

A Complete Landscape for the Price of Envy-Freeness

Zihao Li
Nanyang Technological University
Singapore
zihao004@e.ntu.edu.sg

Shengxin Liu
Harbin Institute of Technology,
Shenzhen
China
sxliu@hit.edu.cn

Xinhang Lu
UNSW Sydney
Australia
xinhang.lu@unsw.edu.au

Biaoshuai Tao
Shanghai Jiao Tong University
China
bstao@sjtu.edu.cn

Yichen Tao
University of Michigan
United States
yghtao@umich.edu

ABSTRACT

We study the efficiency of fair allocations using the well-studied *price of fairness* concept, which quantitatively measures the worst-case efficiency loss when imposing fairness constraints. Previous works provided partial results on the price of fairness with well-known fairness notions such as envy-freeness up to one good (EF1) and envy-freeness up to any good (EFX). In this paper, we give a complete characterization for the price of envy-freeness in various settings. In particular, we first consider the two-agent case under the indivisible-goods setting and present tight ratios for the price of EF1 (for scaled utility) and EFX (for unscaled utility), which resolve questions left open in the literature. Next, we consider the mixed goods setting which concerns a mixture of both divisible and indivisible goods. We focus on envy-freeness for mixed goods (EFM), which generalizes both envy-freeness and EF1, as well as its strengthening called envy-freeness up to any good for mixed goods (EFXM), which generalizes envy-freeness and EFX. To this end, we settle the price of EFM and EFXM by providing a complete picture of tight bounds for two agents and asymptotically tight bounds for n agents, for both scaled and unscaled utilities.

CCS CONCEPTS

• **Theory of computation** → **Algorithmic game theory**.

KEYWORDS

Fair division; Price of fairness; Envy-freeness; Mixed goods; EF1 / EFX / EFM / EFXM

ACM Reference Format:

Zihao Li, Shengxin Liu, Xinhang Lu, Biaoshuai Tao, and Yichen Tao. 2024. A Complete Landscape for the Price of Envy-Freeness. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 9 pages.

1 INTRODUCTION

Fair division is a fundamental topic in algorithmic game theory and has attracted wide attention. In this problem, we need to allocate

some resources among agents in a fair manner. The most classic notion of fairness is *envy-freeness (EF)*, which requires that each agent does not envy another agent’s bundle. In other words, every agent values her own bundle weakly better than the bundle received by any other agent. When the goods are divisible, meaning that they can be divided into arbitrarily small pieces and allocated to different agents (the cake-cutting problem), an envy-free allocation always exists [2]. However, when considering indivisible goods, meaning that each of them should be allocated to an agent in its entirety, an envy-free allocation may not exist. Thus, a relaxation of envy-freeness, called *envy-freeness up to one good (EF1)*, has been proposed to circumvent this issue. An allocation is EF1 if for any pairs of agents i and j , agent i does not envy agent j after removing one item from j ’s bundle. Such an allocation always exists when allocating indivisible goods [41]. Besides EF1, another relaxation of envy-freeness called *envy-freeness up to any good (EFX)* has also received wide attention in recent years [e.g., 1, 4, 30, 31, 34, 37, 46], where an agent does not envy another agent after removing *any* item from the latter agent’s bundle. Moreover, there have been a number of papers on *partial* EFX allocations with good properties [e.g., 19, 27, 32]. Beyond the setting concerning either divisible or indivisible goods, some recent studies have focused on fairly dividing a mixture of both divisible and indivisible resources [15–17, 21, 38, 39, 45]. Among these, Bei et al. [15] first considered the mixed-goods setting and proposed the fairness notion called *envy-freeness for mixed goods (EFM)*, which naturally combines envy-freeness and EF1 together. Specifically, under an EFM allocation, for each agent, if she receives only indivisible goods, no other agents envy her using the EF1 criterion; otherwise, no other agents envy her using the EF criterion. An EFM allocation is guaranteed to exist [15]. By strengthening EF1 to EFX when assessing the envy from others to an agent who receives only indivisible goods, we have a stronger notion than EFM, which is called *envy-freeness up to any good for mixed goods (EFXM)* [15, 45].

In addition to fairness, *efficiency*, or *social welfare*, which refers to the total utility of all the agents towards their bundles, also plays an important role in evaluating an allocation [23, 24]. The *price of fairness*, introduced independently by Bertsimas et al. [20] and Caragiannis et al. [28], is a quantitative measure indicating the loss of social welfare when a given fairness constraint is imposed. More specifically, the price of fairness is defined as the supremum ratio of the maximum social welfare among all allocations to the maximum social welfare among all fair allocations under a particular fairness



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

Table 1: Price of Envy-Freeness for Two Agents with (Un)Scaled Utilities

Price of ...	EF1	EFM / EFXM	EFX
Scaled	$\frac{8}{7}$ (Thm. 3.1)	$\frac{3}{2}$ (Thm. 3.5)	$\frac{3}{2}$ [18]
Unscaled	2 [24]	2 (Thm. 3.4)	2 (Thm. 3.4)

Table 2: Price of Envy-Freeness for n Agents with (Un)Scaled Utilities

Price of ...	EF1	EFM / EFXM	EFX
Scaled	$\Theta(\sqrt{n})$ [8, 18]	$\Theta(\sqrt{n})$ (Thm. 4.1)	$\Theta(\sqrt{n})$ [24]
Unscaled	$\Theta(n)$ [24]	$\Theta(n)$ (Thm. 4.2)	$\Theta(n)$ [24]

property, where the supremum is taken over all possible instances. For divisible goods, the price of envy-freeness is $\Theta(\sqrt{n})$, where n is the number of agents [20]. For indivisible goods, the price of EF1 is $\Theta(\sqrt{n})$ [8, 18]. For the special case where $n = 2$, Bei et al. [18] provided a lower bound of $8/7$ and an upper bound of $\frac{2}{\sqrt{3}}$ for the price of EF1 for scaled utilities, and Bu et al. [24] gave a tight ratio of 2 for unscaled utilities. Regarding the price of EFX for indivisible goods, with $n = 2$ agents, Bei et al. [18] provided a tight bound of $3/2$ under scaled utilities, and Bu et al. [24] presented a lower bound of 2 under unscaled utilities. Bu et al. also showed tight bounds for the price of EFX for n agents for both scaled and unscaled utilities, which are $\Theta(\sqrt{n})$ and $\Theta(n)$, respectively.

In this paper, we close gaps on the price of EF1 and EFX for indivisible-goods allocation when there are two agents [18, 24]. Moreover, for the first time, we provide tight or asymptotically tight bounds on the price of EFM and EFXM in the mixed-goods setting, resolving questions left open in the literature [42]. Our results are shown in Tables 1 and 2. Below, we highlight some interesting features/observations according to our results:

- In all of the settings, tight bounds (or asymptotically tight bounds for an arbitrary number of agents) on the price of envy-freeness are now known.
- For two agents, we show that the price of EF1 is exactly $8/7$, which closes the gap between $8/7$ and $2/\sqrt{3}$ left open in the previous paper by Bei et al. [18].
- The price of EFM is (asymptotically) the same as the price of EFX in all the settings. This is a potential evidence that EFM, although defined in relation to EF1, is more similar to EFX in nature.

1.1 Additional Related Work

While the price of fairness concept captures the efficiency loss in the best fair allocation, Bei et al. [18] introduced the concept of *strong price of fairness*, which captures the efficiency loss in the *worst* allocation. The strong price of fairness has proven to provide meaningful guarantees for fairness notions defined in the form of welfare maximizers, e.g., maximum Nash welfare, maximum

Egalitarian welfare, and leximin. The strong price of fairness is, however, too demanding to yield any non-trivial guarantee for fairness notions of interest in this paper. To be more specific, the strong price of EF1 and the strong price of EFX are ∞ [18]. We thus only focus on the price of fairness in our paper.

The interplay between fairness and efficiency has been extensively studied in the literature of fair division for both divisible and indivisible goods settings [5, 6, 10, 44]. An immediate question is whether a fairness criterion is compatible with Pareto optimality (PO) which is a rather weak economic efficiency measurement. In particular, both envy-freeness and EF1 can be combined with PO by finding the allocation that satisfies maximum Nash welfare in the divisible and indivisible settings respectively [29, 48].

Another related direction considers the problem of maximizing social welfare subject to fairness constraints. For divisible goods, this optimization problem with the envy-freeness constraints for piecewise-constant valuations can be solved optimally [35]. However, this problem has an inapproximability hardness with a polynomial factor if an additional requirement of connectivity on the received piece of cake is imposed [12]. For indivisible goods, the problem with EF1 / EFX criteria is well understood recently [7, 9, 24]. With scaled utilities, Barman et al. [9] gave inapproximability results for general numbers of agents and items while Aziz et al. [7] showed that the problem subject to the EF1 / EFX constraints is NP-hard for some special cases. Bu et al. [24] gave a complete landscape on the computational complexity and approximability of maximizing the social welfare within EF1 / EFX allocations of indivisible goods for both scaled and unscaled utilities.

2 PRELIMINARIES

For any positive integer $t \in \mathbb{N}$, let $[t] := \{1, 2, \dots, t\}$. We consider both indivisible-goods and mixed-goods settings with a set of n agents $N = [n]$. The mixed-goods setting involves a set of m *indivisible goods* $M = \{g_1, g_2, \dots, g_m\}$ and a set of \bar{m} *homogeneous divisible goods* $D = \{d_1, d_2, \dots, d_{\bar{m}}\}$. The indivisible-goods setting can be viewed as a special case where $D = \emptyset$.

Utilities. Denote by $\mathbf{u} = (u_1, u_2, \dots, u_n)$ the *utility profile*, which specifies how the agents value the goods. Each agent $i \in N$ has a non-negative utility $u_i(g)$ for each (in)divisible good $g \in M \cup D$. A *bundle* is a tuple consisting of a (possibly empty) set of indivisible goods $M' \subseteq M$ and an \bar{m} -dimensional vector $\mathbf{x}' = (x'_1, x'_2, \dots, x'_{\bar{m}}) \in [0, 1]^{\bar{m}}$, in which each coordinate specifies the fraction of the corresponding homogeneous divisible good. The bundle is written as (M', \mathbf{x}') . We assume *additive utilities*, i.e., for all $i \in N$, $M' \subseteq M$, and $\mathbf{x}' \in [0, 1]^{\bar{m}}$, the utility of agent i for bundle (M', \mathbf{x}') is defined as:

$$u_i(M', \mathbf{x}') := \sum_{g \in M'} u_i(g) + \sum_{k=1}^{\bar{m}} x'_k \cdot u_i(d_k).$$

We slightly abuse the notation by letting $u_i(M') := \sum_{g \in M'} u_i(g)$ and $u_i(\mathbf{x}') := \sum_{k=1}^{\bar{m}} x'_k \cdot u_i(d_k)$.

An *instance*, written as $\langle N, M, D, \mathbf{u} \rangle$, consists of agents N , indivisible goods M , divisible goods D , and the agents' utility profile \mathbf{u} .

REMARK 2.1 (PIECEWISE-CONSTANT VALUATION ASSUMPTION). *In the paper by Bei et al. [15] where the mixed-goods model was first*

introduced, the divisible part of the resource is modelled by a (heterogeneous) cake on which each agent has a value density function. Our setting with multiple homogeneous divisible goods is equivalent to the setting with a cake where all agents’ value density functions are piecewise-constant. Although being somewhat more restrictive than general valuations, piecewise-constant valuations are standardly assumed in the cake-cutting literature for their ability to represent natural real functions with arbitrarily good precision and to be encoded succinctly [e.g., 5, 6, 12–14, 22, 25, 33, 35, 43].

The detailed discussion on our setting of multiple homogeneous divisible goods can be found in the full version of our paper [40].

Allocations. An allocation is denoted by $\mathcal{A} = (A_1, A_2, \dots, A_n)$, where for each $i \in [n]$, $A_i = (M_i, \mathbf{x}_i = (x_{i1}, x_{i2}, \dots, x_{i\bar{m}}))$ is the bundle allocated to agent i . Allocation \mathcal{A} is feasible if

- for any pair of agents $i \neq j$, $M_i \cap M_j = \emptyset$; and
- for each $\bar{k} = 1, 2, \dots, \bar{m}$, $\sum_{i \in N} x_{i\bar{k}} \leq 1$.

Allocation \mathcal{A} is *complete* if $\bigcup_{i \in N} M_i = M$ and for each $\bar{k} \in [\bar{m}]$, $\sum_{i \in N} x_{i\bar{k}} = 1$; otherwise, this is a *partial* allocation. In this work, we consider feasible allocations and allow partial allocations.

Scaling. Agents’ utilities are *scaled* if for all $i \in N$, $u_i(M \cup D) = 1$, and *unscaled* otherwise. In this work, we consider both scaled and unscaled utilities.

Social Welfare. The (utilitarian) *social welfare* of an allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ is defined as $\text{SW}(\mathcal{A}) := \sum_{i \in N} u_i(A_i)$. The optimal social welfare for an instance I , denoted by $\text{OPT}(I)$, is the maximum social welfare over all allocations of this instance.

Fairness Notions. To better understand the intuition behind fairness notions like EF1, EFM, EFXM and EFX, we start with the definition of envy-freeness [36]. An allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ is said to satisfy *envy-freeness (EF)* if for any pair of agents $i, j \in N$, we have $u_i(A_i) \geq u_i(A_j)$. EF1 relaxes envy-freeness by allowing envies “up to one good”.

Definition 2.2 (EF1 [26, 41]). An allocation (M_1, M_2, \dots, M_n) of indivisible goods M is said to satisfy *envy-freeness up to one good (EF1)* if for any pair of agents $i, j \in N$ and $M_j \neq \emptyset$, there exists a good $g \in M_j$ such that $u_i(M_i) \geq u_i(M_j \setminus \{g\})$.

When allocating indivisible goods, another commonly studied fairness notion is *envy-freeness up to any good (EFX)*, which is substantially stronger than EF1.

Definition 2.3 (EFX [29, 46]). An allocation (M_1, M_2, \dots, M_n) of indivisible goods M is said to satisfy *envy-freeness up to any good (EFX)* if for any pair of agents $i, j \in N$ and $M_j \neq \emptyset$, $u_i(M_i) \geq u_i(M_j \setminus \{g\})$ holds for any good $g \in M_j$.

We now proceed to define EFM in the mixed-goods setting, which, intuitively, requires that an agent compares her bundle to another agent’s bundle using EF1 criterion if the latter bundle consists of only indivisible goods; otherwise, the stronger envy-freeness criterion is invoked.

Definition 2.4 (EFM [15, Definition 2.3]). An allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ of mixed goods is said to satisfy *envy-freeness for mixed goods (EFM)* if for any pair of agents $i, j \in N$,

- if agent j ’s bundle consists of only indivisible goods, i.e., $\mathbf{x}_j = \mathbf{0}$, then $u_i(A_i) \geq u_i(A_j)$ or there exists some (indivisible) good $g \in M_j$ such that $u_i(A_i) \geq u_i(M_j \setminus \{g\}, \mathbf{x}_j)$;
- otherwise, $u_i(A_i) \geq u_i(A_j)$.

Bei et al. [15] also suggested that by substituting the EFX criterion for EF1 in cases where an agent compares her bundle to those containing only indivisible goods, a stronger variant of EFM can be derived. Its formal definition was later introduced in the work of Nishimura and Sumita [45].

Definition 2.5 (EFXM [45]). An allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ of mixed goods is said to satisfy *envy-freeness up to any good for mixed goods (EFXM)* if for any pair of agents $i, j \in N$,

- if agent j ’s bundle consists of only indivisible goods, i.e., $\mathbf{x}_j = \mathbf{0}$, then $u_i(A_i) \geq u_i(A_j)$ or $u_i(A_i) \geq u_i(M_j \setminus \{g\}, \mathbf{x}_j)$ for any (indivisible) good $g \in M_j$;
- otherwise, $u_i(A_i) \geq u_i(A_j)$.

2.1 Price of Fairness

We are now ready to define the central concept of the paper—the price of fairness [8, 18, 20, 28].

Definition 2.6 (Price of Fairness P). For any given fairness criteria P and any instance I , we define

$$\text{price of } P \text{ for instance } I = \frac{\text{OPT}(I)}{\text{SW}(\mathcal{A}^*)},$$

where \mathcal{A}^* is a (partial) allocation that satisfies P and has the maximum social welfare.

The overall *price of P* is then defined as the supremum price of P across all instances.

Note that when we define the price of EF1 and EFX, we consider instance I with $D = \emptyset$.

Partial Allocation and Resource Monotonicity. As a remark, when defining the price of EFM, EFXM and EFX, we include partial allocations into our consideration. To illustrate the idea here, let us first introduce the concept of *resource monotonicity* with respect to social welfare¹ [see, e.g., 24]. Given a fairness property P , we say *resource monotonicity holds* for property P if for any instance, there always exists a complete allocation satisfying P that has a weakly higher social welfare than any other partial allocation satisfying P . Note that when the existence of property P is not guaranteed, resource monotonicity fails for the property.²

Regarding EFX, Bu et al. [24] showed that for two agents, resource monotonicity holds for EFX. In general, however, resource monotonicity *fails* for EFX [24]. Put differently, Bu et al. [24] provided an instance in which a partial EFX allocation has a higher social welfare than any complete EFX allocation. Note also that it is a major open problem whether a complete EFX allocation always exists [3]. We have the following two scenarios:

- (1) if an EFX allocation of indivisible goods always exists, its failure of resource monotonicity suggests that some partial

¹Prior research has also explored the concept of resource monotonicity subject to Pareto optimality. [see, e.g., 47].

²For instance, resource monotonicity fails for envy-freeness when allocating indivisible goods.

EFX allocation may have higher social welfare than any complete EFX allocation; and

(2) otherwise we may not even have a complete EFX allocation.

In either case, independent of the existence of EFX, it is more natural to include partial allocations when defining the price of EFX. Since EFXM is a generalization of EFX, it is also natural to consider partial allocations in its price of fairness definition.

In terms of EFM, its existence is guaranteed [15]; however, we do not know whether resource monotonicity holds for EFM (in the mixed-goods setting) or not. If any partial EFM allocation can be extended to a complete EFM allocation with a weakly higher social welfare, it makes no difference whether or not partial allocations are included when defining the price of EFM. Otherwise, it would be more natural to use the “better” partial allocation for characterizing the price of EFM, as opposed to forcing the allocation being complete. Thus, it is again more natural to include partial allocations into our consideration.

For EF1, as noted in [24], it is easy to see that any partial EF1 allocation can be extended to a complete EF1 allocation with a weakly higher social welfare by carrying on the *envy-graph procedure*. Therefore, we do not need to consider partial allocations for the price of EF1.

3 TWO AGENTS

In this section, we establish tight bounds on the price of EF1 / EFX / EFM / EFXM for two agents with scaled or unscaled utilities.

3.1 Price of EF1 (for Indivisible Goods)

For unscaled utilities, the price of EF1 is exactly 2 due to Bu et al. [24]. For scaled utilities, the price of EF1 is between $\frac{8}{7}$ and $\frac{2}{\sqrt{3}}$ due to Bei et al. [18]. In the following, we close the gap by showing that the price of EF1 is $\frac{8}{7}$.

THEOREM 3.1. *For $n = 2$ and scaled utilities (over indivisible goods), the price of EF1 is $\frac{8}{7}$.*

Some additional notations are introduced here for the sake of clarity during this proof. Denote by T_1 the subset of goods that agent 1 values no less than agent 2, and by T_2 the remaining goods:

$$T_1 = \{g \in M \mid u_1(g) \geq u_2(g)\} \quad \text{and} \quad T_2 = M \setminus T_1.$$

Let the *surplus* of a bundle M' , denoted by $SP(M')$, be how much agent 1 values bundle M' more than agent 2; formally stated as

$$SP(M') = \sum_{g \in M'} (u_1(g) - u_2(g)).$$

We first state and prove a useful proposition, and then provide the proof of Theorem 3.1. Given an allocation (M_1, M_2, \dots, M_n) of indivisible goods M , we say that an agent $i \in N$ *strongly envies* an agent j if and only if for any $g \in M_j$, $u_i(M_i) < u_i(M_j \setminus \{g\})$. It can be seen that given an allocation, if some agent strongly envies some other agent, then the allocation is not EF1. Moreover, if every agent does not strongly envy any other agent, then the allocation is EF1.

PROPOSITION 3.2. *Suppose agent 2 strongly envies agent 1 in the allocation (T_1, T_2) . If T_1 can be partitioned into T'_A and T'_B such that*

agent 2 does not strongly envy agent 1 in the allocation $\mathcal{A} = (T'_A, T'_B \cup T_2)$, then there exists an EF1 allocation \mathcal{A}' with $SW(\mathcal{A}') \geq SW(\mathcal{A})$.

PROOF. If agent 1 does not strongly envy agent 2 in \mathcal{A} , then let $\mathcal{A}' = \mathcal{A}$ and we are done. Otherwise, we apply the following iterative “one-by-one reassignment” process:

- (1) Suppose, at the beginning of the iteration, agent 1 has bundle T'_1 and agent 2 has bundle T'_2 . Select an arbitrary good $g \in T'_2 \cap T_1$. Since agent 1 strongly envies T'_2 , she would still envy the bundle if g is excluded, that is, $u_1(T'_1) < u_1(T'_2 \setminus \{g\})$.
- (2) If $u_2(T'_2 \setminus \{g\}) \geq u_2(T'_1)$, assign good g to agent 1. Otherwise we have $u_2(T'_2 \setminus \{g\}) < u_2(T'_1)$. We now swap the two agents’ bundles, i.e., let agent 1 get bundle T'_2 and agent 2 get bundle T'_1 . Note that given the updated allocation, agent 2 is still EF1 towards agent 1.
- (3) If agent 1 still strongly envies agent 2’s bundle, go back to step 1 and start another iteration of this process.

Then we prove that the following two invariants hold at the end of each iteration: (a) the social welfare does not decrease throughout the process; (b) agent 2 does not strongly envy agent 1’s bundle.

- If agent 2 does not envy agent 1 even when g is excluded from her bundle, moving g to agent 1’s bundle would increase social welfare, because $g \in T_1$, indicating that agent 1 values it more than agent 2; if otherwise, the two agents envy each other’s bundle, swapping their bundle would also increase social welfare, and item g , too, is reassigned to the agent that values it more.
- If agent 2 does not envy agent 1 even when g is excluded from her bundle, adding g to agent 1’s bundle would not lead to agent 2 strongly envying the bundle; if otherwise, both agents would not envy each other’s bundle after swapping, and agent 2 certainly does not strongly envy agent 1’s bundle after g is reassigned.

The two invariants being established, it can be seen that EF1 is guaranteed after the whole one-by-one reassignment process, and the social welfare of the resulting allocation \mathcal{A}' is no less than $SW(\mathcal{A})$. \square

We are now ready to establish Theorem 3.1.

PROOF OF THEOREM 3.1. For brevity, let y be $SP(T_1)$. Since scaled utilities are considered here, we should have that $SP(T_2) = -SP(T_1) = -y$. Note that assigning each item to the agent that values it more, i.e., assigning T_1 to agent 1 and T_2 to agent 2, achieves the optimal social welfare, so for any instance I with two agents, scaled utilities, and only indivisible goods,

$$OPT(I) = u_1(T_1) + u_2(T_2) = SP(T_1) + u_2(T_1) + u_2(T_2) = 1 + y.$$

We divide our proof into three cases.

Case 1: $y \geq \frac{1}{2}$. This case is trivial, since the optimal allocation (allocating each item to the agent who values it more) would have already achieved envy-freeness. More formally speaking, in the optimal allocation, the bundle assigned to agent 1 is T_1 , while agent 2 receives T_2 . Hence, the utility of agent 1’s bundle

$$u_1(T_1) = u_2(T_1) + SP(T_1) \geq y \geq \frac{1}{2}.$$

Since agent 1 has received more than half of the total value of all items from her perspective, there is no chance that she would envy agent 2's bundle. Similarly, for agent 2,

$$u_2(T_2) = u_1(T_2) - SP(T_2) = u_1(T_2) + y \geq \frac{1}{2}.$$

Therefore, neither would agent 2 envy agent 1's bundle, and envy-freeness is proved. Because EF1 can be achieved via an allocation optimizing social welfare, the price of EF1 for such instances is 1.

Case 2: $y \leq \frac{1}{3}$. If the optimal allocation already satisfies EF1, we are done. For this reason, we only consider instances where the optimal allocation violates EF1, and suppose without loss of generality that agent 2 strongly envies agent 1. Note that agent 1 does not strongly envy agent 2's bundle in the optimal allocation, for otherwise exchanging their bundle would lead to an allocation with higher social welfare, contradicting to optimality. Then we let agent 2 partition T_1 into two subsets "as evenly as possible", maximizing the utility of the subset with lower utility from her perspective. Call the two subsets T_A and T_B . Suppose $SP(T_A) \geq SP(T_B)$. We temporarily assign T_A to agent 1, and $T_B \cup T_2$ to agent 2, and it can be proved that, at this time, (a) agent 2 does not strongly envy agent 1; (b) the social welfare is no less than $\frac{y}{2} + 1$.

- If $u_2(T_B) \geq u_2(T_A)$, then agent 2 values her bundle more than agent 1's, and envy-freeness is guaranteed. If $u_2(T_B) \leq u_2(T_A)$, there should exist a certain item $g \in T_A$ such that $u_2(T_B) \geq u_2(T_A \setminus \{g\})$, for otherwise, moving this item to T_B would result in a "more even" partition. Therefore, $u_2(T_B \cup T_2) \geq u_2(T_B) \geq u_2(T_A \setminus \{g\})$.
- The lower bound of social welfare can be established as follows:

$$\begin{aligned} SW(T_A, T_B \cup T_2) &= SP(T_A) + u_2(T_A \cup T_B \cup T_2) \\ &\geq \frac{1}{2}SP(T_1) + u_2(M) \\ &= \frac{y}{2} + 1. \end{aligned}$$

The inequality can be derived by the fact that $SP(T_A) \geq SP(T_B)$.

At this time, agent 1 can still strongly envy agent 2, and we apply Proposition 3.2 to derive a EF1 allocation \mathcal{A}' with social welfare no less than $\frac{y}{2} + 1$. Consequently, for instances with $y \leq \frac{1}{3}$, there exists an allocation \mathcal{A}' with $SW(\mathcal{A}') \geq \frac{y}{2} + 1$ satisfying EF1, ergo the price of EF1

$$\frac{OPT(I)}{SW(\mathcal{A}')} \leq \frac{1+y}{1+\frac{y}{2}} \leq \frac{8}{7}.$$

Case 3: $\frac{1}{3} \leq y \leq \frac{1}{2}$. Again, suppose without loss of generality that agent 2 strongly envies agent 1 under optimal allocation. Let agent 2 partition T_1 into three subsets "as evenly as possible", maximizing the subset with minimum utility from her perspective. Let the three subsets be T_A, T_B and T_C , and $u_2(T_A) \geq u_2(T_B) \geq u_2(T_C)$. Three subcases are studied here.

Subcase 3.1: $u_2(T_C) \geq \frac{1}{6}$. Since $u_2(T_2) = u_1(T_2) - SP(T_2) \geq y \geq \frac{1}{3}$, assigning any one of T_A, T_B , and T_C to agent 2 would result in she claiming more than half of all items' total utility, eliminating the possibility of agent 2 envying agent 1. We assign the subset with smallest surplus, dubbed T_{13} , to agent 2, and the other two,

dubbed T_{11} and T_{12} , to agent 1. Thus, $SP(T_{11} \cup T_{12}) \geq \frac{2}{3}SP(T_1)$. The social welfare of such allocation \mathcal{A} can be lower bounded via the following calculation

$$SW(\mathcal{A}) = SP(T_{11} \cup T_{12}) + u_2(M) \geq 1 + \frac{2}{3}y.$$

Since it is still possible that agent 1 strongly envies agent 2 at this moment, we apply the "one-by-one assignment" process in Proposition 3.2, and derive an allocation \mathcal{A}' satisfying EF1, with social welfare no less than $1 + \frac{2}{3}y$.

Subcase 3.2: $u_2(T_A) \leq \frac{1}{3}$. If any one of T_A, T_B , and T_C is assigned to agent 2, she would not strongly envy agent 1. Assume she gets T_C . For any item $g \in T_B$, $u_2(T_C) \geq u_2(T_B \setminus \{g\})$ should hold, because moving g to T_C would give a more even partition otherwise. Furthermore, $u_2(T_2) \geq y \geq \frac{1}{3} \geq u_2(T_A)$. Thus, $u_2(T_C \cup T_2) \geq u_2(T_A \cup T_B \setminus \{g\})$ for any $g \in T_B$, proving the claim that agent 2 does not strongly envy agent 1. If T_A or T_B is assigned to agent 2 instead, the proof is similar. Again, assigning the subset with smallest surplus to agent 2 would result in an allocation with social welfare no less than $1 + \frac{2}{3}y$. Since it is possible that agent 1 strongly envies agent 2 at this moment, we apply the "one-by-one assignment" process in Proposition 3.2, and derive an allocation \mathcal{A}' satisfying EF1, with social welfare no less than $1 + \frac{2}{3}y$.

Subcase 3.3: $u_2(T_A) \geq \frac{1}{3}$ and $u_2(T_C) \leq \frac{1}{6}$. In this case, all items in T_A must have values no less than $\frac{1}{6}$, for otherwise moving this item to T_C would bring about a more even allocation. Furthermore, observe that if there are two items in T_A , moving one of them to T_C would also result in a more even allocation. Thus, there can only be one item in T_A . Call it g_A . The optimal allocation here already satisfies EF1, because $u_2(T_2) \geq y = \frac{1}{3}$, and $u_2(T_1 \setminus \{g_A\}) = u_2(T_1) - u_2(g_A) \leq \frac{2}{3} - \frac{1}{3} = \frac{1}{3}$.

Concluding the three subcases discussed above, for any instance I with $\frac{1}{3} \leq y \leq \frac{1}{2}$, there exists an EF1 allocation \mathcal{A} with $SW(\mathcal{A}) \geq 1 + \frac{2}{3}y$. Therefore, the price of EF1 for instance I is

$$\frac{OPT(I)}{SW(\mathcal{A})} \leq \frac{1+y}{1+\frac{2}{3}y} \leq \frac{9}{8}.$$

Combining all three cases, the price of EF1 is at most $\frac{8}{7}$. Together with the lower bound provided in Bei et al. [18], we concluded that for two agents, the price of EF1 is exactly $\frac{8}{7}$. \square

3.2 Price of EFX / EFM / EFXM

Regarding EFX, it is known that for scaled utilities, the price of EFX is $\frac{3}{2}$ due to Bei et al. [18] as well as for unscaled utilities, the price of EFX is at least 2 due to Bu et al. [24]. We provide here a matching upper bound and thus conclude that for unscaled utilities, the price of EFX is exactly 2. In addition, we provide a complete picture of tight bounds on the price of EFM and EFXM for two agents with scaled or unscaled utilities.

We start by showing that a variant of the well-known *Cut-and-Choose Algorithm* outputs an EFXM (and thus EFM) allocation with social welfare at least one half of $u_1(M \cup D) + u_2(M \cup D)$. The same idea has also been used to show the price of EFX for two agents when allocating indivisible goods; see Theorem 3.4 of Bei et al. [18]. We slightly tailor the algorithm description to allocating mixed goods.

Algorithm 1: Cut-and-Choose Algorithm**Input:** Fair division instance $\langle [2], M, D, \mathbf{u} \rangle$.**Output:** An EFM (or EFX if $D = \emptyset$) allocation with social welfare at least $\frac{u_1(M \cup D) + u_2(M \cup D)}{2}$.

- 1 Let agent 1 (resp., agent 2) partition the mixed goods into two bundles, denoted by X_1, X_2 (resp., Y_1, Y_2), in the sense that her values for the two bundles are as equal as possible. Assume without loss of generality that $u_1(X_1) \geq u_1(X_2)$, $|u_1(X_1) - u_1(X_2)| \leq |u_2(Y_1) - u_2(Y_2)|$ and that between bundles X_1 and X_2 , all goods of value zero for agent 1, if any, are in bundle X_2 .
- 2 Let agent 2 choose her preferred bundle between X_1 and X_2 , and agent 1 get the other bundle. Denote by $\mathcal{A} = (A_1, A_2)$ the resulting allocation.
- 3 **return** Allocation \mathcal{A}

LEMMA 3.3. *Given any fair division instance $\langle [2], M, D, \mathbf{u} \rangle$, Algorithm 1 computes an EFXM allocation $\mathcal{A} = (A_1, A_2)$ with social welfare $\text{SW}(\mathcal{A}) \geq \frac{u_1(M \cup D) + u_2(M \cup D)}{2}$. If $D = \emptyset$, \mathcal{A} is EFX.*

PROOF. For ease of exposition, we assume without loss of generality that $u_1(X_1) \geq u_1(X_2)$ and $u_2(Y_1) \geq u_2(Y_2)$. Then, according to Algorithm 1, we have

$$u_1(X_1) - u_1(X_2) \leq u_2(Y_1) - u_2(Y_2). \quad (1)$$

Put differently, agent 1's partition of the mixed goods is more equal.

We now show that $\text{SW}(\mathcal{A}) \geq \frac{u_1(M \cup D) + u_2(M \cup D)}{2}$. First, we have $u_1(A_1) \geq u_1(X_2)$. Second, we have $u_2(A_2) \geq u_2(Y_1)$; otherwise, agent 2 could have a more equal partition of the mixed goods, a contradiction to our assumption. The social welfare of allocation \mathcal{A} is lower bounded by

$$\text{SW}(\mathcal{A}) = u_1(A_1) + u_2(A_2) \geq u_1(X_2) + u_2(Y_1) \geq u_1(X_1) + u_2(Y_2),$$

where the last transition is due to Equation (1). It implies that

$$\begin{aligned} \text{SW}(\mathcal{A}) &\geq \frac{u_1(X_2) + u_2(Y_1) + u_1(X_1) + u_2(Y_2)}{2} \\ &= \frac{u_1(M \cup D) + u_2(M \cup D)}{2}, \end{aligned}$$

as desired.

Finally, we show that allocation \mathcal{A} is EFXM. Agent 2 gets her preferred bundle, so she is envy-free and hence EFXM. Regarding agent 1, she is envy-free (and hence EFM) if she receives bundle X_1 . In the case that agent 1 gets bundle X_2 , if agent 1 still has envy after removing some indivisible good or some amount of divisible goods from bundle X_1 , then, by moving the good to bundle X_2 , agent 1 could have created a more equal partition, a contradiction. As a result, allocation \mathcal{A} is EFXM. When $D = \emptyset$, this implies that allocation \mathcal{A} is EFX. \square

We are now ready to show the tight bounds on price of EFX / EFM / EFXM for two agents, and start with the case of agents having unscaled utilities.

THEOREM 3.4. *For $n = 2$ and unscaled utilities, the price of EFX / EFM / EFXM is 2.*

PROOF. The lower bound of 2 for both the price of EFX (and thus EFXM) and the price of EFM (note that EFM generalizes EF1) follows from Theorem F.4 of Bu et al. [24].

We now show the matching upper bound. Consider an arbitrary instance $I = \langle [2], M, D, \mathbf{u} \rangle$. It is easy to see that $\text{OPT}(I) \leq u_1(M \cup D) + u_2(M \cup D)$. Together with Lemma 3.3, we conclude that the price of these three fairness notions is at most 2. \square

We next show the price of EFM / EFXM for two agents with scaled utilities.

THEOREM 3.5. *For $n = 2$ and scaled utilities, the price of EFM and the price of EFXM is $\frac{3}{2}$.*

PROOF. *Lower bound:* Consider the following instance with one indivisible good g_1 and two homogeneous divisible goods d_1, d_2 , and assume that the utilities are as follows:

	g_1	d_1	d_2
Agent 1's value	1/2	1/2 - ϵ	ϵ
Agent 2's value	1/2	ϵ	1/2 - ϵ

The optimal social welfare is $3/2 - 2\epsilon$, achieved by assigning goods g_1 and d_1 to agent 1, and good d_2 to agent 2. On the other hand, in any EFM allocation, no agent can get both the indivisible good g_1 and any positive amount of the divisible goods. Hence, the social welfare of an EFM allocation is at most 1. Taking $\epsilon \rightarrow 0$, we find that the price of EFM is at least $3/2$. Since EFXM is stronger than EFM, this also implies that the price of EFXM is at least $3/2$.

Upper bound: Consider an arbitrary instance. If in an optimal allocation both agents get utility at least $1/2$, this allocation is envy-free (due to the assumptions of additive and scaled utilities) and hence EFM and EFXM; therefore, in this case, the price of EFM is 1. Otherwise, the maximum social welfare is at most $1 + 1/2 = 3/2$. According to Lemma 3.3, Algorithm 1 returns an EFXM (and thus EFM) allocation \mathcal{A} with $\text{SW}(\mathcal{A}) \geq \frac{u_1(M \cup D) + u_2(M \cup D)}{2}$. Since utilities are scaled, we have $\text{SW}(\mathcal{A}) \geq 1$, implying that the price of EFXM and the price of EFM is at most $3/2$. \square

4 ARBITRARY NUMBER OF AGENTS

In this section, we establish asymptotically tight bounds on the price of EFM and the price of EFXM for n agents, and begin with the case that agents' valuations are scaled.

THEOREM 4.1. *For scaled utilities, the price of EFM and the price of EFXM are $\Theta(\sqrt{n})$.*

PROOF. Since EFM generalizes EF1 and EFXM implies EFM, the lower bound $\Omega(\sqrt{n})$ follows from Bei et al. [18].

To show the upper bound $O(\sqrt{n})$, we make use of the result that the price of EFX is $\Theta(\sqrt{n})$ shown by Bu et al. [24]. The high-level idea is as follows. We first split each divisible good d_k into ℓ smaller goods $d_k^1, d_k^2, \dots, d_k^\ell$ of equal size, and we treat each of the ℓ smaller goods as an indivisible good. In other words, we are considering an instance I^ℓ with a total of $m + \bar{m}\ell$ indivisible goods. For each ℓ , we find an EFX allocation that achieves an $O(\sqrt{n})$ -approximation to $\text{OPT}(I^\ell)$. When $\ell \rightarrow \infty$, we have a sequence of EFX allocations

which converges to a “limit allocation” that is EFXM (and thus EFM), and the limit allocation exists due to the compactness of the allocation space. This limit allocation characterizes the upper bound $O(\sqrt{n})$ for the price of EFXM, as the social welfare is a continuous function on the allocation space.

To prove the upper bound formally, we start from defining the instance I^ℓ and the allocation \mathcal{A}^ℓ . In the instance I^ℓ , we have the same set of agents N , and a set of $m + \bar{m}\ell$ indivisible goods which consist of the m goods in M and $\bar{m}\ell$ goods $\{d_k^1, d_k^2, \dots, d_k^\ell\}_{k=1, \dots, \bar{m}}$ as described earlier. The result of Bu et al. [24] indicates that there exists a (partial) EFX allocation \mathcal{A} for instance I^ℓ such that $\frac{\text{OPT}(I^\ell)}{\text{SW}(\mathcal{A})} \leq c\sqrt{n}$ for some constant c and sufficiently large n . Let \mathcal{A}^ℓ be such an allocation. Notice that \mathcal{A}^ℓ is also a valid allocation for the original instance I (where each d_k is divisible), and we will use \mathcal{A}^ℓ for the same allocation in both I^ℓ and I .

Next, we define an allocation \mathcal{A} for the original instance I which is a “limit allocation” for the allocation sequence $\{\mathcal{A}^\ell\}_{\ell=1}^\infty$. To make the notion of limit valid, we need to define a metric space for the set of all allocations, and this is defined in the following natural way. First note that there are $(n+1)^m$ ways to allocate the indivisible goods (each good can be allocated to one of the n agents, or unallocated), which is finite. For each fixed allocation of the indivisible goods, an allocation of the divisible goods $\{d_1, d_2, \dots, d_m\}$ can be naturally described by a point in the following subset of \mathbb{R}^{nm} :

$$\chi = \left\{ (x_{i\bar{k}})_{i=1, \dots, n; \bar{k}=1, \dots, \bar{m}} \in \mathbb{R}^{n\bar{m}} : \sum_{i=1}^n x_{i\bar{k}} \leq 1 \text{ for each } \bar{k} \in [\bar{m}], \right. \\ \left. \text{and } x_{i\bar{k}} \geq 0 \text{ for each } i \in [n] \text{ and } \bar{k} \in [\bar{m}] \right\}.$$

Given two allocations, the distance between them in the metric space is defined as follows:

- if their corresponding allocations for M are different, the distance is ∞ ;
- if their corresponding allocations for M are the same, the distance is defined by the Euclidean distance of the two points in χ describing their allocations for D .

Since χ is closed and bounded and the space of all allocations is a union of finitely many $((n+1)^m$ to be precise) such closed and bounded sets, the Bolzano-Weierstrauss Theorem [11] implies that the allocation space contains at least one allocation that is a limit point for the sequence $\{\mathcal{A}^\ell\}_{\ell=1}^\infty$. Let \mathcal{A} be one such limit point. In the remaining part of the proof, we will conclude the theorem by showing that 1) \mathcal{A} is an EFXM allocation and 2) it satisfies the approximation guarantee $\frac{\text{OPT}(I)}{\text{SW}(\mathcal{A})} = O(\sqrt{n})$.

\mathcal{A} is EFXM. Suppose this is not the case. There exist two agents i and j such that $u_i(A_i) < u_i(A_j)$ and A_j contains some divisible good (i.e., $x_{j\bar{k}} > 0$ for some \bar{k}). We choose a sufficiently small value $\delta_{\bar{k}}$ such that $3\delta_{\bar{k}} \in (0, x_{j\bar{k}})$ and $u_i(A_j) - u_i(A_i) > 3\delta_{\bar{k}} \cdot u_i(d_{\bar{k}})$. In other words, removing an amount $3\delta_{\bar{k}}$ of good $d_{\bar{k}}$ from A_j will not stop agent i from envying agent j . Since u_i is a continuous function (which can be proved by a straightforward application of the definition of continuity given our definition of the metric space) and \mathcal{A} is a limit point of the sequence $\{\mathcal{A}^\ell\}_{\ell=1}^\infty$, by considering a

sufficiently small neighbourhood of \mathcal{A} , there exists ℓ with allocation $\mathcal{A}^\ell = (A_1^\ell, \dots, A_n^\ell)$ such that

- (1) $|u_i(A_i) - u_i(A_i^\ell)| < \delta_{\bar{k}} \cdot u_i(d_{\bar{k}})$,
- (2) $|u_i(A_j) - u_i(A_j^\ell)| < \delta_{\bar{k}} \cdot u_i(d_{\bar{k}})$, and
- (3) $\delta_{\bar{k}} > \frac{1}{\ell}$.

Points 1 and 2 above imply $u_i(A_j^\ell) - u_i(A_i^\ell) > \delta_{\bar{k}} \cdot u_i(d_{\bar{k}})$ under the condition that $u_i(A_j) - u_i(A_i) > 3\delta_{\bar{k}} \cdot u_i(d_{\bar{k}})$. Point 3 further implies that, in the instance I^ℓ , there exists an indivisible item corresponding to a small portion that is smaller than $\delta_{\bar{k}}$ of $d_{\bar{k}}$ whose removal will not stop agent i from envying agent j . This contradicts to our construction that \mathcal{A}^ℓ is EFX.

Approximation guarantee. By our construction of the sequence with the result of Bu et al. [24], for sufficiently large n and a fixed constant c , we have $\frac{\text{OPT}(I^\ell)}{\text{SW}(\mathcal{A}^\ell)} \leq c\sqrt{n}$ for every ℓ . It suffices to show

$$\text{OPT}(I) = \lim_{\ell \rightarrow \infty} \text{OPT}(I^\ell) \quad \text{and} \quad \text{SW}(\mathcal{A}) = \lim_{\ell \rightarrow \infty} \text{SW}(\mathcal{A}^\ell).$$

The second limit follows from the continuity of the function $\text{SW}(\cdot)$, where the continuity can be proved by a straightforward application of the definition of continuity. The first limit follows from the fact that $\text{OPT}(I) = \text{OPT}(I^\ell)$ for each ℓ . To see this, in the optimal allocation, we allocate each divisible good d_k as a whole to a single agent who values it the highest (with tie broken arbitrarily), so it does not matter how each divisible good is sub-divided to multiple indivisible smaller goods. \square

We next show the price of EFM/ EFXM for unscaled utilities.

THEOREM 4.2. *For unscaled utilities, the price of EFM and the price of EFXM are $\Theta(n)$.*

Theorem 4.2 can be proved by using the result that the price of EFX for unscaled valuations is $\Theta(n)$ from Bu et al. [24] and applying the discretization-with-limit technique used in the proof of Theorem 4.1. Here, we give an alternative constructive proof.

When allocating indivisible goods, Lemma 1 of Barman et al. [8] proved that there always exists an EF1 allocation with an absolute welfare guarantee. We show a similar result holds when allocating mixed goods. To be more specific, by slightly tweaking Algorithm 1 (ALG-EF1-ABS) of Barman et al. [8], we can compute a partial EFXM allocation with a similar absolute welfare guarantee:

LEMMA 4.3. *Given any fair division instance (N, M, D, \mathbf{u}) , Algorithm 2 computes a partial EFXM allocation \mathcal{A} with social welfare $\text{SW}(\mathcal{A}) \geq \frac{1}{2n+1} \cdot \sum_{i \in N} u_i(M \cup D)$.*

The proof can be found in the full version of our paper [40]. We give some intuition as follows. At a high level, Barman et al.’s algorithm starts from a maximum weight matching where each agent receives exactly one indivisible good, and then performs the envy-cycle elimination procedure of Lipton et al. [41]. At the end, an EF1 allocation (A_1, A_2, \dots, A_n) of indivisible goods is obtained. In many “natural scenarios”, we have $u_i(A_i) \geq \frac{1}{2}u_i(A_j)$ as removing one indivisible good from A_j eliminates the envy from i to j . This gives us the $2n$ approximation ratio to the optimal social welfare. The inequality $u_i(A_i) \geq \frac{1}{2}u_i(A_j)$ can only fail in the case when the removed indivisible good g from A_j is “large” so that $u_i(\{g\}) > u_i(A_j \setminus \{g\})$. However, the initial allocation with

Algorithm 2: A partial EFXM allocation with an absolute welfare guarantee

Input: Fair division instance $\langle N, M, D, \mathbf{u} \rangle$.

Output: A partial EFXM allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ with social welfare

$$SW(\mathcal{A}) \geq \frac{1}{2n+1} \cdot \sum_{i \in N} u_i(M \cup D).$$

- 1 **if** $|M| \geq n$ **then**
 - 2 $\tilde{M} \leftarrow M$
 - 3 **else**
 - 4 Let \tilde{M} be the union of indivisible goods M and some dummy indivisible goods, for which each agent has value zero, such that $|\tilde{M}| = n$.
 - 5 Consider the weighted bipartite graph $G = (N \cup \tilde{M}, N \times \tilde{M})$ with weight of each edge $(i, g) \in N \times \tilde{M}$ setting as $u_i(g)$. Let π be a maximum-weight matching in G that matches all nodes in N .
 - 6 Construct the partial allocation $\mathcal{A}' = (A'_1, A'_2, \dots, A'_n)$ such that $A'_i = \{\pi(i)\}$ for each $i \in N$.
 - 7 Use Algorithm 2.1 of Chaudhury et al. [32] to extend allocation \mathcal{A}' by allocating the remaining indivisible goods and obtain a partial EFX allocation.
 - 8 Use Algorithm 1 of Bei et al. [15] to allocate the divisible goods and obtain a partial EFXM allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$, where $A_i = (\tilde{M}_i, \mathbf{x}_i)$.
 - 9 **return** Allocation \mathcal{A}
-

the maximum weight matching ensures that the “large indivisible goods” are “reasonably allocated” so that the inequality holds in some average sense. Barman et al. worked out the calculations to make the approximation guarantee $2n$ hold, and we find out that these arguments can be extended to the setting with mixed goods.

PROOF OF THEOREM 4.2. Since EFM generalizes EF1, the desired lower bound of $\Omega(n)$ follows from Theorem 1 of Barman et al. [8]. Since EFXM implies EFM, we obtain the same lower bound of $\Omega(n)$ for the price of EFXM.

We now show the asymptotic matching upper bound for the price of EFXM. Consider an arbitrary instance where $\sum_{i \in N} u_i(M \cup D)$ is a trivial upper bound on the optimal social welfare. Thus Lemma 4.3 implies the desired upper bound of $O(n)$ for the price of EFXM.

Next, we establish the asymptotic matching upper bound for the price of EFM. Given that EFXM implies EFM, we can readily deduce an upper bound of $O(n)$ for the price of EFM. Moreover, if we want to find a complete EFM allocation, we can adapt line 8 of Algorithm 2 in the following way:

- Use Algorithm 1 of Bei et al. [15] to extend the partial EFM allocation \mathcal{A}' by allocating both the remaining indivisible goods and divisible goods into a complete EFM allocation.

Following a similar argument, we can conclude that $SW(\mathcal{A}) \geq \frac{1}{2n} \cdot \sum_{i \in N} u_i(M \cup D)$. Thus, the price of EFM is also $\Theta(n)$. \square

5 CONCLUSION AND FUTURE WORK

In this paper, we have given a complete characterization for the price of envy-freeness in various settings. The bounds we provide

Table 3: An example where EFM is incompatible with Pareto-optimality. Notice that one of the agents needs to receive both d_1 and d_2 to guarantee EFM, but this violates Pareto-optimality. However, this counterexample fails if 0 in the table is replaced by a very small positive value $\varepsilon > 0$.

	g_1	d_1	d_2
Agent 1’s value	1	1/2	0
Agent 2’s value	1	0	1/2

are tight for two agents and asymptotically tight for any number of agents. In particular, we close a gap left open in [18] by showing a tight bound for the price of EF1 for two agents. Furthermore, the price of fairness has been studied for the setting with divisible goods and the setting with indivisible goods, but it is much less understood for allocating mixed divisible and indivisible goods. This paper fills in this missing piece.

For future work, we list two open problems about allocation efficiency which are yet to be understood.

Compatibility between (a Variant of) EFM and Pareto-Optimality. As we mentioned in Section 1.1, Pareto-optimality is compatible with EF1 for the setting with indivisible goods and is compatible with EF for the setting with divisible goods. However, the compatibility of Pareto-optimality with (a variant of) EFM for the setting with mixed divisible and indivisible goods is still largely unknown.³ Although a simple counterexample is shown in Table 3, it is a rather corner case that only happens when an agent has a zero valuation on a divisible good. What if we assume the valuations are positive? What if we consider an arguably more natural relaxed version of EFM where i is allowed to envy j if j ’s bundle only contains divisible goods that has a total value of zero to agent i and the envy can be eliminated by removing one indivisible item from j ’s bundle?

Resource Monotonicity. In Section 2.1, we have mentioned that it is not known whether the resource monotonicity holds for EFM: we do not know if any partial EFM allocation can be extended to a complete EFM allocation with a weakly higher social welfare. Although we have argued that our definition for the price of EFM is more natural by including partial EFM allocations in both cases whether or not resource monotonicity holds, we believe resource monotonicity is an interesting problem by itself.

ACKNOWLEDGMENTS

Biaoshuai Tao was supported by the National Natural Science Foundation of China (No. 62102252). Shengxin Liu was partially supported by the National Natural Science Foundation of China (No. 62102117), by the Shenzhen Science and Technology Program (No. RCBS20210609103900003 and GXWD20231129111306002), and by the Guangdong Basic and Applied Basic Research Foundation (No. 2023A1515011188). Xinhang Lu was partially supported by ARC Laureate Project FL200100204 on “Trustworthy AI”.

³We refer the interested readers to the Section 6 of Bei et al. [15] for more discussions on their preliminary (in)compatibility results.

REFERENCES

- [1] Hannaneh Akrami, Noga Alon, Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, and Ruta Mehta. 2023. EFX: A Simpler Approach and an (Almost) Optimal Guarantee via Rainbow Cycle Number. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 61.
- [2] Noga Alon. 1987. Splitting Necklaces. *Advances in Mathematics* 63, 3 (1987), 247–253.
- [3] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. 2023. Fair Division of Indivisible Goods: Recent Progress and Open Questions. *Artificial Intelligence* 322 (2023), 103965.
- [4] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A. Voudouris. 2021. Maximum Nash Welfare and Other Stories About EFX. *Theoretical Computer Science* 863 (2021), 69–85.
- [5] Yonatan Aumann and Yair Dombb. 2015. The Efficiency of Fair Division with Connected Pieces. *ACM Transactions on Economics and Computation* 3, 4 (2015), 1–16.
- [6] Yonatan Aumann, Yair Dombb, and Avinatan Hassidim. 2013. Computing Socially-Efficient Cake Divisions. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 343–350.
- [7] Haris Aziz, Xin Huang, Nicholas Mattei, and Erel Segal-Halevi. 2023. Computing Welfare-Maximizing Fair Allocations of Indivisible Goods. *European Journal of Operational Research* 307, 2 (2023), 773–784.
- [8] Siddharth Barman, Umang Bhaskar, and Nisarg Shah. 2020. Optimal Bounds on the Price of Fairness for Indivisible Goods. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*. 356–369.
- [9] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. 2019. Fair Division of Indivisible Goods Among Strategic Agents. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 1811–1813.
- [10] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018. Finding Fair and Efficient Allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. 557–574.
- [11] Robert G. Bartle and Donald R. Sherbert. 2000. *Introduction to Real Analysis*. Vol. 2. Wiley New York.
- [12] Xiaohui Bei, Ning Chen, Xia Hua, Biaoshuai Tao, and Endong Yang. 2012. Optimal Proportional Cake Cutting with Connected Pieces. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*. 1263–1269.
- [13] Xiaohui Bei, Ning Chen, Guangda Huzhang, Biaoshuai Tao, and Jiajun Wu. 2017. Cake Cutting: Envy and Truth. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 3625–3631.
- [14] Xiaohui Bei, Guangda Huzhang, and Warut Suksompong. 2020. Truthful Fair Division without Free Disposal. *Social Choice and Welfare* 55, 3 (2020), 523–545.
- [15] Xiaohui Bei, Zihao Li, Jinyan Liu, Shengxin Liu, and Xinhang Lu. 2021. Fair Division of Mixed Divisible and Indivisible Goods. *Artificial Intelligence* 293 (2021), 103436.
- [16] Xiaohui Bei, Shengxin Liu, and Xinhang Lu. 2023. Fair Division with Subjective Divisibility. In *Proceedings of the 19th Conference on Web and Internet Economics (WINE)*. 677.
- [17] Xiaohui Bei, Shengxin Liu, Xinhang Lu, and Hongao Wang. 2021. Maximin Fairness with Mixed Divisible and Indivisible Goods. *Autonomous Agents and Multi-Agent Systems* 35, 2 (2021), 34:1–34:21.
- [18] Xiaohui Bei, Xinhang Lu, Pasin Manurangsi, and Warut Suksompong. 2021. The Price of Fairness for Indivisible Goods. *Theory of Computing Systems* 65, 7 (2021), 1069–1093.
- [19] Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. 2022. Almost Full EFX Exists for Four Agents. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. 4826–4833.
- [20] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. 2011. The Price of Fairness. *Operations Research* 59, 1 (2011), 17–31.
- [21] Umang Bhaskar, A. R. Sricharan, and Rohit Vaish. 2021. On Approximate Envy-Freeness for Indivisible Chores and Mixed Resources. In *Proceedings of the 24th International Conference on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*. 1:1–1:23.
- [22] Steven J. Brams, Michal Feldman, John Lai, Jamie Morgenstern, and Ariel D. Procaccia. 2012. On Maxsum Fair Cake Divisions. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*. 1285–1291.
- [23] Xiaolin Bu, Zihao Li, Shengxin Liu, Jiabin Song, and Biaoshuai Tao. 2023. Fair Division with Allocator’s Preference. In *Proceedings of the 19th Conference on Web and Internet Economics (WINE)*. 77–94.
- [24] Xiaolin Bu, Zihao Li, Shengxin Liu, Jiabin Song, and Biaoshuai Tao. 2023. On the Complexity of Maximizing Social Welfare Within Fair Allocations of Indivisible Goods. *CoRR abs/2205.14296* (2023).
- [25] Xiaolin Bu, Jiabin Song, and Biaoshuai Tao. 2023. On Existence of Truthful Fair Cake Cutting Mechanisms. *Artificial Intelligence* 319 (2023), 103904.
- [26] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- [27] Ioannis Caragiannis, Nick Gravin, and Xin Huang. 2019. Envy-Freeness Up to Any Item with High Nash Welfare: The Virtue of Donating Items. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*. 527–545.
- [28] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2012. The Efficiency of Fair Division. *Theory of Computing Systems* 50, 4 (2012), 589–610.
- [29] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation* 7, 3 (2019), 12:1–12:32.
- [30] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. 2020. EFX Exists for Three Agents. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 1–19.
- [31] Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta, and Pranabendu Misra. 2021. Improving EFX Guarantees Through Rainbow Cycle Number. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. 310–311.
- [32] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. 2021. A Little Charity Guarantees Almost Envy-Freeness. *SIAM J. Comput.* 50, 4 (2021), 1336–1358.
- [33] Yiling Chen, John K. Lai, David C. Parkes, and Ariel D. Procaccia. 2013. Truth, justice, and cake cutting. *Games and Economic Behavior* 77, 1 (2013), 284–297.
- [34] George Christodoulou, Amos Fiat, Elias Koutsoupias, and Alkmini Sgouritsa. 2023. Fair Allocation in Graphs. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 473–488.
- [35] Yuga J. Coehler, John K. Lai, David C. Parkes, and Ariel D. Procaccia. 2011. Optimal Envy-Free Cake Cutting. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*. 626–631.
- [36] Duncan Karl Foley. 1967. Resource Allocation and the Public Sector. *Yale Economics Essays* 7, 1 (1967), 45–98.
- [37] Jugal Garg and Aniket Murhekar. 2023. Computing Fair and Efficient Allocations with Few Utility Values. *Theoretical Computer Science* 962 (2023), 113932.
- [38] Yasushi Kawase, Koichi Nishimura, and Hanna Sumita. 2023. Fair Allocation with Binary Valuations for Mixed Divisible and Indivisible Goods. *CoRR abs/2306.05986* (2023).
- [39] Zihao Li, Shengxin Liu, Xinhang Lu, and Biaoshuai Tao. 2023. Truthful Fair Mechanisms for Allocating Mixed Divisible and Indivisible Goods. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*. 2808–2816.
- [40] Zihao Li, Shengxin Liu, Xinhang Lu, Biaoshuai Tao, and Yichen Tao. 2024. A Complete Landscape for the Price of Envy-Freeness. *CoRR abs/2401.01516* (2024).
- [41] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*. 125–131.
- [42] Shengxin Liu, Xinhang Lu, Mashbat Suzuki, and Toby Walsh. 2024. Mixed Fair Division: A Survey. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. Senior Member Presentation Track. Forthcoming.
- [43] Vijay Menon and Kate Larson. 2017. Deterministic, Strategyproof, and Fair Cake Cutting. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 352–358.
- [44] Aniket Murhekar and Jugal Garg. 2021. On Fair and Efficient Allocations of Indivisible Goods. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5595–5602.
- [45] Koichi Nishimura and Hanna Sumita. 2023. Envy-freeness and Maximum Nash Welfare for Mixed Divisible and Indivisible Goods. *CoRR abs/2302.13342v2* (2023).
- [46] Benjamin Plaut and Tim Roughgarden. 2020. Almost Envy-Freeness with General Valuations. *SIAM Journal on Discrete Mathematics* 34, 2 (2020), 1039–1068.
- [47] Erel Segal-Halevi and Balázs R. Sziklai. 2018. Resource-monotonicity and population-monotonicity in connected cake-cutting. *Mathematical Social Sciences* 95 (2018), 19–30.
- [48] Erel Segal-Halevi and Balázs R. Sziklai. 2019. Monotonicity and Competitive Equilibrium in Cake-Cutting. *Economic Theory* 68, 2 (2019), 363–401.