



Contents lists available at ScienceDirect

## Artificial Intelligence

www.elsevier.com/locate/artint

Fair division of mixed divisible and indivisible goods <sup>☆</sup>Xiaohui Bei <sup>a,\*</sup>, Zihao Li <sup>a</sup>, Jinyan Liu <sup>b</sup>, Shengxin Liu <sup>a</sup>, Xinhang Lu <sup>a</sup><sup>a</sup> School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore<sup>b</sup> School of Computer Science and Technology, Beijing Institute of Technology, PR China

## ARTICLE INFO

## Article history:

Received 5 March 2020

Received in revised form 5 December 2020

Accepted 13 December 2020

Available online 5 January 2021

## Keywords:

Fair division

Resource allocation

Envy-freeness

Social choice

## ABSTRACT

We study the problem of fair division when the set of resources contains both divisible and indivisible goods. Classic fairness notions such as envy-freeness (EF) and envy-freeness up to one good (EF1) cannot be directly applied to this mixed goods setting. In this work, we propose a new fairness notion, *envy-freeness for mixed goods (EFM)*, which is a direct generalization of both EF and EF1 to the mixed goods setting. We prove that an EFM allocation always exists for any number of agents with additive valuations. We also propose efficient algorithms to compute an EFM allocation for two agents with general additive valuations and for  $n$  agents with piecewise linear valuations over the divisible goods. Finally, we relax the envy-freeness requirement, instead asking for  $\epsilon$ -*envy-freeness for mixed goods ( $\epsilon$ -EFM)*, and present an efficient algorithm that finds an  $\epsilon$ -EFM allocation.

© 2020 Elsevier B.V. All rights reserved.

## 1. Introduction

Fair division studies the allocation of scarce resources among interested agents, with the objective of finding an allocation that is fair to all participants involved. Initiated by Steinhaus [43], the study of fair division has since been attracting interest from various disciplines for decades, including among others, mathematics, economics, and computer science [13,39,33,46,34].

The literature of fair division can be divided into two classes, categorized by the type of the resources to be allocated. The first class assumes the resource to be heterogeneous and infinitely divisible. The corresponding problem is commonly known as *cake cutting*. One of the most prominent fairness notions in this setting is *envy-freeness (EF)*. An allocation is said to be envy-free if each agent weakly prefers her own bundle to any other bundle in the allocation. An envy-free allocation with divisible resources always exists [2,45] and can be found via a discrete and bounded protocol [5].

The second class considers the fair allocation of *indivisible* goods. Note that an envy-free allocation may fail to exist in the indivisible goods setting.<sup>1</sup> To circumvent this problem, relaxations of envy-freeness have been studied. One of the commonly considered relaxations is *envy-freeness up to one good (EF1)* [30,16]. An allocation is said to satisfy EF1 if it is possible to eliminate any envy one agent has towards another agent by removing some good from the latter's bundle. An EF1 allocation with indivisible goods always exists and can be found in polynomial time [30,17].

<sup>☆</sup> This paper is an invited revision of a paper which first appeared at the 2020 AAAI Conference on Artificial Intelligence (AAAI-20) [10]. Compared to the conference version, this journal version fixes a bug in the proof of Theorem 3.6 and includes a new section (Section 6) that discusses how to combine the newly proposed fairness notion together with economic efficiency considerations.

\* Corresponding author.

E-mail addresses: xhbei@ntu.edu.sg (X. Bei), zihao004@e.ntu.edu.sg (Z. Li), jyliu.tcs@gmail.com (J. Liu), sxliu@ntu.edu.sg (S. Liu), xinhang001@e.ntu.edu.sg (X. Lu).

<sup>1</sup> Consider the case where there are two agents but only a single valuable good to be allocated.

The vast majority of the fair division literature assumes that the resources either are completely divisible, or consist of only indivisible goods. However, this is not always the case in many real-world scenarios. In inheritance division, for example, the inheritances to be divided among the heirs may contain divisible goods such as land and money, as well as indivisible goods such as houses, cars, and artworks. What fairness notion should one adopt when dividing such mixed type of resources? While EF and EF1 both work well in their respective settings, neither of them can be directly applied to this more general scenario. On the one hand, an EF allocation may not exist, when, for example, all goods are indivisible. On the other hand, the EF1 notion in the mixed goods setting, when interpreted as that each agent does not envy another agent after removing at most one indivisible good from the latter agent's bundle, may also produce unfair allocations. Consider the example where there is an indivisible good and a cake that are both equally valued by two agents. The allocation that divides the cake in half and then gives the indivisible good to one of the agents is EF1 but is arguably unfair. Another tempting solution is to divide the divisible and indivisible resources using EF and EF1 protocols separately and independently, and then combine the two allocations together. This approach, however, also has problems. Consider a simple example where two agents need to divide a cake *and* an indivisible item. EF1 requires to allocate the indivisible item to one of the agents, say agent 1 for example. However, if we then divide the cake using an arbitrary EF allocation, the overall allocation might be unfair to agent 2 who does not receive the indivisible item. In fact, if the whole cake is valued less than the item, it would make more sense to allocate the cake entirely to agent 2. When the cake is valued more than the item, it is still a fairer solution to allocate more cake to agent 2 in order to compensate her disadvantage in the indivisible resource allocation. This demonstrates that it is not straightforward to generalize EF and EF1 to the mixed goods setting. Dividing mixed types of resources calls for a new fairness notion that could unify EF and EF1 together to the new setting in a natural and non-trivial way.

### 1.1. Our results

In this work, we initiate the study of fair division with mixed types of resources. More specifically, we propose a new fairness notion, denoted as *envy-freeness for mixed goods* (or *EFM* for short), that naturally combines EF and EF1 together and works for the setting where the set of resources may contain both divisible and indivisible goods. Intuitively, EFM requires that for each agent, if her allocation consists of only indivisible items, then other agents will compare their bundles to hers using the EF1 criterion; but if this agent's bundle contains *any* positive amount of divisible resources, other agents will compare their bundles to hers using the stricter EF condition. This definition generalizes both EF and EF1 to the mixed goods setting and strikes a natural balance between the two fairness notions.

In Section 3, we first show that with mixed types of goods, an EFM allocation always exists for any number of agents with additive valuations. Our proof is constructive and gives an algorithm for computing such an EFM allocation. The algorithm requires an oracle for computing a *perfect allocation in cake cutting* and can compute an EFM allocation in a polynomial number of steps. In addition, in Section 4, we present two algorithms that could compute an EFM allocation for two special cases without using the perfect allocation oracle: (1) two agents with general additive valuations in the Robertson-Webb model, and (2) any number of agents with piecewise linear valuation functions.

While it is still unclear to us whether in general an EFM allocation can be computed in a finite number of steps in the Robertson-Webb model, in Section 5, we turn our attention to approximations and define the notion of  $\epsilon$ -EFM. We then give an algorithm to compute an  $\epsilon$ -EFM allocation in the Robertson-Webb model with running time *polynomial* in the number of agents  $n$ , the number of indivisible goods  $m$ , and  $1/\epsilon$ , and query complexity *polynomial* in  $n$  and  $1/\epsilon$ . We note that this algorithm does not require a perfect allocation oracle. This is an appealing result in particular due to its polynomial running time complexity. A bounded exact EFM protocol, even if exists, is likely to require a large number of queries and cuts. This is because in the special case when resources are all divisible, EFM reduces to EF in cake cutting, for which the best known protocol [5] has a very high query complexity (a tower of exponents of  $n$ ). This result shows that if one is willing to allow a small margin of errors, such an allocation could be found much more efficiently.

Finally, in Section 6 we discuss EFM in conjunction with efficiency considerations. In particular, to one's surprise, we show that EFM and Pareto optimality (PO) are incompatible. We also propose a weaker version of EFM and discuss the possibilities and difficulties in combining it with PO.

### 1.2. Related work

As we mentioned, most previous works in fair division are from two categories based on whether the resources to be allocated are divisible or indivisible.

When the resources are divisible, the existence of an envy-free allocation is guaranteed [29,21], even with only  $n - 1$  cuts [44,45]. Brams and Taylor [12] gave the first finite (but unbounded) envy-free protocol for any number of agents. Recently, Aziz and Mackenzie [6] gave the first bounded protocol for computing an envy-free allocation with four agents and their follow-up work extended the result to any number of agents [5]. Besides envy-freeness, other classic fairness notions include *proportionality* and *equitability*, both of which have been studied extensively [21,23,22,18,38].

When the resources are indivisible, none of the aforementioned fairness notions is guaranteed to exist, thus relaxations are considered. Among other notions, these include envy-freeness up to one good (EF1), envy-freeness up to any good (EFX), maximin share (MMS), etc. [30,16,17]. An EF1 allocation always exists and can be efficiently computed [30,17]. However, the

existence of an EFX allocation is still open [37], except for several special cases [35,19,4]. As for MMS, an MMS allocation may not always exist; however, an approximation of MMS always exists and can be efficiently computed [28,3,25,24]. A recent paper by Bei et al. [11] also studied the existence, approximation and computation of MMS allocations in the mixed goods setting.

In addition, several works studied fair division with the assumption that resources can be shared among agents. The adjusted-winner (AW) procedure proposed by Brams and Taylor [13] ensures that at most one good must be split in a fair and (economically) efficient division between two agents. Sandomirskiy and Segal-Halevi [41] focused on obtaining a fair and efficient division with minimum number of objects shared between two or more agents. Rubchinsky [40] considered the fair division problem between two agents with both divisible and indivisible items, and introduced three fairness notions with computationally efficient algorithms for finding them. All of the works discussed above assumed that divisible items are *homogeneous*.

Several other works studied the allocation of both indivisible goods and money, with the goal of finding envy-free allocations [31,1,27,32,26,15]. Money can be viewed as a homogeneous divisible good which is valued the same across all agents. In our work, we consider a more general setting with heterogeneous divisible goods. Moreover, these works focused on finding envy-free allocations with the help of a sufficient amount of money, which is again different from our goal in the sense that our method could also be used even in cases where the money is insufficient.

## 2. Preliminaries

We consider a resource allocation setting with both divisible and indivisible goods (mixed goods for short). Denote by  $N = \{1, 2, \dots, n\}$  the set of agents,  $M = \{1, 2, \dots, m\}$  the set of indivisible goods, and  $D = \{D_1, D_2, \dots, D_\ell\}$  the set of  $\ell$  heterogeneous divisible goods or *cakes*. Since the fairness notion we propose below does not distinguish pieces from different cakes, without loss of generality, we assume each cake  $D_i$  is represented by the interval  $[\frac{i-1}{\ell}, \frac{i}{\ell}]$ ,<sup>2</sup> and use a single cake  $C = [0, 1]$  to represent the union of all cakes.<sup>3</sup>

Each agent  $i$  has a non-negative utility  $u_i(g)$  for each indivisible good  $g \in M$ . Agents' utilities for subsets of indivisible goods are additive, meaning that  $u_i(M') = \sum_{g \in M'} u_i(g)$  for each agent  $i$  and subset of goods  $M' \subseteq M$ . Each agent  $i$  also has a density function  $f_i: [0, 1] \rightarrow \mathbb{R}^+ \cup \{0\}$ , which captures how the agent values different parts of the cake. The value of agent  $i$  over a finite union of intervals  $S \subseteq [0, 1]$  is defined as  $u_i(S) = \int_S f_i dx$ .

Denote by  $\mathcal{M} = (M_1, M_2, \dots, M_n)$  the partition of  $M$  into bundles such that agent  $i$  receives bundle  $M_i$ . Denote by  $\mathcal{C} = (C_1, C_2, \dots, C_n)$  the division of cake  $C$  such that  $C_i \cap C_j = \emptyset$  and agent  $i$  receives  $C_i$ , a union of finitely many intervals. An *allocation* of the mixed goods is defined as  $\mathcal{A} = (A_1, A_2, \dots, A_n)$  where  $A_i = M_i \cup C_i$  is the *bundle* allocated to agent  $i$ . Agent  $i$ 's utility for the allocation is then defined as  $u_i(A_i) = u_i(M_i) + u_i(C_i)$ . We assume without loss of generality that agents' utilities are normalized to 1, i.e.,  $u_i(M \cup C) = 1$  for all  $i \in N$ .

Next, we define the fairness notions used in this paper.

**Definition 2.1 (EF).** An allocation  $\mathcal{A}$  is said to satisfy *envy-freeness (EF)* if for any agents  $i, j \in N$ ,  $u_i(A_i) \geq u_i(A_j)$ .

**Definition 2.2 (EF1).** With indivisible goods, an allocation  $\mathcal{A}$  is said to satisfy *envy-freeness up to one good (EF1)* if for any agents  $i, j \in N$  where  $A_j \neq \emptyset$ , there exists  $g \in A_j$  such that  $u_i(A_i) \geq u_i(A_j \setminus \{g\})$ .

Neither EF nor EF1 alone is a suitable definition for mixed goods. In this paper we introduce the following new fairness notion.

**Definition 2.3 (EFM).** An allocation  $\mathcal{A}$  is said to satisfy *envy-freeness for mixed goods (EFM)* if for any agents  $i, j \in N$ ,

- if agent  $j$ 's bundle consists of only indivisible goods, there exists  $g \in A_j$  such that  $u_i(A_i) \geq u_i(A_j \setminus \{g\})$ ;
- otherwise,  $u_i(A_i) \geq u_i(A_j)$ .

It is easy to see that when the goods are all divisible, EFM reduces to EF; when goods are all indivisible, EFM reduces to EF1. Therefore EFM is a natural generalization of both EF and EF1 to the mixed goods setting.

Next, we define  $\epsilon$ -EFM which is a relaxation of EFM. Note that this definition only relaxes the EF condition for the divisible goods; the EF1 condition is not relaxed.

**Definition 2.4 ( $\epsilon$ -EFM).** An allocation  $\mathcal{A}$  is said to satisfy  *$\epsilon$ -envy-freeness for mixed goods ( $\epsilon$ -EFM)* if for any agents  $i, j \in N$ ,

<sup>2</sup> We assume that agents' valuation functions over the cakes are non-atomic. Thus we can view two consecutive cakes as disjoint even if they intersect at one boundary point.

<sup>3</sup> Sometimes we will use an arbitrary interval  $[a, b]$  to denote the resource for simplicity; this can be easily normalized back to  $[0, 1]$ .

- if agent  $j$ 's bundle consists of only indivisible goods, there exists  $g \in A_j$  such that  $u_i(A_i) \geq u_i(A_j \setminus \{g\})$ ;
- otherwise,  $u_i(A_i) \geq u_i(A_j) - \epsilon$ .

Finally, we describe the Robertson-Webb (RW) query model [39], which is a standard model in cake cutting. In this model, an algorithm is allowed to interact with the agents via two types of queries:

- Evaluation: An evaluation query of agent  $i$  on  $[x, y]$  returns  $u_i([x, y])$ .
- Cut: A cut query of  $\beta$  for agent  $i$  from  $x$  returns a point  $y$  such that  $u_i([x, y]) = \beta$ .

### 3. EFM: existence

Although EFM is a natural generalization of both EF and EF1, it is not straightforward whether an EFM allocation would always exist with mixed goods. In this section, we prove through a constructive algorithm that with mixed goods and any number of agents, an EFM allocation always exists.

We first give some definitions which will be helpful for our algorithm and proofs.

*Perfect allocation* Our algorithm will utilize the concept of *perfect allocation* in cake cutting.

**Definition 3.1** (*Perfect allocation*). A partition  $\mathcal{C} = (C_1, C_2, \dots, C_k)$  of cake  $C$  is said to be *perfect* if for all  $i \in N, j \in [k]$ ,  $u_i(C_j) = u_i(C)/k$ .

Intuitively, a perfect allocation in cake cutting divides the cake into  $k$  pieces, such that every agent in  $N$  values these  $k$  pieces equally. It is known that a perfect allocation always exists for any number of agents and any  $k$  [2]. In the following, we will assume that our algorithm is equipped with an oracle  $\text{PERFECTALLOC}(C, k, N)$  that could return us a perfect allocation for any  $k$  and cake  $C$  among all agents in  $N$ .

*Envy graph and addable set* We also make use of the *envy graph* to capture the envy relation among agents in an allocation.

**Definition 3.2** (*Envy graph*). Given an allocation  $\mathcal{A}$ , its corresponding *envy graph*  $G = (N, E_{\text{envy}} \cup E_{\text{eq}})$  is a directed graph, where each vertex represents an agent, and  $E_{\text{envy}}$  and  $E_{\text{eq}}$  consist of the following two types of edges, respectively:

- Envy edge:  $i \xrightarrow{\text{ENVY}} j$  if  $u_i(A_i) < u_i(A_j)$ ;
- Equality edge:  $i \xrightarrow{\text{EQ}} j$  if  $u_i(A_i) = u_i(A_j)$ .

Moreover, a cycle in an envy graph is called an *envy cycle* if it contains at least one envy edge. The concepts of envy edge and equality edge were also used in [27,30].

Given an envy graph, we then define another useful concept called *addable set* which corresponds to a specific set of agents.

**Definition 3.3** (*Addable set*). Given an envy graph, a non-empty set of agents  $S \subseteq N$  forms an *addable set* if,

- there is no envy edge between any pair of agents in  $S$ ;
- there exists neither an envy edge nor an equality edge from any agent in  $N \setminus S$  to any agent in  $S$ .

Moreover, an addable set  $S \subseteq N$  is called a *maximal addable set* if there does not exist any other addable set  $S' \subseteq N$  such that  $S \subset S'$ . The following lemma shows the uniqueness of the maximal addable set in an envy graph.

**Lemma 3.4.** *Given an envy graph, the maximal addable set, if exists, is unique. Moreover, we can find it or decide that none exists in  $O(n^3)$  time.*

**Proof.** Suppose, to the contrary, that there exist two distinct maximal addable sets  $S_1$  and  $S_2$  in the given envy graph. We will show that  $S_1 \cup S_2$  is also an addable set which contradicts the maximality of  $S_1$  and  $S_2$ .

First it is easy to see that there exists neither an envy edge nor an equality edge from any agent in  $N \setminus (S_1 \cup S_2)$  to agents in  $S_1 \cup S_2$  since, otherwise, either  $S_1$  or  $S_2$  is not an addable set.

We next argue that there is no envy edge between any pair of agents in  $S_1 \cup S_2$ . Clearly, according to Definition 3.3, there is no envy edge within each of  $S_1$  and  $S_2$ . The envy edges between  $S_1$  and  $S_2$  also cannot exist because there are no envy edges coming from outside of  $S_1$  or  $S_2$  into any of them. Thus,  $S_1 \cup S_2$  is also an addable set.

Now, we show how to find the unique maximal addable set or decide its non-existence in  $O(n^3)$  time: for each  $j$  which has an incoming envy edge, let  $R_j$  be the collection of vertices (including  $j$ ) that are reachable by  $j$  via the union of envy



**Fig. 1.** Fig. 1a shows an envy graph  $G$  with four vertices (agents)  $a, b, c, d$ . The bundle each agent gets is labeled with  $A_1, A_2, A_3, A_4$  beside. We show in Fig. 1b its corresponding  $G'$ . In this example,  $G$  has an envy cycle (involving vertices  $a, b, c$ ) but no addable set.



**Fig. 2.** After rotating the bundles along the envy cycle in Fig. 1a, we obtain the envy graph in Fig. 2a. The corresponding  $G'$  is shown in Fig. 2b. In this example,  $G$  has addable sets  $\{b\}$ ,  $\{b, c\}$  and  $\{a, b, c\}$  but no envy cycle.

edges and equality edges, and let  $S = N \setminus \bigcup_j R_j$ . We will show that an addable set does not exist in the envy graph if  $S = \emptyset$ . Otherwise,  $S$  is the unique maximal addable set. First,  $S$  is an addable set because any agent in  $S$  does not have any incoming envy edge and is not reachable via the union of envy edges and equality edges from any other agent with an incoming envy edge. In addition,  $S$  is maximal because any agent in  $\bigcup_j R_j$  cannot be in any addable set. Such  $S$  can be found in  $O(n^3)$  time because it takes  $O(n)$  time to check if an agent has an incoming edge, and for any agent  $j$  who has an incoming envy edge, it then takes  $O(n^2)$  time to construct  $R_j$  via, for example, breadth-first search (BFS).  $\square$

Intuitively, agents in the addable set  $S$  can be allocated some cake without creating new envy, since each agent in  $N \setminus S$  values her own bundle strictly more than the bundles of agents in  $S$ . Our next result characterizes the relation between the addable set and the envy cycle.

**Lemma 3.5.** Any envy graph  $G = (N, E_{envy} \cup E_{eq})$  that does not have any envy cycle must have an addable set.

**Proof.** We assume without loss of generality that  $E_{envy} \neq \emptyset$ , since otherwise  $N$  itself is an addable set. Now, we construct graph  $G' = (N', E')$  from  $G$  as follows. Each envy edge  $i \xrightarrow{ENVY} j$  in  $G$  corresponds to a vertex  $v_{ij}$  in  $G'$ . For two envy edges  $i \xrightarrow{ENVY} j$  and  $i' \xrightarrow{ENVY} j'$  in  $G$ , if there exists a path from  $j$  to  $i'$ , we construct an edge  $v_{ij} \rightarrow v_{i'j'}$  in  $G'$ . Note that, if there is an envy edge  $i \xrightarrow{ENVY} j$  and a path from  $j$  to  $i$  in  $G$ , there will be a self-loop  $v_{ij} \rightarrow v_{ij}$  in  $G'$ . We illustrate this transformation using two examples in Figs. 1 and 2.

It is easy to see that a cycle in  $G'$  implies an envy cycle in  $G$ . Thus, by the assumption that there is no envy cycle in  $G$ ,  $G'$  must be acyclic. Then there must exist a vertex  $v_{ij} \in N'$  which is not reachable by any other vertices in  $G'$ . Because  $v_{ij}$  corresponds to the envy edge  $i \xrightarrow{ENVY} j$  in  $G$ , since  $v_{ij}$  cannot be reached by any vertices in  $G'$ , the vertex  $i$  is also not reachable by any  $j'$  which is pointed by an envy edge. We note that, however, this vertex  $i$  may be reachable by other vertices via only equality edges. Thus, we need to not only include agent  $i$  in the addable set but also those agents who are able to reach  $i$  via equality edges.

Let  $S$  be the set containing agent  $i$  and all other agents who can reach  $i$  in the envy graph  $G$  via equality edges. In the following, we show that  $S$  is an addable set. First,  $S$  is non-empty because it at least contains agent  $i$ . Second, by our construction, there is no envy edge between any pair of agents in  $S$ . Third, recall that in envy graph  $G$ , agent  $i$  is not reachable by any  $j'$  which is pointed by an envy edge; thus,  $S$  is also not pointed by any envy edge. Last,  $S$  is also not pointed by any equality edge by our construction of  $S$ . Therefore, according to Definition 3.3,  $S$  must be an addable set.  $\square$

### 3.1. The algorithm

The complete algorithm to compute an EFM allocation is shown in Algorithm 1.

In general, our algorithm always maintains a partial allocation that is EFM. Then, we repeatedly and carefully add resources to the partial allocation, until all resources are allocated. We start with an EF1 allocation of only indivisible goods

**Algorithm 1** EFM algorithm.

---

**Require:** Agents  $N$ , indivisible goods  $M$  and cake  $C$ .

- 1: Find an arbitrary EF1 allocation  $(A_1, A_2, \dots, A_n)$  of  $M$  to  $n$  agents.
- 2: Construct an envy graph  $G = (N, E_{\text{envy}} \cup E_{\text{eq}})$  accordingly.
- 3: **while**  $C \neq \emptyset$  **do**
- 4:   **if** there exists an addable set in  $G$  **then**
- 5:     // cake-adding phase
- 6:     Let  $S$  be the maximal addable set.
- 7:     **if**  $S = N$  **then**
- 8:       Find an EF allocation  $(C_1, C_2, \dots, C_n)$  of  $C$ .
- 9:        $C \leftarrow \emptyset$
- 10:       Add  $C_i$  to bundle  $A_i$  for all  $i \in N$ .
- 11:     **else**
- 12:        $\delta_i \leftarrow \min_{j \in S} (u_i(A_j) - u_i(A_j))$  for each  $i \in N \setminus S$ .
- 13:       **if**  $u_i(C) \leq |S| \cdot \delta_i$  holds for each  $i \in N \setminus S$  **then**
- 14:          $C' \leftarrow C, C \leftarrow \emptyset$
- 15:         **else**
- 16:           Suppose w.l.o.g. that  $C = [a, b]$ . For each agent  $i \in N \setminus S$ , if  $u_i([a, b]) \geq |S| \cdot \delta_i$ , let  $x_i$  be a point such that  $u_i([a, x_i]) = |S| \cdot \delta_i$ ; otherwise, let  $x_i = b$ .
- 17:            $i^* \leftarrow \arg \min_{i \in N \setminus S} x_i$
- 18:            $C' \leftarrow [a, x_{i^*}], C \leftarrow C \setminus C'$
- 19:           **end if**
- 20:           Let  $(C_1, C_2, \dots, C_k) = \text{PERFECTALLOC}(C', k, N)$  where  $k = |S|$ .
- 21:           Add  $C_i$  to the bundle of the  $i$ -th agent in  $S$ .
- 22:           Update envy graph  $G$  accordingly.
- 23:         **end if**
- 24:     **else**
- 25:       // envy-cycle-elimination phase
- 26:       Let  $T$  be an envy cycle in envy graph  $G$ .
- 27:       For each agent  $j \in T$ , give agent  $j$ 's whole bundle to agent  $i$  who points to her in  $T$ .
- 28:       Update envy graph  $G$  accordingly.
- 29:     **end if**
- 30: **end while**
- 31: **return**  $(A_1, A_2, \dots, A_n)$

---

to all agents in Step 1, and construct the corresponding envy graph in Step 2. Then, our algorithm executes in rounds (Steps 3-30). In each round, we try to distribute some cake to the partial allocation while ensuring the partial allocation to be EFM. Such distribution needs to be done carefully because once an agent is allocated with a positive amount of cake, the fairness condition with regard to her bundle changes from EF1 to EF, which is more demanding. We repeat the process until the whole cake is allocated.

In each round of Algorithm 1, depending on whether there is an addable set that can be given some cake in Step 4, we execute either the *cake-adding phase* (Steps 4-23) or the *envy-cycle-elimination phase* (Steps 24-29).

- In the cake-adding phase, we have a maximal addable set  $S$ . By its definition, each agent in  $N \setminus S$  values her own bundle strictly more than the bundles of agents in  $S$ . Thus there is room to allocate some cake  $C'$  to agents in  $S$ . We carefully select  $C'$  to be allocated to  $S$  such that it does not create any new envy among the agents. To achieve this, we choose a piece of cake  $C' \subseteq C$  to be *perfectly* allocated to  $S$  in Steps 12-19 so that no agent in  $N$  will envy agents in  $S$  after distributing  $C'$  in Steps 20-21. More specifically, for each agent  $i \in N \setminus S$ , we determine in Step 12 the largest value  $\delta_i$  to be added to any agent in  $S$  such that  $i$  would still not envy any agent in  $S$ . Then, the way we decide  $x_{i^*}$  in Steps 17-18 ensures that for all agents  $i \in N \setminus S$ ,  $v_i([a, x_{i^*}]) \leq |S| \cdot \delta_i$ . Next, in Step 20, cake  $C' = [a, x_{i^*}]$  is divided into  $|S|$  pieces that are valued equally by all agents in  $N$ . This is to ensure that no agent  $i \in N \setminus S$  values any piece more than  $\delta_i$ .
- In the envy-cycle-elimination phase, i.e., when there does not exist any addable set, we show that in this case there must exist an envy cycle  $T$  in the current envy graph. We can then apply the envy-cycle-elimination technique to reduce some existing envy from the allocation by rearranging the bundles along  $T$ . More specifically, for each agent  $j \in T$ , we give agent  $j$ 's bundle to agent  $i$  who points to her in  $T$  (shown in Step 27).

We remark that when all goods are indivisible, our algorithm performs Steps 1-2 and terminates with an EF1 allocation (which is also EFM). When the whole good is a divisible cake, the algorithm goes directly to Step 8 and ends with an EF allocation of the cake, which is again EFM.

In the following we prove the correctness of this algorithm and analyze its running time.

### 3.2. Analysis

Our main result for the EFM allocation is as follows:

**Theorem 3.6.** *An EFM allocation always exists for any number of agents with additive valuations and can be found by Algorithm 1 in polynomial time with  $O(n^4)$  Robertson-Webb queries and  $O(n^3)$  calls to the PERFECTALLOC oracle.*

To prove Theorem 3.6, we first show that the following invariants are maintained by Algorithm 1 during its run.

---

*Invariants*

- A1.** In each round there is either an addable set for the cake-adding phase or an envy cycle for the envy-cycle-elimination phase.  
**A2.** The partial allocation is always EFM.
- 

**Lemma 3.7.** *Invariant A1 holds during the algorithm's run.*

**Proof.** This invariant is implied directly by Lemma 3.5.  $\square$

**Lemma 3.8.** *Invariant A2 holds during the algorithm's run.*

**Proof.** The partial allocation is clearly EFM after Step 1. Then the allocation is updated in three places in the algorithm: Steps 10 and 21 in the cake-adding phase and Step 27 in the envy-cycle-elimination phase. Given a partial allocation that is EFM, we will show that each of these updates maintains the EFM condition.

First, when we have  $S = N$  in Step 7, i.e., the addable set  $S$  consists of all  $n$  agents, the current envy graph does not contain any envy edge due to the definition of addable set (Definition 3.3). This implies that current partial allocation actually is envy-free. Because all valuation functions are additive, adding another envy-free allocation on top of it in Step 10 results in an envy-free and, hence, EFM allocation.

We next consider Step 21 in the cake-adding phase where a piece of cake is added to the addable set  $S$ . In order to maintain an EFM partial allocation, we need to ensure that this process does not introduce any new envy towards agents in  $S$ . Since we add a perfect allocation in Steps 20-21, envy will not emerge among agents in  $S$ . We also carefully choose the amount of cake to be allocated in Steps 13-19 such that each agent in  $N \setminus S$  weakly prefers her bundle to any bundles that belong to agents in  $S$ . To achieve this, we choose a piece of cake  $C' \subseteq C$  to be perfectly allocated to  $S$  in Steps 12-19 so that no agent in  $N$  will envy agents in  $S$  after distributing  $C'$  in Steps 20-21. More specifically, for each agent  $i \in N \setminus S$ , we determine in Step 12 the largest value  $\delta_i$  to be added to any agent in  $S$  such that  $i$  would still not envy any agent in  $S$ . Then, the way we decide  $x_{i^*}$  in Steps 17-18 ensures that for all agents  $i \in N \setminus S$ ,  $v_i([a, x_{i^*}]) \leq |S| \cdot \delta_i$ . Next, in Step 20, cake  $C' = [a, x_{i^*}]$  is divided into  $|S|$  pieces that are valued equally by all agents in  $N$ . This is to ensure that for any piece of cake  $C_j$  allocated to  $j \in S$ , we have  $u_i(C_j) \leq \delta_i$  for all  $i \in N \setminus S$ . Thus, agents in  $N \setminus S$  continue to not envy agents in  $S$  in Step 21.

Finally, in the envy-cycle-elimination phase, Step 27 eliminates envy edges by rearranging the partial allocation within the envy cycle  $T$ . Since each agent in  $T$  is weakly better off, the partial allocation remains EFM. For agents in  $N \setminus T$ , rearranging the partial allocation that is EFM will not make EFM infeasible. The conclusion follows.  $\square$

*Correctness*

**Lemma 3.9.** *Algorithm 1 always returns an EFM allocation upon termination.*

**Proof.** By Invariant A2, it suffices to prove that all goods are allocated when Algorithm 1 terminates. All indivisible goods are allocated in Step 1. Then the while loop (Steps 3-30) terminates only when the cake is also fully allocated, as desired.  $\square$

*Termination and time complexity* We use the number of envy edges in the envy graph and the size of the maximal addable set as a potential function to bound the running time of this algorithm.

**Lemma 3.10.** *After the algorithm completes a cake-adding phase, the number of envy edges never increases. In addition, if the piece of cake to be allocated is not the whole remaining cake, either (a) the number of envy edges strictly decreases, or (b) the size of the maximal addable set strictly decreases or an addable set no longer exists.*

**Proof.** By Lemma 3.8, the partial allocation is always EFM after a cake-adding phase. In a cake-adding phase, some positive amount of cake is added to every agent in  $S$ . This means after this phase, there would never be any envy edge between agents in  $S$  or from  $N \setminus S$  to  $S$ . The bundles of agents in  $N \setminus S$  remain the same, hence the set of edges among agents in  $N \setminus S$  remains unchanged. Lastly, since only agents in  $S$  are allocated new resources in the cake-adding phase, no new envy edge will be introduced from  $S$  to  $N \setminus S$ . This proves the first part of Lemma 3.10.

For the second part, we only study the situation when the piece of cake to be allocated to agents in  $S$  is not the whole remaining cake (Steps 15-18). Note that the number of envy edges will never increase after a cake-adding phase as proved

above. It suffices to show that if the number of envy edges remains unchanged and an addable set still exists, then the size of the maximal addable set must strictly decrease.

Note that based on how we choose  $i^*$  in Step 17, after the cake-adding phase, at least one equality edge will be generated in the envy graph from agent  $i^*$  to some agent  $j \in S$ . Let  $G$  and  $G'$  be the envy graphs before and after the cake-adding phase, and let  $S$  and  $S'$  be the maximal addable set of  $G$  and  $G'$ , respectively. In the following we will show that  $S' \subset S$ .

We first show  $S' \subseteq S$ . Suppose otherwise, we will show that  $S \cup S'$  is also an addable set in  $G$ , which contradicts to the maximality of  $S$ . The reasons that  $S \cup S'$  is an addable set in  $G$  are as follows.

- (i) We have already proved that compared to  $G$ , there is no new envy edge in  $G'$ . If  $G$  and  $G'$  has the same number of envy edges, they must share exactly the same set of envy edges. Hence, there will be no envy edge pointing to either  $S$  or  $S'$  in  $G$ .
- (ii) If there is an equality edge from  $N \setminus (S \cup S')$  to  $S \cup S'$  in  $G$ , this equality edge cannot be from  $N \setminus (S \cup S')$  to  $S$  because  $S$  is an addable set in  $G$ . Hence, it must be from  $N \setminus (S \cup S')$  to  $S' \setminus S$ . This equality edge remains in  $G'$  because neither the agents in  $N \setminus (S \cup S')$  nor the agents in  $S' \setminus S$  receive any good. However, this is impossible because  $S'$  is an addable set in  $G'$ . In summary, there cannot be any equality edges from  $N \setminus (S \cup S')$  to  $S \cup S'$  in  $G$ .

To further prove  $S' \subset S$ , we recall that according to our algorithm, at least one equality edge, from agent  $i^*$  in  $N \setminus S$  to some agent  $j \in S$ , will be included in  $G'$ . It is then clear that  $j$  cannot be in  $S'$ . This concludes the proof.  $\square$

**Lemma 3.11.** *After the algorithm completes an envy-cycle-elimination phase, the number of envy edges strictly decreases.*

**Proof.** The basic idea of this proof follows from Lipton et al. [30], albeit only strict envy edges were considered in their context. In the envy-cycle-elimination phase, an envy cycle  $T$  is eliminated by giving agent  $j$ 's bundle to agent  $i$  for each edge  $i \xrightarrow{\text{ENVY}} j$  or  $i \xrightarrow{\text{EQ}} j$  in the cycle. First, this process does not affect the bundles of agents in  $N \setminus T$ , hence the set of envy edges among them remains the same. Next, since we only swap bundles in this phase, the number of envy edges from  $N \setminus T$  to  $T$  remains the same. In addition, every agent  $i \in T$  receives a weakly better bundle, meaning that the number of envy edges from  $T$  to  $N \setminus T$  does not increase. Finally, because  $T$  contains at least one envy edge, some agent in  $T$  will receive a strictly better bundle. As a result, although some envy edges between agents in  $T$  may still exist, the total number of envy edges will decrease by at least one.  $\square$

**Lemma 3.12.** *Algorithm 1 terminates in polynomial time with  $O(n^3)$  calls to the PERFECTALLOC oracle and  $O(n^4)$  Robertson-Webb queries.*

**Proof. Calls to the PERFECTALLOC oracle.** By Invariant A1, each round in Algorithm 1 executes either a cake-adding phase or an envy-cycle-elimination phase. According to Lemmas 3.10 and 3.11, the number of envy edges never increases. Thus the number of rounds in which the number of envy edges strictly decreases is bounded by  $O(n^2)$ .

We now upper bound the number of cake-adding phase rounds between any two consecutive rounds that decrease the number of envy edges. If the whole remaining cake is allocated (Step 8), PERFECTALLOC( $C, n, N$ ) is called once and then Algorithm 1 terminates. In the case that a piece of remaining cake is allocated, by Lemma 3.10, the size of the maximal addable set strictly decreases or an addable set no longer exists; in the latter case, the algorithm proceeds to an envy-cycle-elimination phase. Because the size of any addable set is  $O(n)$ , it means that the number of cake-adding phase rounds between any two consecutive rounds that decrease the number of envy edges is  $O(n)$ .

Finally, it follows that Algorithm 1 executes at most  $O(n^2) \cdot O(n) = O(n^3)$  cake-adding phase rounds. Every such round calls the PERFECTALLOC oracle once. Algorithm 1 makes  $O(n^3)$  calls to the PERFECTALLOC oracle.

**Polynomial running time and RW queries.** Note that during the algorithm's run, we add resources to a bundle and rotate bundles among agents, but never split a bundle. For example, the partition of indivisible goods is computed in Step 1 and remains the same since then. To avoid redundant computations, we maintain an  $n$  by  $n$  array to keep track of  $u_i(A_j)$  for all  $i, j \in N$  and update them as necessary.

In Step 1, finding an EF1 allocation of indivisible goods can be done in  $O(mn \log m)$  via the round-robin algorithm [17]. The implementation details are as follows. We first compute the sorted order of goods according to each agent's valuation, which takes  $O(nm \log m)$  time overall. Next, in each agent's turn, we keep looking for the next unallocated good in that agent's sorted list. This step takes  $O(mn)$  time in total. Therefore, the overall running time of the round-robin algorithm is dominated by  $O(mn \log m)$ .

Next, in Step 2, the overall time to construct the corresponding envy graph is  $O(n^2)$ .

We now consider the while loop. According to Lemma 3.4, we can find the maximal addable set or decide its non-existence in time  $O(n^3)$ . In the case that we need to perform an envy-cycle-elimination, an envy cycle  $T$  can be found in the following way. Fix an agent  $i$ , we can first spend  $O(n)$  time scanning all outgoing edges and ignore those equality edges. Then, we apply depth-first search (DFS) starting from vertex  $i$ . If there is a back edge pointing to vertex  $i$ , then there must be an envy cycle with at least one envy edge, say, e.g.,  $i \xrightarrow{\text{ENVY}} j$ , for some  $j \in N$ . This takes  $O(n^2)$  time since DFS dominates the time complexity. Since there are  $O(n)$  agents, overall, this step can be implemented in  $O(n^3)$  time.



In the following, we discuss the steps in each phase at length.

**Cake-adding phase** When we have  $S = N$  satisfied in Step 7, we implement an EF allocation by calling `PERFECTALLOC(C, n, N)`. It takes  $O(n)$  time to update the allocation. Algorithm 1 then terminates.

It takes  $O(n^2)$  time in Step 12 to compute  $\delta_i$  for all  $i \in N \setminus S$ . Steps 13 and 16 need  $O(n)$  evaluation and cut queries respectively. Once  $C'$  is determined in Step 18, we can make  $O(n)$  evaluation queries from all  $n$  agents over  $C'$ . Because we use a perfect allocation of  $C'$ , we can directly compute  $u_i(C')/|S|$  for all  $i \in N$  to obtain the value increment of each agent in the addable set. It then takes  $O(n^2)$  time to update all agents' valuations of all bundles after Step 21. After this, updating an envy graph also takes  $O(n^2)$  time.

Since we only make RW queries in this phase, we summarize here that Algorithm 1 makes  $O(n^4)$  Robertson-Webb queries, because there are  $O(n^3)$  cake-adding phases (stated earlier in this proof) and each such phase needs at most  $O(n)$  RW queries.

**Envy-cycle-elimination phase** Since we maintain an array as the reference for agents' valuations over the current bundles, we can rotate the bundles as well as update the array and the envy graph in time  $O(n^2)$ .

The remaining steps can be implemented in time  $O(n)$ . Overall, Algorithm 1 runs in time  $O(mn \log m + n^6)$ , where the  $n^6$  term comes from the  $O(n^3)$  total number of `while` loops and  $O(n^3)$  time to run each loop.  $\square$

Finally the correctness of Theorem 3.6 is directly implied by Lemma 3.9 and Lemma 3.12.

*Bounded protocol in the RW model* Even though we showed that Algorithm 1 can produce an EFM allocation, it is not a bounded protocol in the RW model. This is because our algorithm utilizes an oracle that can compute a perfect allocation of any piece of cake. However, while a perfect allocation always exists, it is known that such an allocation cannot be implemented with a finite number of queries in the RW model, even if there are only two agents [39]. Whether there exists a bounded protocol in the RW model to compute an EFM allocation remains a very interesting open question. Note that the perfect allocation oracle cannot be implemented even with a finite number of queries, therefore it is even an open question to find a finite EFM protocol.

A natural and tempting approach to get a bounded EFM protocol would be to use an envy-free allocation, for which a bounded protocol in the RW model is known [5], to replace the perfect allocation in Step 20. Note that doing so would not create any new envy edge within set  $S$ . Then, in order to not create any envy edge from  $N \setminus S$  to  $S$ , we need to restrict the total value of the piece of cake allocated to  $S$  to not exceed  $\delta_i$  (Step 12) for every agent  $i \in N \setminus S$ . However, when doing so, we will not be able to quantify the progress of the algorithm like in Lemma 3.10. Specifically, we can no longer guarantee that either the number of envy edges strictly decreases or the size of the maximal addable set strictly decreases. This is because we are not guaranteed the equality edge from agent  $i^*$  to some agent  $j \in S$  as we rely on in the proof of Lemma 3.10. In other words, we cannot show the algorithm will always terminate in bounded steps. Interestingly, in sharp contrast, in Section 5 we will show that an approximate envy-free protocol, instead of an  $\epsilon$ -perfect protocol, is enough to give an efficient  $\epsilon$ -EFM algorithm. We will discuss this phenomenon in further detail in Section 5.

In the next two sections, we present two bounded protocols to compute an EFM allocation for two special cases, and another bounded (polynomial time) protocol to compute an  $\epsilon$ -EFM allocation in the general case.

#### 4. EFM allocation in special case

In this section, we show two special cases where an EFM allocation can be computed in polynomial time without using the perfect allocation oracle. One is the 2-agent case with general valuations while the other deals with the  $n$ -agent case but each agent has a structured density function for the cake.

##### 4.1. Two agents

We first show that with only two agents, an EFM allocation can be found using a simple cut-and-choose type of algorithm. We start with a partition  $(M_1, M_2)$  of all indivisible goods such that agent 1 is EF1 with respect to either bundle. Without loss of generality, we assume that  $u_1(M_1) \geq u_1(M_2)$ . Next agent 1 adds the cake into  $M_1$  and  $M_2$  so that the two bundles are as close to each other as possible. Note that if  $u_1(M_1) > u_1(M_2 \cup C)$ , agent 1 would add all cake to  $M_2$ . If  $u_1(M_1) \leq u_1(M_2 \cup C)$ , agent 1 has a way to make the two bundles equal. We then give agent 2 her preferred bundle and leave to agent 1 the remaining bundle.

**Theorem 4.1.** *Algorithm 2 returns an EFM allocation in the case of two agents in polynomial time.*

**Proof. Correctness.** It is obvious that all goods are allocated. We next show that the allocation returned is EFM. Agent 2 is guaranteed EF (thus EFM) since she gets her preferred bundle between  $A_1$  and  $A_2$ . In the following, we focus on agent 1. If  $u_1(M_1) \leq u_1(M_2 \cup C)$  holds, agent 1 is indifferent between bundles  $A_1$  and  $A_2$ , so either  $A_1$  or  $A_2$  makes her EF (thus

**Algorithm 2** EFM allocation for two agents.**Require:** Agents 1 and 2, indivisible goods  $M$  and cake  $C$ .

- 1: Divide  $M$  into two parts,  $M_1, M_2$ , such that agent 1 is EF1 with respect to either bundle. Assume w.l.o.g. that  $u_1(M_1) \geq u_1(M_2)$  (otherwise we can swap  $M_1$  and  $M_2$ ).
- 2: **if**  $u_1(M_1) \leq u_1(M_2 \cup C)$  **then**
- 3: Let agent 1 partition the cake into two pieces  $(C_1, C_2)$ , such that  $u_1(M_1 \cup C_1) = u_1(M_2 \cup C_2)$ .
- 4: Let  $(A_1, A_2) = (M_1 \cup C_1, M_2 \cup C_2)$ .
- 5: **else**
- 6: Let  $(A_1, A_2) = (M_1, M_2 \cup C)$ .
- 7: **end if**
- 8: Give agent 2 her preferred bundle among  $A_1, A_2$ . Give agent 1 the remaining bundle.

EFM). In the case that  $u_1(M_1) > u_1(M_2 \cup C)$  holds, agent 1 is EF if she receives  $A_1$ , and is EFM if she gets  $A_2$  because  $A_1$  consists of only indivisible goods and there exists some good  $g$  in  $A_1$  such that  $u_1(A_2) \geq u_1(M_2) \geq u_1(A_1 \setminus \{g\})$ .

**Polynomial time.** To obtain the initial partition  $(M_1, M_2)$ , we can let two copies of agent 1 run the round-robin algorithm on the indivisible items. This step can be done easily in polynomial time. The remaining steps only take constant running time and a constant number of RW queries. The conclusion follows.  $\square$

*A stronger EFM notion* With two agents, an *envy-freeness up to any good (EFX)* allocation, in which no agent prefers the bundle of another agent following the removal of *any* single good, always exists [35]. This result can be carried over to show the existence of a stronger EFM notion in the mixed goods setting, in which an agent is EFX towards any agent with only indivisible goods, and EF towards the rest. Such an allocation can be obtained by using an EFX partition (with respect to agent 1) instead of an EF1 partition in Step 1 of Algorithm 2. Moreover, with any number of agents, whenever an EFX allocation exists among indivisible goods,<sup>4</sup> we can start with such an EFX allocation in Step 1 of Algorithm 1. The cake-adding phase maintains the EFM condition and does not introduce new envy. Thus Algorithm 1 will also produce an allocation with this stronger notion of EFM.

#### 4.2. Any number of agents with piecewise linear functions

In the second case, we consider an arbitrary number of agents when agents' valuation functions over the cake are *piecewise linear*.

**Definition 4.2.** A valuation density function  $f_i$  is *piecewise linear* if the interval  $[0, 1]$  can be partitioned into a finite number of intervals such that  $f_i$  is linear on each interval.

Piecewise linear function is a generalization of both *piecewise uniform function* and *piecewise constant function*, each of which has been considered in several previous fair division works [8,20,9]. In this case, we do not use the RW model, but rather assume that the valuation functions are provided to us in *full information*.

The only obstacle in converting Algorithm 1 to a bounded protocol is the implementation of the perfect allocation oracle for cake cutting. When agents have piecewise linear functions, Chen et al. [20] showed that a perfect allocation can be computed efficiently in polynomial time. This fact, combined with Theorem 3.6, directly implies the following result.

**Corollary 4.3.** For any number of agents with piecewise linear density functions over the cake, an EFM allocation can be computed in polynomial time.

### 5. $\epsilon$ -EFM: algorithm

In this section, we focus on  $\epsilon$ -EFM, a relaxation of the EFM condition. Despite the computational issues with finding bounded exact EFM protocols, we will show that there is an efficient algorithm in the RW model that computes an  $\epsilon$ -EFM allocation for general density functions with running time polynomial in  $n, m$  and  $1/\epsilon$ .

Since the difficulty in finding a bounded EFM protocol in the RW model lies in computing perfect allocations of a cake (Section 3), one might be tempted to simply use a bounded  $\epsilon$ -Perfect Allocation protocol to replace the exact procedure. Here a partition  $\mathcal{C} = (C_1, C_2, \dots, C_k)$  of cake  $C$  is said to be  $\epsilon$ -perfect if for all  $i \in N, j \in [k], |u_i(C_j) - u_i(C)/k| \leq \epsilon$ . However, although a bounded  $\epsilon$ -perfect protocol exists in the RW model [39,14], all known protocols have running time exponential

<sup>4</sup> EFX exists for three agents [19] or  $n$  agents with identical valuations [35], but the existence of EFX remains open for four or more agents with additive valuations.

in  $1/\epsilon$ .<sup>5</sup> It is still an open question to find an  $\epsilon$ -perfect allocation with both query and time complexity polynomial in  $1/\epsilon$ . Therefore, to design an efficient  $\epsilon$ -EFM protocol, extra work needs to be done to circumvent this issue.

We next define the relaxed version of EF and envy graph.

**Definition 5.1** ( $\epsilon$ -EF). An allocation  $\mathcal{A}$  is said to satisfy  $\epsilon$ -envy-freeness ( $\epsilon$ -EF) if for all agents  $i, j \in N$ ,  $u_i(A_i) \geq u_i(A_j) - \epsilon$ .

**Definition 5.2** ( $\epsilon$ -envy graph). Given an allocation  $\mathcal{A}$  and a parameter  $\epsilon$ , the  $\epsilon$ -envy graph is defined as  $G(\epsilon) = (N, E_{\epsilon\text{-envy}} \cup E_{\epsilon\text{-eq}})$ , where every vertex represents an agent, and  $E_{\epsilon\text{-envy}}$  and  $E_{\epsilon\text{-eq}}$  consist of the following two types of edges, respectively:

- $\epsilon$ -envy edge:  $i \xrightarrow{\epsilon\text{-ENVY}} j$  if  $u_i(A_i) < u_i(A_j) - \epsilon$ ;
- $\epsilon$ -equality edge:  $i \xrightarrow{\epsilon\text{-EQ}} j$  if  $u_i(A_j) - \epsilon \leq u_i(A_i) \leq u_i(A_j)$ .

Given an  $\epsilon$ -envy graph, a cycle is said to be an  $\epsilon$ -envy cycle if it contains at least one  $\epsilon$ -envy edge. We also note that when  $\epsilon = 0$ , the  $\epsilon$ -envy graph degenerates into the envy graph defined in Section 3.

### 5.1. The algorithm

The complete algorithm to compute an  $\epsilon$ -EFM allocation is shown in Algorithm 3. Similarly to Algorithm 1, our algorithm adds resources to the partial allocation iteratively. We always maintain the partial allocation to be  $\hat{\epsilon}$ -EFM where  $\hat{\epsilon}$  is updated increasingly and would never exceed  $\epsilon$ . This will ensure that the final allocation is  $\epsilon$ -EFM.

---

#### Algorithm 3 $\epsilon$ -EFM algorithm.

---

**Require:** Agents  $N$ , indivisible goods  $M$ , cake  $C$ , and parameter  $\epsilon$ .

```

1:  $\hat{\epsilon} \leftarrow \frac{\epsilon}{4}$ ,  $\gamma \leftarrow \frac{\epsilon^2}{8n}$ 
2: Find an arbitrary EF1 allocation  $(A_1, A_2, \dots, A_n)$  of  $M$  to  $n$  agents.
3: Construct an  $\hat{\epsilon}$ -envy graph  $G(\hat{\epsilon}) = (N, E_{\hat{\epsilon}\text{-envy}} \cup E_{\hat{\epsilon}\text{-eq}})$  accordingly.
4: while  $C \neq \emptyset$  do
5:   if there exists an addable set  $S$  then
6:     // cake-adding phase
7:     if  $S = N$  then
8:       Let  $(C_1, C_2, \dots, C_n) = \frac{\epsilon}{4}$ -EFALLOC( $C, N$ ).
9:        $C \leftarrow \emptyset$ 
10:       $\hat{\epsilon} \leftarrow \hat{\epsilon} + \epsilon/4$ 
11:      Add  $C_i$  to bundle  $A_i$  for all  $i \in N$ .
12:     else
13:       if  $\max_{i \in N \setminus S} u_i(C) \leq \hat{\epsilon}$  then
14:          $C' \leftarrow C, C \leftarrow \emptyset$ 
15:       else
16:         Suppose w.l.o.g. that  $C = [a, b]$ . For each agent  $i \in N \setminus S$ , if  $u_i([a, b]) \geq \hat{\epsilon}$ , let  $x_i$  be a point such that  $u_i([a, x_i]) = \hat{\epsilon}$ ; otherwise, let  $x_i = b$ .
17:          $i^* \leftarrow \arg \min_{i \in N \setminus S} x_i$ 
18:          $C' \leftarrow [a, x_{i^*}], C \leftarrow C \setminus C'$ 
19:       end if
20:       Let  $(C_1, C_2, \dots, C_k) = \gamma$ -EFALLOC( $C', S$ ) where  $k = |S|$ .
21:        $\hat{\epsilon} \leftarrow \hat{\epsilon} + \gamma$ 
22:       Add  $C_i$  to the bundle of the  $i$ -th agent in  $S$ .
23:       Update  $\hat{\epsilon}$ -envy graph  $G(\hat{\epsilon})$  accordingly.
24:     end if
25:   else
26:     // envy-cycle-elimination phase
27:     Let  $T$  be an  $\hat{\epsilon}$ -envy cycle in the  $\hat{\epsilon}$ -envy graph.
28:     For each agent  $j \in T$ , give agent  $j$ 's whole bundle to agent  $i$  who points to her in  $T$ .
29:     Update  $\hat{\epsilon}$ -envy graph  $G(\hat{\epsilon})$  accordingly.
30:   end if
31: end while
32: return  $(A_1, A_2, \dots, A_n)$ 

```

---

Like Algorithm 1, starting with an EF1 allocation of indivisible goods to all agents in Step 2, Algorithm 3 then executes in rounds (Steps 4–31). Even though each round still executes either a cake-adding phase or an envy-cycle-elimination phase, the execution details are different from Algorithm 1.

<sup>5</sup> Brânzei and Nisan [14] showed that an  $\epsilon$ -perfect allocation can be computed in  $O(n^3/\epsilon)$  RW queries. However, although the query complexity is polynomial, the protocol still requires an exponential running time because it finds the correct partition of the small pieces into bundles via an exhaustive enumeration, of which no polynomial time algorithm is known.

- In the cake-adding phase, instead of allocating some cake to an addable set  $S$  in a way that is perfect, we resort to a  $\gamma$ -EF allocation, where  $\gamma$  will be fixed later in Algorithm 3. In the following, we will utilize an algorithm  $\gamma$ -EFALLOC( $C, S$ ) that could return us a  $\gamma$ -EF allocation for any set of agents  $S$  and cake  $C$ . Note that, for any  $\bar{\epsilon} > 0$ , the algorithm  $\bar{\epsilon}$ -EFALLOC can be implemented with both running time and query complexity polynomial in the number of agents involved and  $1/\bar{\epsilon}$  [36]. We also update  $\hat{\epsilon}$  to a larger number, say  $\hat{\epsilon} + \gamma$ , in order to avoid generating  $\hat{\epsilon}$ -envy edges due to cake-adding.
- In the envy-cycle-elimination phase, we eliminate an  $\hat{\epsilon}$ -envy cycle, instead of an envy cycle, by rearranging the current partial allocation.

### 5.2. Analysis

Our main result for the  $\epsilon$ -EFM allocation is as follows:

**Theorem 5.3.** *An  $\epsilon$ -EFM allocation can be found by Algorithm 3 with running time  $O(n^4/\epsilon + mn \log m)$ ,  $O(n^3/\epsilon)$  Robertson-Webb queries, and  $O(n/\epsilon)$  calls to the approximate EFALLOC oracle.*

To prove Theorem 5.3, we first show that each round (Steps 4-31 in Algorithm 3) maintains the following invariants during the run of the algorithm.

---

#### Invariants

- B1.** In each round there is either an addable set for the cake-adding phase or an  $\hat{\epsilon}$ -envy cycle for the envy-cycle-elimination phase.
  - B2.** The partial allocation is always  $\hat{\epsilon}$ -EFM with the current  $\hat{\epsilon}$ .
- 

We next prove these invariants in the following.

**Lemma 5.4.** *Invariant B1 holds during the algorithm's run.*

**Proof.** The proof is similar to the proof of Lemma 3.7, except that we consider the  $\hat{\epsilon}$ -envy edge instead of the envy edge.  $\square$

**Lemma 5.5.** *Invariant B2 holds during the algorithm's run.*

**Proof.** First, it is worth noting that at the beginning, when indivisible goods are allocated, the allocation is EF1 and therefore EFM. We then note that the partial allocation is only updated in Steps 11 and 22 in the cake-adding phase as well as in Step 28 in the envy-cycle-elimination phase. Given a partial allocation that is  $\hat{\epsilon}$ -EFM, we will show that each of these updates maintains  $\hat{\epsilon}$ -EFM with the updated  $\hat{\epsilon}$ , which completes the proof of Lemma 5.5. We note that  $\hat{\epsilon}$  is only updated in the cake-adding phase and is non-decreasing during the algorithm's run.

For analysis in the envy-cycle-elimination phase, the proof is identical to that of Lemma 3.8 in the case of envy-cycle-elimination phase.

We then discuss the cases in the cake-adding phase. For the updated partial allocation in Step 11, we allocate the remaining cake  $C$  in a way that is  $\frac{\epsilon}{4}$ -EF which implies that  $u_i(C_i) \geq u_i(C_j) - \epsilon/4$  holds for any pair of agents  $i, j \in N$ . Given the partial allocation that is  $\hat{\epsilon}$ -EFM, we have

$$u_i(A_i \cup C_i) \geq u_i(A_j \cup C_j) - \hat{\epsilon} - \epsilon/4.$$

Thus it is clear that no  $\hat{\epsilon}$ -envy edge will be generated among agents in  $S$  if we update  $\hat{\epsilon}$  to  $\hat{\epsilon} + \epsilon/4$  in Step 10.

For the updated partial allocation in Step 22, we allocate some cake  $C'$  to agents in  $S$  in a way that is  $\gamma$ -EF. By a similar argument to the case above, we know that no  $\hat{\epsilon}$ -envy edge will be generated among agents in  $S$  if we update  $\hat{\epsilon}$  to  $\hat{\epsilon} + \gamma$ . Then, for any agent  $i \in N \setminus S$ , we have  $u_i(A_i) > u_i(A_j)$  where  $j \in S$ . As  $u_i(C') \leq \hat{\epsilon}$ , we have

$$u_i(A_i) > u_i(A_j \cup C'_j) - \hat{\epsilon},$$

where  $C'_j$  is the piece of cake allocated to agent  $j \in S$ . It means that again no  $\hat{\epsilon}$ -envy edge will be generated from  $N \setminus S$  to  $S$  when we update  $\hat{\epsilon}$  to  $\hat{\epsilon} + \gamma$ . Last, it is obvious that for any pair of agents in  $N \setminus S$ , we do not generate any  $(\hat{\epsilon} + \gamma)$ -envy edge because the bundles of these agents remain the same. We note that  $\hat{\epsilon}$  is updated to  $\hat{\epsilon} + \gamma$  in Step 21 in order to make sure that there is no introduced  $\hat{\epsilon}$ -envy edge in the updated  $\hat{\epsilon}$ -envy graph.  $\square$

#### Correctness

**Lemma 5.6.** *In the cake-adding phase, if the piece of cake to be allocated is not the whole remaining cake, the sum of all agents' valuations on the remaining cake decreases by at least  $\hat{\epsilon}$ .*

**Proof.** We consider the cake-adding phase in Steps 5-24. If the piece of cake  $C'$  to be allocated is not the whole remaining cake, there exists an agent  $i^*$  in  $N \setminus S$  such that  $u_{i^*}(C') = \hat{\epsilon}$  according to Steps 16-18. Thus the lemma follows.  $\square$

We are now ready to show the correctness of Algorithm 3.

**Lemma 5.7.** Algorithm 3 always returns an  $\epsilon$ -EFM allocation upon termination.

**Proof.** By Invariant B2, it suffices to prove that all goods are allocated and  $\hat{\epsilon}$  is at most  $\epsilon$  when Algorithm 3 terminates. All indivisible goods are allocated in Step 2. Then the while loop (Steps 4-31) terminates only when the cake is also fully allocated, as desired.

We now turn our attention to  $\hat{\epsilon}$ . First,  $\hat{\epsilon}$  is initialized to be  $\epsilon/4$  and never decreases during the algorithm's run. If the whole remaining cake is allocated, there is at most one execution of the cake-adding phase (Steps 7-11). Moreover,  $\hat{\epsilon}$  is increased by  $\epsilon/4$  in Step 10. We will show later in this proof that this increment would not let  $\hat{\epsilon}$  exceed  $\epsilon$ . We then focus on the case where the remaining cake is not fully allocated. There are at most  $n/\hat{\epsilon} \leq 4n/\epsilon$  executions of the cake-adding phase (Steps 12-23) according to Lemma 5.6 and the fact that agents' utilities are normalized to 1. In addition,  $\hat{\epsilon}$  is increased by  $\gamma$  in each cake-adding phase in Step 21. Thus,  $\hat{\epsilon}$  is upper bounded by  $\epsilon/4 + 4n/\epsilon \cdot \gamma + \epsilon/4 = \epsilon$  due to  $\gamma = \frac{\epsilon^2}{8n}$ . It follows that the final allocation is  $\epsilon$ -EFM.  $\square$

### Termination and time complexity

**Lemma 5.8.** In the envy-cycle-elimination phase, the social welfare  $\sum_{i \in N} u_i(A_i)$ , increases by at least  $\hat{\epsilon}$ .

**Proof.** We eliminate an  $\hat{\epsilon}$ -envy cycle  $T$  which contains at least one  $\hat{\epsilon}$ -envy edge in the envy-cycle-elimination phase (Steps 25-30). Step 28 eliminates the cycle by giving agent  $j$ 's bundle to agent  $i$  for each edge  $i \xrightarrow{\hat{\epsilon}\text{-ENVY}} j$  or  $i \xrightarrow{\hat{\epsilon}\text{-EQ}} j$  in cycle  $T$ . None of the agents involved in  $T$  is worse off and at least one agent in  $T$  is better off by at least  $\hat{\epsilon}$  by the definition of  $\hat{\epsilon}$ -envy edge in Definition 5.2. Since agents outside cycle  $T$  do not change their bundles, we complete the proof.  $\square$

**Lemma 5.9.** Algorithm 3 has running time  $O(n^4/\epsilon + mn \log m)$  and invokes  $O(n/\epsilon)$  calls to the approximate EFALLOC oracle and  $O(n^3/\epsilon)$  Robertson-Webb queries.

**Proof.** Since most parts of Algorithm 3 are similar to those in Algorithm 1, and we have discussed their time complexities in the proof of Lemma 3.12, we will focus on the steps that affect the time complexity for Algorithm 3 in this proof.

Similar to Algorithm 1, Algorithm 3 takes  $O(mn \log m + n^2)$  time to perform Steps 2 and 3. Afterwards, by Invariant B1, Algorithm 3 executes either a cake-adding phase or an envy-cycle-elimination phase in each round, and it takes  $O(n^3)$  time to check which phase to go into each time. Next, recall that agents' utilities are normalized to 1 and  $\hat{\epsilon}$  is always no less than  $\epsilon/4$ . This means there are at most  $O(n/\epsilon)$  cake-adding rounds by Lemma 5.6 and at most  $O(n/\epsilon)$  envy-cycle-elimination rounds by Lemma 5.8.

In the following, we discuss the steps in each phase in details.

**Cake-adding phases** When we have  $S = N$  in Step 7, we invoke  $\frac{\epsilon}{4}$ -EFALLOC once, use  $O(n)$  time to update the allocation, and terminate the algorithm.

To determine the piece of cake  $C'$  to be allocated later, we need  $O(n)$  evaluation queries in Step 13, and  $O(n)$  cut queries in Step 16 if the condition check in Step 13 fails. In each cake-adding phase, we invoke the  $\gamma$ -EFALLOC oracle with  $\gamma = \frac{\epsilon^2}{8n}$  once (Step 20). In order to update agents' valuations for each bundle, we invoke  $O(n^2)$  evaluation queries to obtain all agents' valuations of all pieces in  $(C_1, C_2, \dots, C_k)$ , and then use  $O(n^2)$  time to update the envy graph.

Summarizing everything, we conclude that Algorithm 3 makes  $O(n^3/\epsilon)$  Robertson-Webb queries, and  $O(n/\epsilon)$  calls to the  $\frac{\epsilon^2}{8n}$ -EFALLOC oracle in total in all cake-adding phases.

**Envy-cycle-elimination phases** Since we are keeping track of all agents' valuations for all bundles, it takes no Robertson-Webb queries and  $O(n^2)$  time to rearrange the bundles as well as update the  $\hat{\epsilon}$ -envy graph. Overall all envy-cycle-elimination phases take  $O(n^3/\epsilon)$  running time.

The remaining steps can be implemented in time  $O(n)$ . The overall time complexity of our algorithm is  $O(n/\epsilon \cdot n^3 + n/\epsilon \cdot n^2 + mn \log m + n^2) = O(n^4/\epsilon + mn \log m)$ . The overall Robertson-Webb query complexity is  $O(n^3/\epsilon)$ , and the number of calls to the  $\gamma$ -EFALLOC oracle is  $O(n/\epsilon)$ .  $\square$

Finally the correctness of Theorem 5.3 is directly implied by Lemmas 5.7 and 5.9.

Note that at the end of Section 3, we explained why an exact envy-free oracle may not be helpful to obtain an EFM allocation. However, as we showed in this section, the approximate envy-free oracle does help to obtain an  $\epsilon$ -EFM allocation.

Lemma 5.6 provides the key difference. In particular, the error allowed in the  $\epsilon$ -EFM condition ensures that agents' welfare for the remaining cake is reduced by at least an amount of  $\hat{\epsilon}$ . This claim, however, makes no sense when discussing exact envy-freeness. Furthermore, Algorithm 3 introduces additional error into the EFM condition on top of the error that comes from the approximate envy-free oracle. As a result, even if Algorithm 3 was paired with an exact envy-free oracle, it would still not produce an exact EFM allocation.

### 6. EFM and efficiency

In this section, we discuss how to combine EFM with efficiency considerations. In particular, we focus on the well-studied efficiency notion of *Pareto optimality*.

**Definition 6.1 (PO).** An allocation  $\mathcal{A}$  is said to satisfy *Pareto optimality (PO)* if there is no allocation  $\mathcal{A}'$  that Pareto-dominates  $\mathcal{A}$ , i.e., satisfies  $u_i(\mathcal{A}'_i) \geq u_i(\mathcal{A}_i)$  for all  $i \in N$  and at least one inequality is strict.

**Definition 6.2 (fPO [7]).** An allocation  $\mathcal{A}$  is said to satisfy *fractional Pareto optimality (fPO)* if it is not Pareto dominated by any fractional allocation.<sup>6</sup>

As Barman et al. [7] noted, an fPO allocation is also PO but not vice versa. It is known that with divisible resources, an allocation that is both envy-free and PO always exists [47]. With indivisible goods, an allocation satisfying both EF1 and fPO (and hence PO) also exists [7]. Perhaps to our surprise, in the following we show via a counter-example that with mixed types of goods, EFM and PO are no longer compatible.

**Example 6.3 (EFM is not compatible with PO).** Consider an instance with two agents, one indivisible good, and one cake. Agents' valuation functions are listed below.

	Indivisible good	A cake $C = [0, 1]$
Agent 1	0.6	$u_1(C) = 0.4$ with uniform density over $[0, 0.5]$
Agent 2	0.6	$u_2(C) = 0.4$ with uniform density over $[0.5, 1]$

It is obvious that in any EFM allocation, one agent will get the indivisible good and the entire cake has to be allocated to the other agent. However, such an allocation cannot be PO since the agent with the cake has no value for half of it, and giving that half to the other agent would make that agent better off without making the first agent worse off.

This counter-example relies on the fact that in the definition of EFM, if some agent  $i$ 's bundle contains any positive amount of cake, another agent  $j$  will compare her bundle to  $i$ 's bundle using the stricter EF condition, even if agent  $j$  has value zero over  $i$ 's cake. This may seem counter-intuitive, because when  $j$  has no value over  $i$ 's cake, removing that cake from  $i$ 's bundle will not help eliminate agent  $j$ 's envy. To this end, one may consider the following weaker version of EFM.

**Definition 6.4 (Weak EFM).** An allocation  $\mathcal{A}$  is said to satisfy *weak envy-freeness for mixed goods (weak EFM)*, if for any agents  $i, j \in N$ ,

- if agent  $j$ 's bundle consists of indivisible goods, and
  - either no divisible good,
  - or divisible goods that yield value 0 to agent  $i$ , i.e.,  $u_i(C_j) = 0$ ,

there exists an indivisible good  $g \in A_j$  such that  $u_i(A_i) \geq u_i(A_j \setminus \{g\})$ ;

- otherwise,  $u_i(A_i) \geq u_i(A_j)$ .

From the definition, it is easy to see that EFM implies weak EFM, which means all existence results for EFM established in Section 3 can be carried over to weak EFM.

This weaker version of EFM precludes the incompatibility result in Example 6.3. Nevertheless, we show in the following example that weak EFM is incompatible with fPO.

**Example 6.5 ((Weak) EFM is incompatible with fPO).** Consider an instance with two agents, one indivisible good and two homogeneous divisible goods. Agents' valuation functions are listed below.

<sup>6</sup> In a fractional allocation, an agent may get a fractional share of an *indivisible* good. We refer to [7] for its formal definition.

	Indivisible good	Divisible good 1	Divisible good 2
Agent 1	2	1	2
Agent 2	2	2	1

Because the valuations are symmetric, we can assume without loss of generality that in an EFM allocation, the indivisible good is given to agent 1. We also observe that in any EFM allocation, we cannot allocate all divisible goods to a single agent. This means that both agents' bundles must contain some divisible good, which then implies that both agents need to be envy-free towards the other agent's bundle. Next, via two simple linear programs one can compute the maximum utility of each agent in EFM allocations: giving the indivisible good and one half of divisible good 2 to agent 1 gives her a maximum utility 3; giving divisible good 1 and three quarters of divisible good 2 to agent 2 gives her a maximum utility 2.75. We note that the maximum utilities for the two agents are achieved under different allocations. However, even putting these two maximum utilities together, it is dominated by the utilities guaranteed by the fractional allocation in which agent 1 gets divisible good 2 and half of the indivisible good while agent 2 gets divisible good 1 and the other half of the indivisible good, which will give both agents a utility of 3. This means that any EFM allocation is not fPO in this problem instance.

Could there always exist an allocation that satisfies both weak EFM and PO? We do not know the answer, and believe this is a very interesting open question. One tempting approach to answer this open question is to consider the maximum Nash welfare (MNW) allocation. This is the allocation that maximizes the Nash welfare  $\prod_{i \in N} u_i(A_i)$  among all allocations.<sup>7</sup> It has been shown that an MNW allocation enjoys many desirable properties in various settings. In particular, an MNW allocation is always envy-free and PO in cake cutting [42], and EF1 and PO for indivisible resource allocation [17]. It is therefore a natural idea to conjecture that it also satisfies EFM and PO for mixed goods. Unfortunately, this is not the case. Here we give such a counter-example.

**Example 6.6** (MNW does not imply (weak) EFM). Consider the following instance with two agents, two indivisible goods and one homogeneous cake. Agents' valuation functions are listed below.

	Indivisible good 1	Indivisible good 2	A homogeneous cake
Agent 1	0.4	0.4	0.2
Agent 2	0.499	0.499	0.002

We discuss the following cases to find the MNW allocation.

- When both indivisible goods are given to agent 1, giving the whole cake to agent 2 maximizes the Nash welfare, which is  $(0.4 + 0.4) \times 0.002 = 0.0016$ .
- When both indivisible goods are given to agent 2, giving the whole cake to agent 1 maximizes the Nash welfare, which is  $0.2 \times (0.499 + 0.499) = 0.1996$ .
- When each agent gets exactly one indivisible good, in the Nash welfare maximizing allocation, denoted by  $\mathcal{A}$ , agent 1 receives an indivisible good and the entire cake, and agent 2 receives the other indivisible good. The Nash welfare of  $\mathcal{A}$  is  $(0.4 + 0.2) \times 0.499 = 0.2994$ . This is also the overall MNW allocation for this instance.

However, allocation  $\mathcal{A}$  is not weak EFM, because agent 1's bundle contains some cake that yields positive value to agent 2, and agent 2 is envious of agent 1. It is also worth noting that there is a simple envy-free and PO allocation for this instance: each agent gets one indivisible good and one half of the cake, with Nash welfare  $(0.4 + 0.1) \times (0.499 + 0.001) = 0.25$ .

Note that the compatibility of weak EFM and PO remains an open question even for the special case with indivisible goods and a single homogeneous divisible good (e.g. money), even though this case is well-studied when there is enough money.

### 7. Conclusion and future work

This work is concerned with fair division of a mixture of divisible and indivisible goods. To this end, we introduce the envy-freeness for mixed goods (EFM) fairness notion, which generalizes both EF and EF1 to the mixed goods setting. We show that an EFM allocation always exists for any number of agents. We also provide bounded protocols to compute an EFM allocation in special cases, and an  $\epsilon$ -EFM allocation in the general setting in time  $\text{poly}(n, m, 1/\epsilon)$ .

<sup>7</sup> In the case where the maximum Nash welfare is 0, an allocation is an MNW allocation if it gives positive utility to a set of agents of maximal size and moreover maximizes the product of utilities of the agents in that set.

It remains an important open question whether there exists a bounded, or even finite protocol in the RW model that computes an EFM allocation in the general setting for any number of players. With regard to  $\epsilon$ -EFM, although our algorithm runs in time  $\text{poly}(n, m, 1/\epsilon)$ , it remains an open question to design an algorithm that runs in time  $\text{poly}(n, m, \log(1/\epsilon))$ .

Besides envy-freeness, one could also generalize other fairness notions to the mixed goods setting. How well would this notion behave with mixed goods in terms of its existence and approximation? Overall, we believe that fair division in the mixed goods setting encodes a rich structure and creates a new research direction that deserves to be pursued for future work.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgements

The authors acknowledge the helpful comments by the reviewers. We are in particular grateful to an anonymous reviewer for providing Example 6.5 to us.

This work is supported in part by the Ministry of Education, Singapore, under its Academic Research Fund Tier 1 (RG23/20), and by an RGC grant (HKU 17203717E).

### References

- [1] Ahmet Alkan, Gabrielle Demange, David Gale, Fair allocation of indivisible goods and criteria of justice, *Econometrica* 59 (4) (1991) 1023–1039.
- [2] Noga Alon, Splitting necklaces, *Adv. Math.* 63 (3) (1987) 247–253.
- [3] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, Amin Saberi, Approximation algorithms for computing maximin share allocations, *ACM Trans. Algorithms* 13 (4) (2017) 52:1–52:28.
- [4] Georgios Amanatidis, Georgios Bimpas, Aris Filos-Ratsikas, Alexandros Hollender, Alexandros A. Voudouris, Maximum Nash welfare and other stories about EFX, in: *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 2020, pp. 24–30.
- [5] Haris Aziz, Simon Mackenzie, A discrete and bounded envy-free cake cutting protocol for any number of agents, in: *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2016, pp. 416–427.
- [6] Haris Aziz, Simon Mackenzie, A discrete and bounded envy-free cake cutting protocol for four agents, in: *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, 2016, pp. 454–464.
- [7] Siddharth Barman, Sanath Kumar Krishnamurthy, Rohit Vaish, Finding fair and efficient allocations, in: *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 2018, pp. 557–574.
- [8] Xiaohui Bei, Ning Chen, Xia Hua, Biaoshuai Tao, Endong Yang, Optimal proportional cake cutting with connected pieces, in: *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, 2012, pp. 1263–1269.
- [9] Xiaohui Bei, Guangda Huzhang, Warut Suksompong, Truthful fair division without free disposal, *Soc. Choice Welf.* 55 (2020) 523–545.
- [10] Xiaohui Bei, Zihao Li, Jinyan Liu, Shengxin Liu, Xinhang Lu, Fair division of mixed divisible and indivisible goods, in: *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020, pp. 1814–1821.
- [11] Xiaohui Bei, Shengxin Liu, Xinhang Lu, Hongao Wang, Maximin fairness with mixed divisible and indivisible goods, *CoRR*, arXiv:2002.05245, 2020, <https://arxiv.org/abs/2002.05245>.
- [12] Steven J. Brams, Alan D. Taylor, An envy-free cake division protocol, *Am. Math. Mon.* 102 (1) (1995) 9–18.
- [13] Steven J. Brams, Alan D. Taylor, *Fair Division: From Cake-Cutting to Dispute Resolution*, Cambridge University Press, 1996.
- [14] Simina Brânzei, Noam Nisan, The query complexity of cake cutting, *CoRR*, arXiv:1705.02946, 2017, <https://arxiv.org/abs/1705.02946>.
- [15] Johannes Brustle, Jack Dippel, Vishnu V. Narayan, Mashbat Suzuki, Adrian Vetta, One dollar each eliminates envy, in: *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 2020, pp. 23–39.
- [16] Eric Budish, The combinatorial assignment problem: approximate competitive equilibrium from equal incomes, *J. Polit. Econ.* 119 (6) (2011) 1061–1103.
- [17] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, Junxing Wang, The unreasonable fairness of maximum Nash welfare, *ACM Trans. Econ. Comput.* 7 (3) (2019) 12:1–12:32.
- [18] Katarína Cechlářová, Eva Pillárová, On the computability of equitable divisions, *Discrete Optim.* 9 (4) (2012) 249–257.
- [19] Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, EFX exists for three agents, in: *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 2020, pp. 1–19.
- [20] Yiling Chen, John K. Lai, David C. Parkes, Ariel D. Procaccia, Truth, justice, and cake cutting, *Games Econ. Behav.* 77 (1) (2013) 284–297.
- [21] Lester E. Dubins, Edwin Henry Spanier, How to cut a cake fairly, *Am. Math. Mon.* 68 (1) (1961) 1–17.
- [22] Jeff Edmonds, Kirk Pruhs, Balanced allocations of cake, in: *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006, pp. 623–634.
- [23] Shimon Even, Azaria Paz, A note on cake cutting, *Discrete Appl. Math.* 7 (3) (1984) 285–296.
- [24] Jugal Garg, Setareh Taki, An improved approximation algorithm for maximin shares, in: *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 2020, pp. 379–380.
- [25] Mohammad Ghodsi, MohammadTaghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, Hadi Yami, Fair allocation of indivisible goods: improvements and generalizations, in: *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 2018, pp. 539–556.
- [26] Daniel Halpern, Nisarg Shah, Fair division with subsidy, in: *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*, 2019, pp. 374–389.
- [27] Flip Klijn, An algorithm for envy-free allocations in an economy with indivisible objects and money, *Soc. Choice Welf.* 17 (2) (2000) 201–215.
- [28] David Kurokawa, Ariel D. Procaccia, Junxing Wang, Fair enough: guaranteeing approximate maximin shares, *J. ACM* 65 (2) (2018) 8:1–8:27.
- [29] A.A. Liapounoff, Sur les fonctions-vecteurs complètement additives, *Izv. Akad. Nauk SSSR, Ser. Mat.* 4 (6) (1940) 465–478, <http://mi.mathnet.ru/izv3907>.
- [30] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, Amin Saberi, On approximately fair allocations of indivisible goods, in: *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, 2004, pp. 125–131.



- [31] Eric S. Maskin, On the fair allocation of indivisible goods, in: George R. Feiwel (Ed.), *Arrow and the Foundations of the Theory of Economic Policy*, Palgrave Macmillan, UK, 1987, pp. 341–349, chapter 11.
- [32] Marc Meertens, Jos Potters, Hans Reijnierse, Envy-free and Pareto efficient allocations in economies with indivisible goods and money, *Math. Soc. Sci.* 44 (3) (2002) 223–233.
- [33] Hervé Moulin, *Fair Division and Collective Welfare*, MIT Press, 2003.
- [34] Hervé Moulin, Fair division in the Internet age, *Annu. Rev. Econ.* 11 (1) (2019) 407–441.
- [35] Benjamin Plaut, Tim Roughgarden, Almost envy-freeness with general valuations, *SIAM J. Discrete Math.* 34 (2) (2020) 1039–1068.
- [36] Ariel D. Procaccia, Cake cutting algorithms, in: Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, Ariel D. Procaccia (Eds.), *Handbook of Computational Social Choice*, Cambridge University Press, 2016, pp. 311–329, chapter 13.
- [37] Ariel D. Procaccia, An answer to fair division's most enigmatic question: technical perspective, *Commun. ACM* 63 (4) (2020) 118.
- [38] Ariel D. Procaccia, Junxing Wang, A lower bound for equitable cake cutting, in: *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 2017, pp. 479–495.
- [39] Jack Robertson, William Webb, *Cake-Cutting Algorithm: Be Fair If You Can*, A K Peters/CRC Press, 1998.
- [40] Alexander Rubchinsky, Brams-Taylor model of fair division for divisible and indivisible items, *Math. Soc. Sci.* 60 (1) (2010) 1–14.
- [41] Fedor Sandomirskiy, Erel Segal-Halevi, Fair division with minimal sharing, *CoRR*, arXiv:1908.01669, 2019, <http://arxiv.org/abs/1908.01669>.
- [42] Erel Segal-Halevi, Balázs R. Sziklai, Monotonicity and competitive equilibrium in cake-cutting, *Econ. Theory* 68 (2) (2019) 363–401.
- [43] Hugo Steinhaus, The problem of fair division, *Econometrica* 16 (1) (1948) 101–104.
- [44] Walter Stromquist, How to cut a cake fairly, *Am. Math. Mon.* 87 (8) (1980) 640–644.
- [45] Francis Edward Su, Rental harmony: Sperner's lemma in fair division, *Am. Math. Mon.* 106 (10) (1999) 930–942.
- [46] William Thomson, Introduction to the theory of fair allocation, in: Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, Ariel D. Procaccia (Eds.), *Handbook of Computational Social Choice*, Cambridge University Press, 2016, pp. 261–283, chapter 11.
- [47] Dietrich Weller, Fair division of a measurable space, *J. Math. Econ.* 14 (1) (1985) 5–17.