the fact that a Nash Social welfare allocation is always EF1 (Caragiannis *et al.* 2019). It is worth mentioning




Figure 1: Flowchart of the 0.73-EFR allocation algorithm

that Caragiannis *et al.* (Caragiannis *et al.* 2019) presented an example to show that Nash Social welfare allocation is not necessarily EFX. However, their example is still EFR. The difference between these two examples can be seen as an evidence for the distinction between EFR and EFX.

As noted, the best approximation guarantee for EFX is 0.61 by Amanatidis *et al.* (Amanatidis, Markakis, and Ntokos 2020). Since every EFX allocation is also EFR, this result also provides a 0.61-approximation algorithm for EFR. In this paper, we improve this ratio to 0.73.

Theorem 1. There exists an algorithm that finds a 0.73-EFR allocation. In addition, such an allocation can be found in polynomial time.

In order to prove Theorem 1, we propose a three-step algorithm that finds a 0.73-EFR allocation in polynomial time. Roughly speaking, in the first two steps, we allocate valuable (i.e., large) items while preserving the 0.73-EFR property. Next, we use an envy-cycle based procedure to allocate the rest of the items. Figure 1 shows a flowchart of our method.

The first challenge to address is the method by which we must allocate large items in the first step. Interestingly, we introduce a special type of matching allocation with intriguing properties which makes it ideal for our algorithm. We call such an allocation a *Nash Social Welfare Matching*.

Nash Social Welfare Matching. In the first step of the algorithm, we allocate one item to each agent such that the product of the utilities of the agents is maximized. The idea of Nash social welfare matching has been used in (Cole and Gkatzelis 2018; Nguyen and Rothe 2014; Garg, Kulkarni, and Kulkarni 2020) previously. The interesting fact about this allocation is that, not only does it allocate large items, but it also provides very useful information about the value of the rest of the items. In Section we broadly discuss such allocations and their properties. However, to shed light on their usefulness, assume that after a Nash Social Welfare Matching, agent i envies agent j with a ratio $\alpha > 1$, meaning that he thinks the value of the good allocated to agent j is α times greater than the value of his item. In that case, we can immediately conclude that the item allocated to agent j is α times more valuable to him (agent j) than any remaining item; otherwise, we could improve the utility product by allocating the most valuable remaining item to agent j and giving his former item to agent i (and of course, freeing agent i 's former item). In addition, we can express the same proposition for the value of the item allocated to agent i for agent j : the value of this item for agent j is at most $1/\alpha$ of the item allocated to agent j himself. The above statement can be generalized to the arguments that include more than two agents. With this aim, we introduce several new

concepts, including *envy-ratio graph* (a complete weighted graph that represents the envy-ratios between agents), *improving cycles*, and *envy-rank*.

It is worth mentioning that the main challenge in many fair allocation problems for different fairness criteria (e.g., MMS, EFX) is allocating valuable items. The structure of such matchings makes them ideal for allocating these items. We strongly believe that using Nash Social Welfare matching is not only useful for our algorithm, but can also be seen as a strong tool in the way of finding fair allocations related to the other fairness notions, especially maximin-share. In Appendix C we show how to use NSW matching to obtain an algorithm with the approximation ratio of $(\phi - 1) \simeq 0.61$ for EFX. The approximation ratio of our algorithm matches the state-of-the-art $(\phi - 1)$ approximation result by Amanatidis et al. (Amanatidis, Birmpas, and Markakis 2018).

Related Work

Fair allocation of a divisible resource (known as *cake cutting*) was first introduced by Steinhaus (Steinhaus 1948) in 1948, and has been the subject of intensive studies ever since. We refer the reader to (Brams and Taylor 1996) and (Robertson and Webb 1998) for an overview of different fairness notions and their related results. Proportionality and Envy-freeness are among the most well-established notions for cake cutting. As mentioned, the literature of cake cutting admits strong positive results for these two notions (see (Steinhaus 1948) for details).

Since neither EF nor proportionality or any approximation of these notions can be guaranteed for indivisible goods, several relaxations have been introduced for these two notions in the past decade. These relaxations include EF1 and EFX for envy-freeness and maximin-share (Budish 2011) for proportionality. Nash Social Welfare (NSW) is also another important notion in allocation of indivisible goods which is a good balance between fairness and optimality.

Apart from the results mentioned in the introduction for EFX and EF1, there are other studies related to these notions (Chan et al. 2019; Barman, Krishnamurthy, and Vaish 2018; Barman, KrishnaMurthy, and Vaish 2018; Barman, Krishnamurthy, and Vaish 2018; Caragiannis, Gravin, and Huang 2019; Caragiannis et al. 2019; Chaudhury et al. 2018). In particular, Barman et al. (Barman, Krishnamurthy, and Vaish 2018) propose a pseudo-polynomial time algorithm that finds an EF1 and Pareto efficient allocation. They also show that any EF1 and Pareto efficient allocation approximates Nash Social Welfare with a factor of 1.45. In contrast to EF1, our knowledge of EFX and NSW beyond additive valuations is limited. For EFX, the only positive results for general valuations is the work of Plaut and Roughgarden (Plaut and Roughgarden 2020) which provides a 1/2-EFX allocation. For NSW, Garg et al. (Garg, Kulkarni, and Kulkarni 2020) prove an $O(n \log n)$ approximation guarantee for submodular valuations. Recently this factor is improved to $O(n)$ (Chaudhury, Garg, and Mehta 2020). In a recent paper, Amanatidis et al. (Amanatidis et al. 2020) establish that a maximum Nash welfare allocation is always EFX as long as there are two possible values for the goods. They also prove that this implication is no longer true for

three or more distinct values.

Maximin-share is one of the most well-studied notions in the recent years. In a pioneering study, Kurokawa et al. (Kurokawa, Procaccia, and Wang 2018) provide an approximation algorithm with the factor of 2/3 for maximin-share, which is improved to 3/4 by Ghodsi et al. (Ghodsi et al. 2018). Beyond additivity, Barman et al. (Barman and Krishna Murthy 2017) show that a simple round robin algorithm can guarantee 1/10-MMS for submodular valuations, and Ghodsi et al. provide approximation guarantees for submodular (1/3), XOS (1/5) and subadditive ($1/\log n$) valuations. In addition, several notions are ramified from maximin-share, including weighted maximin-share (WMMS) (Farhadi et al. 2019), pairwise maximin-share (PMMS) (Caragiannis et al. 2019), and groupwise maximin-share (GMMS) (Barman et al. 2018). Several studies consider the relation between these notions and seek to find an allocation that guarantees a subset of them simultaneously. For example, Amanatidis et al. (Amanatidis, Birmpas, and Markakis 2018) investigate the connections between EF1, EFX, maximin share, and pairwise maximin share. They show that any EF1 allocation is also a 1/n-MMS and a 1/2-PMMS allocation. They also prove that any EFX allocation is a 4/7-MMS and a 2/3-PMMS allocation.

Preliminaries and Basic Observations

Fair allocation problem. An instance of a fair allocation problem consists of a set of n agents, a set \mathcal{M} of m goods, and a valuation profile $V = \{v_1, v_2, \dots, v_n\}$. Each v_i is a function of the form $2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ which specifies the preferences of agent $i \in [n]$ over the goods. Throughout the paper, we assume that a valuation function v_i satisfies the following conditions.

- **Normalization:** $v_i(\emptyset) = 0$.
- **Monotonicity:** $v_i(S) \leq v_i(T)$ whenever $S \subseteq T$.
- **Additivity:** $v_i(S) = \sum_{b \in S} v_i(\{b\})$.

An allocation of a set S of goods is an n -partition $\mathcal{A} = \langle \mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \rangle$ of S , where \mathcal{A}_i is the bundle allocated to agent i . Allocation is complete, if $S = \mathcal{M}$ and is partial otherwise. Since we are interested in the allocations that allocate the whole set of items, the final allocation must be complete.

Fairness criteria. Given an instance of the fair division problem and an allocation \mathcal{A} , an agent i envies another agent j , if he strictly prefers \mathcal{A}_j over his bundle \mathcal{A}_i . An allocation is then said to be *envy-free* (EF), if no agent envies another, i.e., for every pair $i, j \in [n]$ of agents we have $v_i(\mathcal{A}_i) \geq v_i(\mathcal{A}_j)$. As mentioned, envy-freeness is too strong to be guaranteed in an allocation of indivisible items. Therefore, two relaxations of this notion are introduced, namely *envy-free up to one good* (EF1) and *envy-free up to any good* (EFX).

Definition 2. An allocation \mathcal{A} is called

- *envy-free up to one good* (EF1) if for all i, j we have $v_i(\mathcal{A}_i) \geq \min_{b \in \mathcal{A}_j} v_i(\mathcal{A}_j \setminus \{b\})$,

- envy-free up to any good (EFX) if for all i, j we have $v_i(\mathcal{A}_i) \geq \max_{b \in \mathcal{A}_j} v_i(\mathcal{A}_j \setminus \{b\})$.

Even though these two notions seem to be somewhat related, there is a huge discrepancy between the current results obtained for them. It is shown that even for instances with general valuations, an EF1 allocation always exists, and can be computed in polynomial time (Lipton et al. 2004). In contrast, whether or not an EFX allocation always exists is still an open problem, even for additive valuations.

In this paper, we introduce another relaxation of envy-freeness, namely *envy-free up to a random good*. Let D_j be a uniform distribution over the items of \mathcal{A}_j that selects each item with probability $1/|\mathcal{A}_j|$.

Definition 3. Allocation \mathcal{A} is envy-free up to a random good (EFR) if for all i, j where $\mathcal{A}_j \neq \emptyset$, we have

$$v_i(\mathcal{A}_i) \geq \mathbb{E}_{b \sim D_j} \left[v_i(\mathcal{A}_j \setminus \{b\}) \right].$$

Clearly, EFR lies in between EFX and EF1: EFX is a stronger notion of fairness than EFR, and EFR is stronger than EF1. In Example 4, we show one structural difference between EF1 and EFR: in contrast to EF1, EFR is not implied by an allocation that maximizes Nash social welfare.

	1	2	3	4	5
v_1	3	3	1	1	1
v_2	5	5	1	4	3

Figure 2: Agents' valuations over items

Example 4. Consider an instance of the fair allocation problem with 5 items and 2 agents with the valuations represented in Figure 2. The unique allocation that maximizes the NSW allocates the first 3 items to the first agent, and the other 2 items to the second agent. Let \mathcal{A} be this allocation. Since there are 3 items in the first agent's bundle, we have

$$\begin{aligned} \mathbb{E}_{b \sim D_1} \left[v_2(\mathcal{A}_1 \setminus \{b\}) \right] &= \frac{1}{3} \cdot (v_2(\mathcal{A}_1 \setminus \{1\}) \\ &\quad + v_2(\mathcal{A}_1 \setminus \{2\}) + v_2(\mathcal{A}_1 \setminus \{3\})) \\ &= \frac{22}{3} \geq v_2(\mathcal{A}_2) = 7. \end{aligned}$$

Hence, this allocation is not EFR.

Finally, approximate versions of EFX and EFR are defined as follows.

Definition 5. For a constant $c \leq 1$, an allocation \mathcal{A} is called

- c -approximate envy-free up to any good (c -EFX), if for all i, j we have

$$v_i(\mathcal{A}_i) \geq c \cdot \max_{b \in \mathcal{A}_j} v_i(\mathcal{A}_j \setminus \{b\}),$$

- c -approximate envy-free up to a random good (c -EFR) if for all i, j we have

$$v_i(\mathcal{A}_i) \geq c \cdot \mathbb{E}_{b \sim D_j} \left[v_i(\mathcal{A}_j \setminus \{b\}) \right].$$

Note that Example 4 also shows that the maximum NSW allocation does not guarantee better than $\frac{21}{22}$ approximation of EFR.

Envy-ratio Graph. Envy-ratio graph is in fact a generalization of the envy-graph introduced by Lipton et al. (Lipton et al. 2004). Suppose that at some stage of our algorithm we have a partial allocation \mathcal{A} . We define a graph called *envy-ratio graph* to be a complete weighted digraph with the following construction: each vertex corresponds to an agent, and for each ordered pair (i, j) , there is a directed edge from vertex i to vertex j with the weight $w_{i,j} = v_i(\mathcal{A}_j)/v_i(\mathcal{A}_i)$.

Assuming each agent has a non-zero value for each good, for every i, j we have $w_{i,j} \in [0, \infty)$. Note that $w_{i,j} \leq 1$ implies that agent i does not envy agent j , whereas $w_{i,j} > 1$ indicates agent i envies agent j . The higher the value of $w_{i,j}$ is, the more envious agent i is to the bundle of agent j . Indeed, the well-known envy-graph is a subgraph of envy-ratio graph containing only the edges with $w_{i,j} > 1$.

Nash Social Welfare (NSW) Matching. Nash social welfare, originally proposed by Nash (Nash Jr 1950), is defined to be the geometric mean of the agents' valuations. An allocation that maximizes Nash social welfare is known to have desirable properties. For example, such allocations are proved to be EF1 and Pareto optimal. Roughly, Nash social welfare maximizing allocations can be seen as a trade-off between the egalitarian and the utilitarian.

In the first step of the algorithm, we allocate one item to each agent such that the Nash social welfare of the agents is maximized. More formally, define Nash Social Welfare matching of $[m]$ to be a partial allocation $\mathcal{A} = \langle \mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \rangle$, such that $\prod_i v_i(\mathcal{A}_i)$ is maximized and for every i we have $|\mathcal{A}_i| = 1$.

Similar to Nash social welfare allocations, Nash social welfare matchings exhibit beautiful properties which greatly help us in designing our algorithm. One simple property of such allocations is shown in Observation 7. Before we state Observation 7, we need to define concepts of *improving* and *strictly improving* cycles.

Definition 6. Let $c = i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_k \rightarrow i_1$ be a cycle in the envy-ratio graph. Then, c is an improving cycle, if

$$w_{i_1, i_2} \times w_{i_2, i_3} \times \dots \times w_{i_{k-1}, i_k} \times w_{i_k, i_1} > 1.$$

Furthermore, we say a cycle c is a strictly improving cycle, if c is an improving cycle and for every $(i \rightarrow j) \in c$, $w_{i,j} > 1$ holds.

We note that strictly improving cycle is an essential concept in all envy-cycle elimination methods (Lipton et al. 2004; Barman and Krishna Murthy 2017; Chaudhury et al. 2020; Amanatidis, Markakis, and Ntokos 2020). These methods typically rotate the shares over strictly improving cycles to enhance social welfare. However, to the best of our knowledge, no previous work made use of improving cycles.

Observation 7. Suppose that we allocate one item to each agent using Nash social welfare matching. Then, the envy-ratio graph admits no improving cycle.

The proof of the mentioned observation is available in Appendix A. A particularly useful case of Observation 7 is for the cycles of length 2, which we state in Corollary 8.

Corollary 8 (of Observation 7). Suppose that for two agents i, j we have $v_i(\mathcal{A}_j) \geq r \cdot v_i(\mathcal{A}_i)$. Then, we have $v_j(\mathcal{A}_i) \leq v_j(\mathcal{A}_j)/r$.

Definition 9. Suppose that we allocate one item to each agent using Nash social welfare matching. We define the envy-rank of an agent i , denoted by r_i as

$$r_i = \max_{j_0, j_1, \dots, j_k} \prod_{z=1}^k w_{j_z, j_{z-1}},$$

where $j_0 = i$.

Roughly speaking, let p be a path leading to vertex i in the envy-ratio graph such that the product of the weights of the edges in p is maximum. Then, the envy-rank of agent i equals the product of the weights of the edges in p . Note that by Observation 7 we can assume w.l.o.g. that p is a simple path (i.e., p includes no duplicate vertices).

Observation 10. p is a simple path.

Proof. Assume p is not simple and let c be a cycle in p . By Observation 7 we know that c cannot be an improving cycle. Therefore, the product of the weight of the edges of $p \setminus c$ is at least as large as that of p . \square

To get a better understanding of these definitions, take a look at Example 11.

Example 11. Consider an instance of the fair allocation problem with 4 items, 4 agents, and a valuation profile $V = \{v_1, v_2, v_3, v_4\}$ illustrated in Figure 3a. Let \mathcal{A} be the allocation that allocates item i to agent i . The envy-ratio graph and the envy graph of \mathcal{A} are shown in Figure 3b and Figure 3c respectively. This allocation is not envy-free. However, it is both EFX and EFR since each agent receives only one item.

As we mentioned before, the envy-rank of an agent can be seen as the maximum product of the weights of the edges in a path leading to that agent. For instance, consider agent 1. The envy-rank of this agent is 3 which is the product of the weights of the edges in the path $3 \rightarrow 2 \rightarrow 1$. Also consider the cycle $1 \rightarrow 3 \rightarrow 2 \rightarrow 1$. This cycle is an improving cycle. Therefore, allocation \mathcal{A} is not a NSW matching. The allocation can be improved by moving the items alongside this cycle which leads to a new allocation $\mathcal{A}' = \langle \{3\}, \{1\}, \{2\}, \{4\} \rangle$. The envy-ratio graph of \mathcal{A}' can be seen in Figure 3d.

We finish our discussion in this section by mentioning some properties of envy-rank values. The proofs of the following observations are available in Appendix A.


Observation 12. Suppose that allocation \mathcal{A} allocates one item to each agent using a Nash social welfare matching. Then for every pair of agents i and j , we have

$$\frac{v_i(\mathcal{A}_j)}{v_i(\mathcal{A}_i)} \leq \min \left\{ r_j, \frac{r_j}{r_i} \right\}.$$


In addition to Observation 7, Nash social welfare matchings admit another important and elegant property, which we state in Observation 13. This observation provides upper bounds on the value of remaining goods and can be of independent interest for various fair allocation problems.

	1	2	3	4
v_1	8	2	4	3
v_2	4	2	0	2
v_3	0	3	2	2
v_4	1	6	3	9


(a) Agents' valuations over items



(b) Envy-ratio graph



(c) Envy graph



(d) Envy-ratio graph after eliminating a cycle.

Figure 3: An example to illustrate envy-ratio graph.

ALGORITHM 1: The outline of the 0.73-EFR algorithm.

Parameters : $\varphi = \sqrt{3} + 1$.
// Step 1
Allocate NSW matching;
Let r_i be envy-rank of an agent i . Divide the agents into groups G_1, G_2, G_3 as follows. Agent i belongs to G_1 if $r_i > \varphi$, belongs to G_2 if $2 < r_i \leq \varphi$, and belongs to G_3 if $r_i \leq 2$;
// Step 2
Let \mathcal{O} be a topological ordering of the agents with respect to the envy-graph;
foreach $i \in G_3$ *ordered by* \mathcal{O} **do**
| Ask agent i to pick his most valuable remaining item;
end
foreach $i \in G_2$ *ordered by* \mathcal{O} **do**
| Ask agent i to pick his most valuable remaining item;
end
foreach $i \in G_1$ *ordered by* \mathcal{O} **do**
| Ask agent i to pick his most valuable remaining item;
end
// Step 3
while *the allocation is not complete* **do**
| Eliminate all directed cycles in the envy-graph;
| Let s be an arbitrary source in the envy-graph;
| Ask agent s to pick his most valuable remaining item;
end
return the allocation;

Observation 13. Suppose that we allocate one item to each agent using a Nash social welfare matching. Then, for each agent i and any unallocated item b we have

$$v_i(b) \leq \min \left\{ v_i(\mathcal{A}_i), \frac{v_i(\mathcal{A}_i)}{r_i} \right\}.$$

An Approximate EFR Allocation

In this section, we present our algorithm for finding a 0.73-EFR allocation. Our algorithm is divided into 3 steps, namely NSW matching, allocation refinement, and envy-graph based allocation. In the first step, we allocate to each agent one item using a Nash social welfare matching and accordingly divide the agents into three groups based on their envy-rank. Next, in the second step we allocate a set of goods to the agents in each group, and finally in the third step we allocate the rest of the items using the classic envy-cycle elimination method. The outline of our algorithm is represented in Algorithm 1.

Step 1

In the first step, we allocate one item to each agent using a NSW matching. We first show that this allocation can be found in polynomial time. The proof is available in Appendix A.

Observation 14. NSW matching can be found in polynomial time.

Let \mathcal{A} be an NSW matching and fix a parameter $\varphi = \sqrt{3} + 1$. Based on the envy-rank of the agents, we divide them into three groups G_1, G_2 , and G_3 as follows.

- Agent i belongs to G_1 if $r_i > \varphi$.
- Agent i belongs to G_2 if $2 < r_i \leq \varphi$.
- Agent i belongs to G_3 if $r_i \leq 2$.

Note that the envy-rank of an agent i which is r_i can be found in polynomial time. In order to find a path leading to vertex i with the maximum product of the weights of the edges, we can take the logarithm of the weights of edges, and find a path leading to i with the maximum summation. Since the envy-ratio graph does not have an improving cycle, after taking the logarithm of the weights of the edges, the resulting graph is without a positive cycle. Therefore, we have to find a longest path leading to i in a directed graph without a positive cycle which can be done in polynomial time using Bellman-Ford algorithm.

Considering the envy-ranks of agents, by Observation 13, we know that for every remaining item b the following properties hold.

- (Property 1): For every agent $i \in G_1$ we have $v_i(b) < v_i(\mathcal{A}_i)/\varphi$.
- (Property 2): For every agent $i \in G_2$ we have $v_i(b) < v_i(\mathcal{A}_i)/2$.

Intuitively, if each remaining item is worth less than $v_i(\mathcal{A}_i)/\varphi$ to every agent i , then we can guarantee the approximation factor of $1/(1 + 1/\varphi)$ in the third step. This property holds for the agents in G_1 ; however, this is not the case for agents in G_2 and G_3 . In the second step, we seek to allocate a set of items to the agents in G_2 and G_3 so that the same property holds for these agents. Note that alongside this property, the final partial allocation after the second step must be fair (i.e., 0.73-EFR).

Step 2

In the second step, we allocate one item to each agent in G_2 and two items to each agent in G_3 . Algorithm 1 shows the method by which we allocate these items to the agents in G_2 and G_3 . Let \mathcal{O} be a topological ordering of the agents with respect to the envy-graph. We order the agents in G_3 according to \mathcal{O} and ask them one by one to pick their most valuable remaining good. We then again ask agents in G_3 to pick one more item according to the same topological ordering \mathcal{O} . Afterwards, we order the agents in G_2 according to \mathcal{O} and ask them one by one to add the most desirable remaining item to their bundles.

We now show that at the end of Step 2 the following conditions hold. The proof can be found in Appendix B.

Claim 15. At the end of Step 2 the following conditions hold.

- The allocation is EFR with respect to the agents in G_1 .
- The allocation is (3/4)-EFR with respect to the agents in G_2 .
- The allocation is (2/φ)-EFR with respect to the agents in G_3 .

Since $2/\varphi < 3/4$, the allocation by the end of Step 2 is (2/φ)-EFR.

Step 3

In the third step, we use the envy-graph to allocate the remaining unallocated items. We repeat the following steps until all the goods are allocated.

- Find and eliminate all the directed cycles from the envy-graph. In order to eliminate all cycles in the envy-graph, we repeatedly find a directed cycle. Let $i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_k \rightarrow i_1$ be a cycle in envy-graph. By definition, each agent i_j envies agent $i_{(j \bmod k)+1}$, i.e.,

$$v_{i_j}(\mathcal{A}_{i_j}) < v_{i_j}(\mathcal{A}_{i_{(j \bmod k)+1}}),$$

where \mathcal{A} is the current allocation. We then exchange the allocations of the agents that are in the cycle such that each agent i_j receives $\mathcal{A}_{i_{(j \bmod k)+1}}$. Note that this exchanging does not change bundles. Furthermore, the utility of each agent does not decrease. Hence, if the allocation is α -EFR before the exchange, it remains α -EFR after it (*Lemma 6.1* in (Plaut and Roughgarden 2020)). Also, exchanging these allocations decreases the number of edges in the envy-graph. Thus, we eventually find an allocation such that its corresponding envy-graph is acyclic.

- Give an item to an agent that no-one envies. In the previous step we showed that we can always find an allocation such that its corresponding envy-graph is acyclic. Therefore, there should be a vertex in the envy-graph with no incoming edges. Let i be the agent corresponding to this vertex. Since i has no incoming edges in the envy-graph, no other agent envies i . At this step, we ask agent i to pick his most valued item among all remaining goods.

The following Lemma shows the approximation guarantee of our algorithm. The proof can be found in Appendix A.

Lemma 16. Suppose that we are given a partial α -EFR allocation \mathcal{A} such that for every agent i and every remaining item b , we have $v_i(b) \leq \alpha' \cdot v_i(\mathcal{A}_i)$ for some constant $\alpha' \leq 1$. Then, the resulting allocation after performing the method mentioned above is $\min\{\alpha, \frac{1}{1+\alpha'}\}$ -EFR.

We now show that at the beginning of Step 3, the valuation of every remaining item is small for all agents. The proof is available in Appendix A.

Observation 17. Let \mathcal{A} be the allocation after Step 2. Then for an agent i and every remaining item b we have

- If $i \in G_1$, $v_i(b) \leq v_i(\mathcal{A}_i)/\varphi$.
- If $i \in G_2$, $v_i(b) \leq v_i(\mathcal{A}_i)/3$.
- If $i \in G_3$, $v_i(b) \leq v_i(\mathcal{A}_i)/3$.

It follows from the observation above that for every agent i the valuation of every remaining item is at most $v_i(\mathcal{A}_i)/\varphi$ after the second step of our algorithm. Recall that our allocation by the end of Step 2 is $(2/\varphi)$ -EFR. Therefore, using Lemma 16, the allocation at the end of Step 3 is $\min\{\frac{2}{\varphi}, \frac{1}{1+1/\varphi}\}$ -EFR. Since $\varphi = \sqrt{3} + 1$, we have

$$\frac{2}{\varphi} = \frac{1}{1+1/\varphi} = \sqrt{3} - 1.$$

Therefore our final allocation is $\sqrt{3} - 1 \approx 0.73$ -EFR. This, coupled with the fact that all the steps can be implemented in polynomial time follows Theorem 1.

Theorem 1. There exists an algorithm that finds a 0.73-EFR allocation. In addition, such an allocation can be found in polynomial time.

Conclusions and Future Directions

The envy-free relaxations recently have received significant attention in the field of fair division. Numerous recent papers considered EFX as their fairness notion and provided various results regarding this notion. However, the main problem is still open, and it is unknown if the EFX allocation always exists, while we can always guarantee EF1, a weaker fairness notion, with a simple round-robin algorithm. This difference motivated us to propose a new fairness notion that lies somewhere in between EF1 and EFX. We proposed the EFR notion and provided a polynomial-time 0.72-approximation EFR algorithm. This approximation factor improves the best-known approximation factor for EFX. The first future direction is to improve the approximation guarantee of EFR and provide more properties for further differentiating it from EFX. The next future direction is to consider the connection of EFR with other fairness notions. Finally, it is also interesting to provide an algorithm that guarantees both EFR approximation and efficiency at the same time.

References

- Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2020. Maximum Nash Welfare and Other Stories About EFX. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence*, 24–30.
- Amanatidis, G.; Birmpas, G.; and Markakis, E. 2018. Comparing approximate relaxations of envy-freeness. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 42–48.
- 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.
- Barman, S.; Biswas, A.; Krishnamurthy, S. K.; and Narahari, Y. 2018. Groupwise maximin fair allocation of indivisible goods. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Barman, S.; and Krishna Murthy, S. K. 2017. Approximation algorithms for maximin fair division. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 647–664.
- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 557–574.
- 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 (ACM-EC)*, 557–574. ACM Press.

- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Greedy Algorithms for Maximizing Nash Social Welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 7–13.
- Brams, S. J.; and Taylor, A. D. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press.
- 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 *Proceedings of the 2019 ACM Conference on Economics and Computation*, 527–545.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation* 7(3): 12.
- Chan, H.; Chen, J.; Li, B.; and Wu, X. 2019. Maximin-Aware Allocations of Indivisible Goods. In Elkind, E.; Veloso, M.; Agmon, N.; and Taylor, M. E., eds., *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, 1871–1873. International Foundation for Autonomous Agents and Multiagent Systems. URL <http://dl.acm.org/citation.cfm?id=3331947>.
- Chaudhury, B. R.; Cheung, Y. K.; Garg, J.; Garg, N.; Hoefler, 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 2018, December 11-13, 2018, Ahmedabad, India*, 25:1–25:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Chaudhury, B. R.; Garg, J.; and Mehlhorn, K. 2020. EFX exists for three agents. *arXiv preprint arXiv:2002.05119*.
- Chaudhury, B. R.; Garg, J.; and Mehta, R. 2020. Fair and Efficient Allocations under Subadditive Valuations. *arXiv preprint arXiv:2005.06511*.
- Chaudhury, B. R.; Kavitha, T.; Mehlhorn, K.; and Sgouritsa, A. 2020. A Little Charity Guarantees Almost Envy-Freeness. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms*, 2658–2672.
- Cole, R.; and Gkatzelis, V. 2018. Approximating the Nash Social Welfare with Indivisible Items. *SIAM Journal on Computing, Volume 47* 1211–1236.
- Farhadi, A.; Ghodsi, M.; Hajiaghayi, M. T.; Lahaie, S.; Pennock, D.; Seddighin, M.; Seddighin, S.; and Yami, H. 2019. Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research* 64: 1–20.
- Foley, D. K. 1967. Resource allocation and the public sector. *Yale Econ Essays* 7(1): 45–98.
- Garg, J.; Kulkarni, P.; and Kulkarni, R. 2020. Approximating Nash social welfare under submodular valuations through (un) matchings. In *Proceedings of the fourteenth annual ACM-SIAM symposium on discrete algorithms*, 2673–2687. SIAM.
- Ghodsi, M.; HajiAghayi, M.; 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*, 539–556.
- Kurokawa, D.; Procaccia, A. D.; and Wang, J. 2018. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM* 65(2): 8.
- 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*, 125–131.
- Nash Jr, J. F. 1950. The bargaining problem. *Econometrica: Journal of the Econometric Society* 155–162.
- Nguyen, T. T.; and Rothe, J. 2014. Minimizing Envy and Maximizing Average Nash Social Welfare in the Allocation of Indivisible Goods. *Discrete Applied Mathematics* 54–68.
- Plaut, B.; and Roughgarden, T. 2020. Almost Envy-Freeness with General Valuations. *SIAM J. Discret. Math.* 34(2): 1039–1068. doi:10.1137/19M124397X. URL <https://doi.org/10.1137/19M124397X>.
- Robertson, J.; and Webb, W. 1998. *Cake-cutting algorithms: Be fair if you can*. AK Peters/CRC Press.
- Steinhaus, H. 1948. The problem of fair division. *Econometrica* 16(1).