Truthful Fair Mechanisms for Allocating Mixed Divisible and Indivisible Goods

Zihao Li[1] , **Shengxin Liu**[2] , **Xinhang Lu**[3] and **Biaoshuai Tao**[4]
[1]Nanyang Technological University
[2]Harbin Institute of Technology, Shenzhen
[3]UNSW Sydney
[4]Shanghai Jiao Tong University
zihao004@e.ntu.edu.sg, sxliu@hit.edu.cn, xinhang.lu@unsw.edu.au, bstao@sjtu.edu.cn

Abstract

We study the problem of designing truthful and fair mechanisms when allocating a mixture of divisible and indivisible goods. We first show that there does not exist an EFM (envy-free for mixed goods) and truthful mechanism in general. This impossibility result holds even if there is only one indivisible good and one divisible good and there are only two agents. Thus, we focus on some more restricted settings. Under the setting where agents have binary valuations on indivisible goods and identical valuations on a single divisible good (e.g., money), we design an EFM and truthful mechanism. When agents have binary valuations over both divisible and indivisible goods, we first show there exist EFM and truthful mechanisms when there are only two agents or when there is a single divisible good. On the other hand, we show that the mechanism maximizing Nash welfare cannot ensure EFM and truthfulness simultaneously.

1 Introduction

Fair allocation problem considers how to fairly allocate scarce resources among interested agents (see excellent books or surveys by, e.g., [Brams and Taylor, 1995; Robertson and Webb, 1998; Moulin, 2019; Suksompong, 2021; Amanatidis *et al.*, 2022]). This problem has gained substantial attentions in various fields including computer science, mathematics, and economics, due to the ubiquity in numerous real-world scenarios (e.g., school choices [Abdulkadiroğlu *et al.*, 2005], course allocations [Budish and Cantillon, 2012], and allocating computational resources [Ghodsi *et al.*, 2011]).

The literature of fair allocation problem can be categorized by the type of resources being allocated. The first line of work studies the allocation of *divisible* goods, where the famous fairness criterion *envy-freeness* has been extensively studied [Foley, 1967; Aziz and Mackenzie, 2016]. In an envy-free allocation, each agent weakly prefers her own bundle than any other agent's bundle. The second group studies the allocation of *indivisible* goods, in which an envy-free allocation may fail to exist. A common practice to circumvent the issue is to consider relaxed notions such as *envy-freeness up to one good (EF1)* in which agent i's envy towards agent j could be eliminated if we (hypothetically) remove a good in agent j's bundle [Lipton *et al.*, 2004; Budish, 2011].

In addition to fairness, *truthfulness* is an important consideration, given that agents report their private preferences over the resources. Roughly speaking, a mechanism is said to be truthful if each agent cannot benefit by misreporting her preference. The truthfulness aspect of fair allocation has been addressed in a number of recent papers, e.g., [Bogomolnaia and Moulin, 2004; Caragiannis *et al.*, 2009; Mossel and Tamuz, 2010; Maya and Nisan, 2012; Kurokawa *et al.*, 2013; Aziz and Ye, 2014; Brânzei and Miltersen, 2015; Li *et al.*, 2015; Amanatidis *et al.*, 2016; Menon and Larson, 2017; Bei *et al.*, 2017; Bei *et al.*, 2020]. When the resource to be allocated is a cake (i.e., a *heterogeneous* divisible good), the seminal work by [Chen *et al.*, 2013] designed the first truthful envy-free mechanism when each agent's valuation is piecewise-uniform. On the other hand, very recently, [Bu *et al.*, 2023] showed that for piecewise-constant valuations (which is a more general type of valuation functions than piecewise-uniform functions), there does not exist a (deterministic) truthful and envy-free mechanism. For indivisible goods setting, [Amanatidis *et al.*, 2017] provided a characterization of truthful mechanisms for two agents, and further showed that truthfulness and EF1 are incompatible even for two agents and five indivisible goods. This negative result, however, does not hold any more for some restricted cases. With binary valuations, [Halpern *et al.*, 2020] and [Babaioff *et al.*, 2021] independently designed truthful and EF1 mechanisms by using maximum Nash welfare with lexicographic tie-breaking.

The aforementioned results paved the way for understanding the interplay between truthfulness and fairness for the fair allocation problem with either divisible or indivisible goods. However, when the resources contain a mixture of both, the study of designing truthful and fair allocation mechanisms is mostly absent, which is our focus in this paper. The only exception we know of is the work by [Goko *et al.*, 2022]. They concerned indivisible goods allocation and designed a truthful and fair mechanism that achieves envy-freeness by *subsidizing* each agent with at most 1, the maximum marginal value of each good for each agent.

We adopt a different perspective than that of [Goko *et al.*, 2022]. To be more specific, in our setting, the divisible and indivisible goods to be allocated are both *fixed* in advance [Bei *et al.*, 2021a; Bei *et al.*, 2021b; Bhaskar *et al.*, 2021; Lu *et al.*, 2023; Nishimura and Sumita, 2023]. In a setting with mixed divisible and indivisible goods (*mixed goods* for short), [Bei *et al.*, 2021a] proposed a new fairness notion called *envy-freeness for mixed goods (EFM)* that generalizes both envy-freeness and EF1, and showed constructively that an EFM allocation always exists for any number of agents. Can we go one step further by designing *truthful and EFM* mechanisms when allocating mixed goods?

1.1 Our Results

We study the problem of designing truthful and EFM mechanisms when allocating mixed divisible and indivisible goods to agents who have additive valuations over the goods. To the best of our knowledge, this is the first work examining the compatibility of truthfulness and EFM. Two variants of EFM are considered in this paper, and we use $\text{EFM}_{\geq 0}$ and $\text{EFM}_{>0}$ to distinguish them (see Section 2 for their formal definitions). Intuitively speaking, $\text{EFM}_{\geq 0}$ requires that the envy-free criterion is imposed even if the envied bundle contains a *positive amount* of divisible goods and the EF1 criterion is used otherwise. Slightly differently, $\text{EFM}_{>0}$ only requires to impose the envy-free criterion if the envied bundle contains divisible goods with *positive value*. It can be verified that $\text{EFM}_{\geq 0}$ implies $\text{EFM}_{>0}$. While in the following we mostly present our results regarding to $\text{EFM}_{>0}$, some of the results can be extended to the case of $\text{EFM}_{\geq 0}$.

We start by giving in Section 3 a strong impossibility result showing that truthfulness and $\text{EFM}_{>0}$ are incompatible even if there are only two agents and there are a single indivisible good and a single divisible good. Since we can normalize the valuations so that agents' valuations on the indivisible good are 0 or 1, as a corollary to the impossibility result, truthfulness and $\text{EFM}_{>0}$ are incompatible for two agents with *binary* valuations on indivisible goods. Truthfulness and $\text{EFM}_{>0}$, however, are compatible if we further restrict the expressiveness of agents' valuations on the divisible goods.

First, in Section 4, we design a truthful and $\text{EFM}_{>0}$ mechanism when agents have binary valuations over indivisible goods and an *identical* valuation (not necessarily binary) over a single divisible good. In addition, the allocations output by our mechanism satisfy some nice efficiency properties including leximin and Maximum Nash welfare (MNW). Next, in Section 5, we consider the case where agents have binary valuations over all the goods. Specifically, we design truthful and $\text{EFM}_{>0}$ mechanisms when (i) there are *two* agents (and an arbitrary number of goods), or (ii) the mixed goods consist of an arbitrary number of indivisible goods and a single divisible good. Technically speaking, in general, our mechanisms first make use of the truthful and EF1 mechanism of [Halpern *et al.*, 2020] to allocate the (binary) indivisible goods, and next design different methods to allocate the divisible good(s) in the three different scenarios described above.

2 Preliminaries

Let $[s] := \{1, \ldots, s\}$. We use $N = [n]$ to denote the set of n agents. The set of the goods is denoted by (G, D), where $G = \{g_1, \ldots, g_m\}$ is the set of m *indivisible goods* and $D = \{d_1, \ldots, d_{\overline{m}}\}$ is the set of $\overline{m}$ *divisible goods*. Each divisible good is *homogeneous*, meaning that an agent's value on each divisible good only depends on the *fraction* of this divisible good allocated to her. (We will formally define the valuations of the agents later.) Denote by $\mathcal{A} = (A_1, \ldots, A_n)$ an *allocation*, where we assign bundle A_i to agent i. Each A_i is composed by a pair $(G_i, \mathbf{x}_i)$, where G_i is a subset of the indivisible goods allocated to agent i and $\mathbf{x}_i = (x_{i1}, \ldots, x_{i\overline{m}})$ specifies how divisible goods are allocated to agent i—specifically, $x_{i\overline{k}}$ denotes the fraction of (homogeneous) divisible good $d_{\overline{k}}$ allocated to agent i. Naturally, an allocation $\mathcal{A} = ((G_1, \mathbf{x}_1), \ldots, (G_n, \mathbf{x}_n))$ must satisfy that $(G_1, \ldots, G_n)$ is a partition of G and that $\sum_{i=1}^{n} x_{i\overline{k}} = 1$ for each $\overline{k} = 1, \ldots, \overline{m}$ (in particular, we have assumed each divisible good has 1 unit of amount).

We assume that each agent $i \in N$ has an additive valuation function v_i and call $(v_1, \ldots, v_n)$ a *valuation profile*. That is, agent i's value on a bundle $(G_j, \mathbf{x}_j)$ is given by $v_i(G_j, \mathbf{x}_j) = \sum_{g_k \in G_j} v_i(g_k) + \sum_{\overline{k}=1}^{\overline{m}} x_{j\overline{k}} \cdot v_i(d_{\overline{k}})$, where $v_i(g_k)$ is agent i's value on the indivisible good g_k and $v_i(d_{\overline{k}})$ is agent i's value on the divisible good $d_{\overline{k}}$. We slightly abuse the notation by letting $v_i(G_j) = \sum_{g_k \in G_j} v_i(g_k)$ and $v_i(\mathbf{x}_j) = \sum_{\overline{k}=1}^{\overline{m}} x_{j\overline{k}} \cdot v_i(d_{\overline{k}})$. We say that agents' valuations are *binary* if $v_i(g_k) \in \{0, 1\}$ and $v_i(d_{\overline{k}}) \in \{0, 1\}$ for every i, k and $\overline{k}$. Correspondingly, we say that agents have *binary valuations on indivisible goods* if $v_i(g_k) \in \{0, 1\}$ for every i and k and agents have *binary valuations on divisible goods* if $v_i(d_{\overline{k}}) \in \{0, 1\}$ for every i and $\overline{k}$.

An allocation is *envy-free* if each agent believes (based on her own valuation) her allocated bundle is weakly more valuable than that of every other agent's. In our case with mixed goods, this means $v_i(G_i, \mathbf{x}_i) \geq v_i(G_j, \mathbf{x}_j)$ holds for every agents i and j. An envy-free allocation may not exist even if there are only indivisible goods (i.e., $D = \emptyset$). For allocating only indivisible goods, a commonly adopted relaxation of envy-freeness is *envy-freeness up to one item* (EF1), and it is well-known that an EF1 allocation always exists [Lipton *et al.*, 2004; Budish, 2011].

Definition 2.1 (EF1)**.** For $D = \emptyset$, given a valuation profile $(v_1, \ldots, v_n)$, an allocation $(G_1, \ldots, G_n)$ is EF1 if for any pair of $i, j \in N$, there exists a good $g \in G_j$ such that $v_i(G_i) \geq v_i(G_j \setminus \{g\})$.

With mixed goods, we adopt the fairness notion called *envy-freeness for mixed goods (EFM)* [Bei *et al.*, 2021a]. There are two variants of EFM, and we use $\text{EFM}_{\geq 0}$ and $\text{EFM}_{>0}$ to distinguish them. As we will see shortly, $\text{EFM}_{\geq 0}$ implies $\text{EFM}_{>0}$. We start by presenting the intuition of $\text{EFM}_{>0}$: For any pair of agents $i, j \in N$, agent i should not envy agent j, as stated in Point 2 of Definition 2.2 below, with the exception that the divisible part allocated to agent j is worthless to agent i, in which case the EF1 condition holds for the indivisible part (see Point 1 of Definition 2.2).

Definition 2.2 ($\text{EFM}_{>0}$)**.** Given a valuation profile $(v_1, \ldots, v_n)$, an allocation $((G_1, \mathbf{x}_1), \ldots, (G_n, \mathbf{x}_n))$ is $\text{EFM}_{>0}$ if the followings hold for any pair of $i, j \in N$.

1. If $v_i(\mathbf{x}_j) = 0$ and $G_j \neq \emptyset$, then there exists a good $g \in G_j$ such that $v_i(G_i, \mathbf{x}_i) \geq v_i(G_j \setminus \{g\}, \mathbf{x}_j)$.
2. Otherwise, $v_i(G_i, \mathbf{x}_i) \geq v_i(G_j, \mathbf{x}_j)$.

$\text{EFM}_{\geq 0}$, on the other hand, imposes envy-free condition as long as agent j's bundle has any positive amount of divisible goods, even if agent i values the divisible part at 0:

Definition 2.3 ($\text{EFM}_{\geq 0}$)**.** Given a valuation profile $(v_1, \dots, v_n)$, an allocation $((G_1, \mathbf{x}_1), \dots, (G_n, \mathbf{x}_n))$ is $\text{EFM}_{\geq 0}$ if the followings hold for any pair of $i, j \in N$.

1. If $\mathbf{x}_j = \mathbf{0}$ and $G_j \neq \emptyset$, then there exists a good $g \in G_j$ such that $v_i(G_i, \mathbf{x}_i) \geq v_i(G_j \setminus \{g\}, \mathbf{x}_j)$.
2. Otherwise, $v_i(G_i, \mathbf{x}_i) \geq v_i(G_j, \mathbf{x}_j)$.

It is easy to see that $\text{EFM}_{\geq 0}$ implies $\text{EFM}_{>0}$. It also directly follows from the above definitions that when there are only indivisible goods (i.e., $D = \emptyset$), $\text{EFM}_{>0}$ and $\text{EFM}_{\geq 0}$ reduce to EF1; when there are only divisible goods (i.e., $G = \emptyset$), $\text{EFM}_{>0}$ and $\text{EFM}_{\geq 0}$ reduce to envy-freeness.

A *mechanism* is a function $\mathcal{M}$ that maps the set of n valuation functions $(v_1, \dots, v_n)$ to an allocation $((G_1, \mathbf{x}_1), \dots, (G_n, \mathbf{x}_n))$. We only consider deterministic mechanisms in this paper. In the game-theoretical setting, each agent i submits a valuation v_i' to $\mathcal{M}$ which may or may not be her true valuation v_i. A mechanism $\mathcal{M}$ is $\text{EFM}_{>0}$/$\text{EFM}_{\geq 0}$ if, upon receiving every input $(v_1', \dots, v_n')$, it outputs an allocation that is $\text{EFM}_{>0}$/$\text{EFM}_{\geq 0}$ w.r.t. $(v_1', \dots, v_n')$. A mechanism $\mathcal{M}$ is *truthful* if it is each agent i's dominant strategy to truthfully report her valuation v_i. Formally, let v_i be the true valuation function of an arbitrary agent i and v_i' be an arbitrary valuation function, for any $n-1$ valuation functions $v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n$ of the remaining $n-1$ agents, we have $v_i(G_i, \mathbf{x}_i) \geq v_i(G_i', \mathbf{x}_i')$, where $(G_i, \mathbf{x}_i)$ is the bundle allocated to agent i by $\mathcal{M}$ when receiving the input $(v_1, \dots, v_{i-1}, v_i, v_{i+1}, \dots, v_n)$ and $(G_i', \mathbf{x}_i')$ is the bundle allocated to agent i by $\mathcal{M}$ when receiving the input $(v_1, \dots, v_{i-1}, v_i', v_{i+1}, \dots, v_n)$.

We make the standard *free-disposal* assumption, see, e.g., [Chen *et al.*, 2013; Bei *et al.*, 2017; Halpern *et al.*, 2020], which assumes that a good g_k (resp., $d_{\overline{k}}$) is discarded by the mechanism if $v_i(g_k) = 0$ (resp., $v_i(d_{\overline{k}}) = 0$) for all $i \in N$. Without this assumption, we may run into uninteresting technicality. This discussion along with all ommitted proofs can be found in the full version of our paper [Li *et al.*, 2023].

2.1 Maximum Nash Welfare and Leximin

We now proceed to review the concepts and some properties of *Maximum Nash Welfare* (MNW) allocations and *leximin* allocations, which will be useful in our paper.

The definitions of MNW and leximin allocations apply to general fair division settings. For each $i \in [n]$, if we are only allocating indivisible goods, then $A_i = G_i$; if we are only allocating divisible goods, then $A_i = \mathbf{x}_i$; for mixed indivisible and divisible goods, $A_i = (G_i, \mathbf{x}_i)$.

Definition 2.4 (MNW)**.** Given a valuation profile $(v_1, \dots, v_n)$, an allocation $(A_1, \dots, A_n)$ is a *Maximum Nash Welfare (MNW)* allocation if it first maximizes the number of the agents receiving positive values, i.e., $|\{i \in [n] : v_i(A_i) > 0\}|$, and, subject to this, maximizes the product of the positive utilities, i.e., $\prod_{i: v_i(A_i) > 0} v_i(A_i)$.

When allocating only divisible goods, an MNW allocation is always envy-free [Varian, 1974]. When allocating only indivisible goods, an MNW allocation is always EF1 [Caragiannis *et al.*, 2019]. However, with mixed goods, an MNW allocation may not be $\text{EFM}_{>0}$ [Bei *et al.*, 2021a].

Definition 2.5 (Leximin)**.** Given two vectors $\mathbf{s}_1, \mathbf{s}_2 \in \mathbb{R}^n$, let $\mathbf{s}_1'$ and $\mathbf{s}_2'$ be the vectors obtained by sorting $\mathbf{s}_1$ and $\mathbf{s}_2$ in ascending order respectively. We say that $\mathbf{s}_1$ *leximin-dominates* $\mathbf{s}_2$ if there exists $i \in [n]$ such that $\mathbf{s}_1'$ and $\mathbf{s}_2'$ are identical for the first $i-1$ entries and the i-th entry of $\mathbf{s}_1'$ is greater than the i-th entry of $\mathbf{s}_2'$. Given a valuation profile $(v_1, \dots, v_n)$, an allocation $(A_1, \dots, A_n)$ is a *leximin* allocation if $(v_1(A_1), \dots, v_n(A_n))$ is maximum in the total order induced by leximin-domination among all allocations.

In other words, a leximin allocation maximizes the minimum among the agents' utilities; among all such allocations, it considers those maximizing the second smallest utility, and so on.

When allocating indivisible goods *with binary valuations*, [Aziz and Rey, 2020] and [Halpern *et al.*, 2020] showed the equivalence of *Maximum Nash Welfare* (MNW) allocations and *leximin* allocations. In addition, [Halpern *et al.*, 2020] showed that the mechanism that outputs the MNW/leximin allocation with a consistent lexicographic tie-breaking rule is truthful. For the purpose of our paper, we state Halpern *et al.*'s mechanism (referred to as MNW^{tie}) below.

Theorem 2.6 ([Halpern *et al.*, 2020], MNW^{tie})**.** *For allocating only indivisible goods with binary valuations, there exists a truthful mechanism that always outputs an allocation that is both MNW and leximin.*

3 General Impossibility Results

Unfortunately, truthfulness and $\text{EFM}_{>0}$ are incompatible in general. In this section, we prove that there does not exist a truthful and $\text{EFM}_{>0}$ mechanism even under the very restricted settings where there are only one indivisible good and one divisible good and there are only two agents.

Our impossibility result for mixed divisible and indivisible goods holds for a minimum number of agents (which is 2) and a minimum number of goods (which is 2). This is in contrast with [Amanatidis *et al.*, 2017]'s 2-agent-5-good impossibility result for only indivisible goods.

Theorem 3.1. *There does not exist a truthful and $\text{EFM}_{>0}$ (and thus $\text{EFM}_{\geq 0}$) mechanism even when there are only two agents and the set of goods consists of one indivisible good and one divisible good.*

Proof. It suffices to prove the statement for $\text{EFM}_{>0}$. Consider the instance with two agents $\{1, 2\}$, one indivisible good and one divisible good. Both agents have value 1 on the indivisible good, and have value a and b on the divisible good, where $b > a > 1$.

Firstly, we prove that all the $\text{EFM}_{>0}$ allocations must allocate the indivisible good to agent 1. Suppose for the sake of contradiction that the indivisible good is allocated to agent 2.

To guarantee EFM$_{>0}$, agent 2 must get at least a fraction $\frac{b-1}{2b}$ from the divisible good in order to not envy agent 1, and agent 2 must get at most a fraction of $\frac{a-1}{2a}$ from the divisible good in order to avoid that agent 1 envies agent 2. This is impossible as $\frac{a-1}{2a} < \frac{b-1}{2b}$.

Therefore, the possible EFM$_{>0}$ allocations can be described as follows. Agent 1 receives the indivisible good and a fraction x of the divisible good, and agent 2 receives a fraction $1 - x$ of the divisible good. For the reasons similar as above, we must have $\frac{a-1}{2a} \le x \le \frac{b-1}{2b}$ to guarantee EFM$_{>0}$. If the mechanism outputs an allocation with $x = \frac{a-1}{2a}$, agent 1 can misreport her valuation by increasing a, which increases agent 1's received value as x increases. If the mechanism outputs an allocation with $x > \frac{a-1}{2a}$, agent 2 can misreport her valuation by decreasing b so that $\frac{a-1}{2a} < \frac{b-1}{2b} < x$. In this case, x will decrease and agent 2 receives more value. □

When there is only one indivisible good and one divisible good, we can assume w.l.o.g. that agents' valuations on the indivisible good are binary (as we did in the proof of Theorem 3.1) or agents' valuations on the divisible good are binary. This is because we can normalize the valuations of the agents. Thus, Theorem 3.1 implies the following corollary.

Corollary 3.2. *There does not exist a truthful and EFM$_{>0}$ (and thus EFM$_{\geq 0}$) mechanism even when there are two agents and agents have binary valuations on either the indivisible goods or the divisible goods.*

As a remark, although we focus on deterministic mechanisms here, Theorem 3.1 and Corollary 3.2 continue to hold for randomized mechanisms (that are universally EFM$_{>0}$ and truthful in expectation).[1] The proof is almost the same: We must allocate the indivisible good to agent 1 in order to guarantee universally EFM$_{>0}$, and x in the proof becomes the *expected* fraction of the divisible good allocated to agent 1.

The strong impossibility results suggest that truthfulness and EFM$_{>0}$ may only be compatible in more restrictive settings. We confirm our intuitions in the affirmative in the following sections by considering (1) the setting with a single divisible good of identical value to all agents and multiple indivisible goods on which agents have binary valuations (Section 4), and (2) the setting where agents' valuations are binary for *both* indivisible and divisible goods (Section 5).

4 Binary Valuations on Indivisible Goods and Identical Valuation on Single Divisible Good

This section considers the setting where agents' valuations on indivisible goods are binary and there is one divisible good on which agents have an identical valuation. This describes the natural scenario where we are allocating a set of indivisible goods and some amount of *money*. Here, the divisible good is just a sum of money, and an agent's value on each indivisible good is described by the amount of money the agent is willing to pay for the good. We will see that there exists a truthful and EFM$_{\geq 0}$ (and thus EFM$_{>0}$) mechanism under this setting.

[1]Loosely speaking, *universally EFM$_{>0}$* mechanisms randomize over deterministic EFM$_{>0}$ mechanisms. A randomized mechanism is said to be *truthful in expectation* if misreporting a valuation function cannot increase the expected utility of an agent.

Mechanism 1 A truthful EFM$_{\geq 0}$ mechanism for binary valuations on indivisible good and identical valuation on a single divisible good

use $\mathrm{MNW}^{\mathrm{tie}}$ (Theorem 2.6) to compute an MNW/leximin allocation $(G_1, \dots, G_n)$ of G
initialize $x_i \leftarrow 0$ for each agent i
while $y := 1 - \sum_{i=1}^{n} x_i > 0$:
  let $T_1 = \arg\min_{i \in [n]} \{v_i(G_i, x_i)\}$
  $\Delta \leftarrow \infty$
  if $T_1 \neq [n]$ **then**
    let $T_2 = \arg\min_{i \in [n] \setminus T_1} \{v_i(G_i, x_i)\}$
    $\Delta \leftarrow v_j(G_j, x_j) - v_i(G_i, x_i)$ for $i \in T_1, j \in T_2$
  for each $i \in T_1$, $x_i \leftarrow x_i + \min\{\frac{\Delta}{u}, \frac{y}{|T_1|}\}$
return allocation $((G_1, x_1), \dots, (G_n, x_n))$

For the ease of notation, we will use $x_1, \dots, x_n$ to denote the fractions of the (unique) divisible good allocated to the n agents. Thus, each agent's allocated share is denoted by (G_i, x_i). We will use u to denote each agent's value on the (unique) divisible good.

Our mechanism is shown in Mechanism 1. The mechanism first allocates the indivisible goods using $\mathrm{MNW}^{\mathrm{tie}}$ (Theorem 2.6) and then iteratively allocates the divisible good. In each iteration, the mechanism identifies a set T_1 of agents who receive minimum values in the current partial allocation. The mechanism then attempts to compensate agents in T_1 with some fraction of the unallocated divisible good. This is done by a "water-filling" process: We allocate the divisible good to the agents in T_1 at an equal rate, until the divisible good is fully allocated, in which case the mechanism terminates, or until the value received by each agent in T_1 reaches the second minimum value received among all agents (i.e., the value received by an agent in T_2), in which case the mechanism proceeds to the next iteration.

Theorem 4.1. *Mechanism 1 is EFM$_{\geq 0}$ and truthful. Moreover, it always outputs allocations that are leximin and MNW.*

The following three subsections aim to prove this theorem.

4.1 Mechanism 1 Is EFM$_{\geq 0}$

We first present some simple yet important observations.

Proposition 4.2. *Let $((G_1, x_1), \dots, (G_n, x_n))$ be the allocation output by Mechanism 1. The followings hold.*

1. *For any agent i, $v_i(G_i, x_i) = |G_i| + u \cdot x_i$.*
2. *For any agents i and j, $v_i(G_i, x_i) \geq v_j(G_i, x_i)$.*

Proof. Since agents have binary valuations on indivisible goods and the identical valuation u on the unique divisible good, we have $v_j(G_i, x_i) \leq |G_i| + u \cdot x_i$ for any agents i and j. Therefore, Point 1 implies Point 2, and it remains to show Point 1. To show Point 1, notice that $v_i(G_i, x_i) \neq |G_i| + u \cdot x_i$ is only possible when there exists $g \in G_i$ such that $v_i(g) = 0$. Suppose this is the case for the sake of contradiction. If there exists another agent j with $v_j(g) = 1$, then moving g from i's bundle to j's increases the Nash welfare, which contradicts to that $(G_1, \ldots, G_n)$ is an MNW allocation. If the good g is worth 0 to all agents, then g is discarded by the free-disposal assumption, which contradicts to $g \in G_1$. $\square$

Lemma 4.3. *Mechanism 1 always outputs EFM$_{\geq 0}$ (and thus EFM$_{>0}$) allocations.*

Proof. We will prove by induction that the partial allocation is EFM$_{\geq 0}$ after each while-loop iteration. For the base step, the MNW allocation $(G_1, \ldots, G_n)$ for the indivisible goods is EF1 by [Caragiannis *et al.*, 2019], and EFM$_{\geq 0}$ is satisfied since the allocation of the divisible good has not been started.

For the inductive step, suppose the partial allocation $((G_1, x_1), \ldots, (G_n, x_n))$ is EFM$_{\geq 0}$ before a while-loop iteration. After one while-loop iteration, each agent in T_1 receives an extra fraction $z := \min\{\frac{\Delta}{u}, \frac{y}{|T_1|}\}$ of the divisible good. It suffices to show that j does not envy i for any agent $i \in T_1$ and any other agent j. We discuss two cases: $j \in T_1$ and $j \notin T_1$.

If $j \in T_1$, we have $v_j(G_j, x_j) = v_i(G_i, x_i) \geq v_j(G_i, x_i)$ before the execution of this while-loop iteration, where the first equality is due to our definition of T_1 in Mechanism 1 and the second inequality is due to Point 2 of Proposition 4.2. This implies that agent j does not envy agent i before the while-loop iteration. Agent j will not envy agent i after the while-loop iteration, as both agents receive the same amount z of the divisible good, which is worth the same value $u \cdot z$ to both agents.

If $j \notin T_1$, we have $v_j(G_j, x_j) \geq \Delta + v_i(G_i, x_i) \geq \Delta + v_j(G_i, x_i)$ before the execution of this while-loop iteration, where, again, the first equality is due to our definition of T_1 in Mechanism 1 and the second inequality is due to Point 2 of Proposition 4.2. This implies that agent j will not envy agent i after the while-loop iteration if the portion of the divisible goods allocated to agent i is worth at most Δ. This is true as $u \cdot z \leq u \cdot \frac{\Delta}{u} = \Delta$. $\square$

4.2 Mechanism 1 Is Truthful

We first define a type of mechanisms called *water-filling mechanisms*. A water-filling mechanism starts from an allocation $(G_1, \ldots, G_n)$ of the indivisible goods and then proceeds to allocate the unique divisible good by the "water-filling process" defined by the while-loop in Mechanism 1. Our Mechanism 1 is a *particular* water-filling mechanism by specifying that the allocation of the indivisible goods $(G_1, \ldots, G_n)$ is output by $\mathrm{MNW}^{\mathrm{tie}}$.

Definition 4.4. Given a valuation profile $(v_1, \ldots, v_n)$, an allocation $((G_1, x_1), \ldots, (G_n, x_n))$ *satisfies the water-filling property* if

1. for any two agents i and j with $x_i > 0$ and $x_j > 0$, we have $v_i(G_i, x_i) = v_j(G_j, x_j)$, and
2. for any two agents i and j with $x_i = 0$ and $x_j > 0$, we have $v_i(G_i, x_i) \geq v_j(G_j, x_j)$.

It is straightforward to check that the allocation output by any water-filling mechanism satisfies the water-filling property. Given a valuation profile $\mathbf{v} = (v_1, \ldots, v_n)$ and an allocation $A = ((G_1, x_1), \ldots, (G_n, x_n))$ satisfying the water-filling property, we define the *potential* $\phi(A, \mathbf{v})$ by the "height of the water level": $\phi(A, \mathbf{v}) = v_i(G_i, x_i)$, where i is an arbitrary agent with $x_i > 0$. When $x_i = 0$ for all $i \in [n]$, $\phi(A, \mathbf{v}) = \min_{i \in [n]} v_i(G_i)$.

We will first prove a proposition, Proposition 4.7, which follows from that an MNW allocation of indivisible goods with binary valuations is always *Lorenz dominating* [Babaioff *et al.*, 2021]. Before stating the proposition, we will first state Babaioff *et al.*'s result.

Definition 4.5. Given a valuation profile $(v_1, \ldots, v_n)$, an allocation $(A_1, \ldots, A_n)$ is *Lorenz dominating* if, for any $k \in [n]$, the sum of the k smallest values of $v_1(A_1), \ldots, v_n(A_n)$ is weakly larger than the sum of the k smallest values of $v_1(A'_1), \ldots, v_n(A'_n)$ for any allocation $(A'_1, \ldots, A'_n)$.

A Lorenz dominating allocation may not exist. However, if it does, it is easy to see that a Lorenz dominating allocation is always leximin. [Babaioff *et al.*, 2021] (implicitly) proved the following theorem.

Proposition 4.6 ([Babaioff *et al.*, 2021])**.** *Consider the allocation of indivisible goods (i.e., $D = \emptyset$). For any binary valuation profile on indivisible goods, an MNW/leximin allocation is Lorenz dominating.*

In particular, Halpern *et al.*'s mechanism $\mathrm{MNW}^{\mathrm{tie}}$ always outputs Lorenz dominating allocations. We then get the following proposition from the Lorenz domination.

Proposition 4.7. *Fix a valuation profile $(v_1, \ldots, v_n)$. Mechanism 1 outputs an allocation A that maximizes the potential $\phi(A, \mathbf{v})$ among all allocation satisfying the water-filling property.*

Lemma 4.8. *Mechanism 1 is truthful.*

Proof. Let $\mathcal{M}$ be Mechanism 1. Consider a valuation profile $\mathbf{v} = (v_1, \ldots, v_n)$ and suppose agent 1 misreports her valuation to v'_1. Let $\mathbf{v}'$ be the valuation profile $(v'_1, v_2, \ldots, v_n)$. Let $A = ((G_1, x_1), \ldots, (G_n, x_n))$ be the allocation output by $\mathcal{M}$ when agent 1 reports truthfully, and let $A' = ((G'_1, x'_1), \ldots, (G'_n, x'_n))$ be the allocation output by $\mathcal{M}$ when agent 1 reports v'_1. By the truthfulness of $\mathrm{MNW}^{\mathrm{tie}}$, we have $v_1(G_1) \geq v_1(G'_1)$. Let $A^\dagger = ((G'_1, x^\dagger_1), \ldots, (G'_n, x^\dagger_n))$ be the allocation obtained by applying the water-filling process (the while-loop in $\mathcal{M}$) to the start-up allocation $(G'_1, \ldots, G'_n)$ with the true valuation profile $\mathbf{v}$ considered. Both A and $A^\dagger$ satisfy the water-filling property w.r.t. $\mathbf{v}$, and A' satsifies the water-filling property w.r.t. $\mathbf{v}'$.

If $x'_1 = 0$, we have $v_1(G_1, x_1) \geq v_1(G_1) \geq v_1(G'_1) = v_1(G'_1, x'_1)$, and the truthfulness of $\mathcal{M}$ holds trivially. Thus, we assume $x'_1 > 0$ from now on.

By Proposition 4.7, $\phi(A, \mathbf{v}) \geq \phi(A^\dagger, \mathbf{v})$. We have $v_1(G_1, x_1) \geq \phi(A, \mathbf{v})$ according to the definition of the potential function ϕ. To conclude that $\mathcal{M}$ is truthful, it suffices to show that $\phi(A^\dagger, \mathbf{v}) \geq v_1(G'_1, x'_1)$.

Since we have assumed $x'_1 > 0$, $v'_1(G'_1, x'_1) = \phi(A', \mathbf{v}')$. Let $\delta = v'_1(G'_1, x'_1) - v_1(G'_1, x'_1)$. It then remains to show that $\phi(A^\dagger, \mathbf{v}) + \delta \geq \phi(A', \mathbf{v}')$.

By Point 1 of Proposition 4.2, we have $v'_1(g) = 1$ for each $g \in G'_1$. Moreover, it is clear that δ equals to the number of goods g in G'_1 with $v_1(g) = 0$ and $v'_1(g) = 1$. In particular, $\delta \geq 0$. Now, consider the two water-filling processes corresponding to $\phi(A^\dagger, \mathbf{v})$ and $\phi(A', \mathbf{v}')$. When the "height of the water level" reaches $\phi(A^\dagger, \mathbf{v})$, the first process terminates, while an additional amount δ of water is yet to be split among one or more agents in the second process. Thus, the "height of the water level" for the second process can be further increased by at most δ. Therefore, $\phi(A^\dagger, \mathbf{v}) + \delta \geq \phi(A', \mathbf{v}')$. $\square$

4.3 Leximin and MNW

We show that the allocation output by Mechanism 1 is Lorenz dominating.

Proposition 4.9. *Given a valuation profile $\mathbf{v} = (v_1, \dots, v_n)$, the allocation $A = ((G_1, x_1), \dots, (G_n, x_n))$ output by Mechanism 1 is Lorenz dominating.*

Proof. Suppose there is an allocation $A' = ((G'_1, x'_1), \dots, (G'_n, x'_n))$ such that A does not Lorenz dominate A'. Let $a_1, \dots, a_n$ and $b_1, \dots, b_n$ be orderings of the n agents such that $v_{a_1}(G_{a_1}, x_{a_1}) \leq \dots \leq v_{a_n}(G_{a_n}, x_{a_n})$ and $v_{b_1}(G'_{b_1}, x'_{b_1}) \leq \dots \leq v_{b_n}(G'_{b_n}, x'_{b_n})$, respectively. Let k be the smallest index such that

$$\sum_{i=1}^{k} v_{a_i}(G_{a_i}, x_{a_i}) < \sum_{i=1}^{k} v_{b_i}(G'_{b_i}, x'_{b_i}). \tag{1}$$

We can first assume w.l.o.g. that A' satisfies the water-filling property. If not, we can adjust A' by applying the water-filling process to the start up allocation $(G'_1, \dots, G'_n)$. It is easy to see that the adjusted allocation Lorenz dominates the original allocation.

By Proposition 4.7, we have $\phi(A, \mathbf{v}) \geq \phi(A', \mathbf{v})$. Let ℓ and ℓ' be the numbers of agents with $x_i > 0$ and $x'_i > 0$, respectively. We must have $k > \ell'$. Otherwise, $\sum_{i=1}^{k} v_{b_i}(G'_{b_i}, x'_{b_i}) = k \cdot \phi(A', \mathbf{v}) \leq k \cdot \phi(A, \mathbf{v})$. Since $\phi(A, \mathbf{v})$ is a lower bound to each $v_i(G_i, x_i)$, Equation (1) cannot be true. Since $k > \ell'$, we have

$$\sum_{i=1}^{k} v_{b_i}(G'_{b_i}, x'_{b_i}) = u + \sum_{i=1}^{k} v_{b_i}(G'_{b_i}). \tag{2}$$

We consider two cases: $\ell \leq k$ and $\ell > k$. If $\ell \leq k$, we must have $\sum_{i=1}^{k} v_{a_i}(G_{a_i}, x_{a_i}) = u + \sum_{i=1}^{k} v_{a_i}(G_{a_i})$. This, together with Equations (1) and (2), implies $\sum_{i=1}^{k} v_{a_i}(G_{a_i}) < \sum_{i=1}^{k} v_{b_i}(G'_{b_i})$. However, this contradicts to Proposition 4.6.

If $\ell > k$, we have $v_{a_i}(G_{a_i}, x_{a_i}) = \phi(A, \mathbf{v}) = v_{a_k}(G_{a_k}, x_{a_k}) < v_{b_k}(G'_{b_k}, x'_{b_k}) \leq v_{b_i}(G'_{b_i}, x_{b'_i})$ for each $i = k+1, \dots, \ell$, where the only strict inequality in the middle is due to our assumption that k is the smallest index. This implies Equation (1) continues to hold if we summing over the first ℓ terms instead of the first k terms:

$$\sum_{i=1}^{\ell} v_{a_i}(G_{a_i}, x_{a_i}) < \sum_{i=1}^{\ell} v_{b_i}(G'_{b_i}, x'_{b_i}). \tag{3}$$

Since $\ell > k > \ell'$, we have $\sum_{i=1}^{\ell} v_{a_i}(G_{a_i}, x_{a_i}) = u + \sum_{i=1}^{\ell} v_{a_i}(G_{a_i})$ and $\sum_{i=1}^{\ell} v_{b_i}(G'_{b_i}, x'_{b_i}) = u + \sum_{i=1}^{\ell} v_{b_i}(G'_{b_i})$. These, together with Equation (3), give a contradiction to Proposition 4.6 again. $\square$

Mechanism 2 A truthful EFM$_{>0}$ mechanism for binary valuations on both divisible and indivisible goods with two agents

use $\mathrm{MNW}^{\mathrm{tie}}$ (Theorem 2.6) to compute an MNW/leximin allocation (G_1, G_2) of G
initialize $\mathbf{x}_i \leftarrow \mathbf{0}$ for each agent $i \in \{1, 2\}$
if $v_1(G_1) \neq v_2(G_2)$ **then**
  let $i^* = \arg\min_{i \in \{1,2\}} \{v_i(G_i)\}$
  if $\exists \bar{k}^* \in [\bar{m}]$ s.t. $v_{i^*}(d_{\bar{k}^*}) = 1$ **then** $x_{i^* \bar{k}^*} \leftarrow 1$
for each $\bar{k} = 1$ to $\bar{m}$ with $\bar{k} \neq \bar{k}^*$ **do**
  if $v_1(d_{\bar{k}}) = v_2(d_{\bar{k}}) = 1$ **then** $x_{1\bar{k}} \leftarrow 0.5$, $x_{2\bar{k}} \leftarrow 0.5$
  else
    if $v_1(d_{\bar{k}}) = 1$ **then** $x_{1\bar{k}} \leftarrow 1$
    if $v_2(d_{\bar{k}}) = 1$ **then** $x_{2\bar{k}} \leftarrow 1$
return allocation $((G_1, \mathbf{x}_1), (G_2, \mathbf{x}_2))$

Lorenz domination leads to the following lemma.

Lemma 4.10. *Mechanism 1 always outputs allocations that are both leximin and MNW.*

5 Binary Valuations for Both Divisible and Indivisible Goods

In this section, we focus on the setting where agents' valuations on mixed goods are binary. We start from the simplest case where there are only two agents. We first show that no deterministic mechanism can ensure truthfulness and always output an MNW/leximin allocation. We illustrate this by Example 5.1. This is in sharp contrast to Halpern et al.'s mechanism $\mathrm{MNW}^{\mathrm{tie}}$ and Mechanism 1.

Example 5.1. Consider the instance with 2 agents and an indivisible good g_1 and two divisible goods d_1 and d_2. We have $v_1(g_1) = v_2(g_1) = 1$. For d_1, let $v_1(d_1) = 1$ and $v_2(d_1) = 0$; for d_2, let $v_1(d_2) = 0$ and $v_2(d_2) = 1$.

In the allocation returned by a mechanism that can always output an MNW/leximin allocation, each agent i receives the corresponding divisible good d_i and one of them also receives the indivisible good g_1. Without loss of generality, we assume agent 1 receives g_1 and d_1. If the actual valuation of agent 1 is positive towards all three goods and she reports her valuation truthfully, she will only receive d_1 and a half of d_2 to achieve an MNW/leximin allocation. Thus, she has an incentive to misreport the valuation we stated in the table to earn a benefit.

Remark 5.2. Even when allowing randomized mechanisms, we can also show there is no truthful mechanism that always outputs an MNW/leximin allocation with a general number of agents (where the proof can be found in the full version).

Such an example shows that directly returning an MNW allocation cannot work for our setting. Instead, we adopt the mechanism $\mathrm{MNW}^{\mathrm{tie}}$ (Theorem 2.6) for only indivisible goods and design a truthful mechanism which can always output an EFM$_{>0}$ allocation when there are only two agents.

Mechanism 2 first allocates indivisible goods according to $\mathrm{MNW}^{\mathrm{tie}}$. If the valuations of two agents over their own bundles are different, we can allocate one valuable divisible good to the one with a smaller value to eliminate the possible envy, if such a divisible good exists. Each remaining divisible good

Mechanism 3 A truthful $\text{EFM}_{>0}$ mechanism for binary valuations on indivisible goods and a single divisible good

```
1: use MNW^tie (Theorem 2.6) to compute an MNW/leximin allocation (G_1, ..., G_n) of G
2: initialize x_i ← 0 for each agent i ∈ [n]
3: let T = arg min_{i∈[n]} {v_i(G_i)}
4: for each i ∈ T do
5:   x_{i1} ← 1/|T|
6: return allocation ((G_1, x_1), ..., (G_n, x_n))
```

is allocated evenly to all the agents who value it positively. We have the following result for Mechanism 2:

Theorem 5.3. *Mechanism 2 is $\text{EFM}_{>0}$ and truthful.*

Proof. We first show the output allocation $\mathcal{A} = (A_1, A_2)$ is $\text{EFM}_{>0}$. Since the allocation (G_1, G_2) output by MNW^{tie} is EF1, so it is also $\text{EFM}_{>0}$. If $v_1(G_1) = v_2(G_2)$, such an allocation is envy-free. If $v_1(G_1) \neq v_2(G_2)$, after Steps 4-5, either the present allocation is envy-free, or $v_{i^*}(D) = 0$, which means the output allocation is always $\text{EFM}_{>0}$ for agent v_{i^*} no matter how divisible goods are allocated. In the remaining steps, each divisible good is allocated evenly to the agents who positively value them, which will not destroy the envy-freeness. Hence the final allocation is $\text{EFM}_{>0}$.

We next show Mechanism 2 is truthful. We assume we receive the allocation $\mathcal{A} = (A_1, A_2)$ given the true valuation profile (v_1, v_2) and the allocation $\mathcal{A}' = (A'_1, A'_2)$ when agent 1 misreports her valuation to v'_1.

Since the allocation of divisible goods after Step 6 is to evenly allocate all goods to the agents who value them, misreporting can make no benefit in this part. For Steps 4-5, the decision depends only on the valuations of the indivisible goods, there is also no gain from misreporting the valuations on divisible goods. So the only possible way to benefit is to misreport the valuations on indivisible goods.

As MNW^{tie} is truthful, we get $v_1(G_1) \geq v_1(G'_1)$. The first case is when $v_1(G_1) \geq v_1(G'_1) + 1$. By misreporting, the largest benefit is from the change of i^* from 2 to 1 during Steps 4-5. Such benefit is upper bounded by $2 \times 0.5 = 1$, since agent 1 will receive half of this good if she values the allocated good positively and such a good is allocated after Step 6. Hence there is no need to misreport in this case.

For the second case where $v_1(G_1) = v_1(G'_1)$, we first have $v'_1(G'_1) \geq v_1(G'_1)$ since the leximin allocation with binary valuations. The benefit of the misreporting is also from the change of i^* during Steps 4-5. If i^* is 2 under (v_1, v_2) and does not exist under (v'_1, v_2), then $v_2(G'_2) = v'_1(G'_1) \geq v_1(G'_1) = v_1(G_1) > v_2(G_2)$, this violates the leximin property of (G_1, G_2) under (v_1, v_2). If i^* is 2 under (v_1, v_2) and 1 under (v'_1, v_2), the leximin property of (G_1, G_2) under (v_1, v_2) is also violated from the similar argument. Similar analysis also holds for the case where i^* does not exist under (v_1, v_2) and is 1 under (v'_1, v_2). Since all three cases which can benefit agent 1 cannot occur, misreporting the valuations has no benefit. Thus, this mechanism is truthful. □

If we consider the setting with more than two agents, we can also design an $\text{EFM}_{>0}$ and truthful mechanism for allocating indivisible goods and a single divisible good, which is shown in Mechanism 3. It is worth noting that, in this setting, agents value the divisible good at either 1 or 0, which is different from the setting in Section 4 where all agents have an identical value over the single divisible good. In this mechanism, after allocating all indivisible goods according to MNW^{tie}, we allocate the only divisible good evenly to all agents with the smallest $v_i(G_i)$.

Theorem 5.4. *Mechanism 3 is $\text{EFM}_{\geq 0}$ and truthful.*

Proof. Since MNW^{tie} returns an EF1 allocation and the remaining of our mechanism is just to allocate the single divisible good evenly to agents with the smallest $v_i(G_i)$, no agent in T will envy others in T. Because there is only one divisible good, no agent in $[n] \setminus T$ will envy the agents in T after allocating this good. Thus, the output allocation is $\text{EFM}_{\geq 0}$.

We then prove Mechanism 3 is truthful. We assume we get the allocation $\mathcal{A} = (A_1, \ldots, A_n)$ given the true valuation profile $(v_1, \ldots, v_n)$, and the allocation $\mathcal{A}' = (A'_1, \ldots, A'_n)$ when agent 1 misreports her valuation to v'_1.

From the truthfulness of MNW^{tie}, we have $v_1(G_1) \geq v_1(G'_1)$. Since there is only one divisible good, there is no incentive to misreport if $v_1(G_1) \geq v_1(G'_1) + 1$. We then consider the case where $v_1(G_1) = v_1(G'_1)$. If agent 1 can benefit from misreporting, this means that after misreporting, agent 1 is in the set T at Step 3, and when agent 1 truthfully reports v_1, either she is not in the set T or the size of the set T is larger than that after misreporting.

Since MNW^{tie} is leximin, under binary valuations, we have $v'_1(G'_1) \geq v_1(G'_1) = v_1(G_1)$. Because, after misreporting the valuation, agent 1 is in the set T which contains all agents with the smallest $v_i(G_i)$. Both two cases mentioned above violate the leximin property of the allocation $(G_1, \ldots, G_n)$ of G under the valuation $(v_1, \ldots, v_n)$. Thus, this mechanism is truthful. □

Remark 5.5. This mechanism is no longer $\text{EFM}_{\geq 0}$ and $\text{EFM}_{>0}$ when there is more than one divisible good, because some envy may occur from the agent outside the set T to the agent in T after allocating multiple divisible goods.

6 Conclusion

In this paper, we have studied truthful and fair (i.e., EFM) mechanisms when allocating mixed divisible and indivisible goods. Our strong impossibility result shows that truthfulness and EFM are incompatible even if there are only two agents and two goods. We then designed truthful and EFM mechanisms in various restricted settings.

In future research, it would be intriguing to completely determine the compatibility between truthfulness and EFM when agents have binary valuations over all goods. Given by our impossibility result, another future direction is to consider weaker notions of truthfulness, e.g., *maximin strategy-proofness* [Brams *et al.*, 2006], *not obviously manipulation (NOM)* [Troyan and Morrill, 2020; Ortega and Segal-Halevi, 2022], and *risk-averse truthfulness* [Bu *et al.*, 2023].

Acknowledgments

Shengxin Liu is the corresponding author. The research of Biaoshuai Tao was supported by the National Natural Science Foundation of China (Grant No. 62102252). The research of Shengxin Liu was partially supported by the National Natural Science Foundation of China (Grant No. 62102117), by the Shenzhen Science and Technology Program (Grant No. RCBS20210609103900003), and by the Guangdong Basic and Applied Basic Research Foundation (Grant No. 2023A1515011188), and was partially sponsored by CCF-Huawei Populus Grove Fund (Grant No. CCF-HuaweiLK2022005). Xinhang Lu was partially supported by ARC Laureate Project FL200100204 on "Trustworthy AI".

References

[Abdulkadiroğlu *et al.*, 2005] Atila Abdulkadiroğlu, Parag A. Pathak, and Alvin E. Roth. The new york city high school match. *American Economic Review*, 95(2):364–367, 2005.

[Amanatidis *et al.*, 2016] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. On truthful mechanisms for maximin share allocations. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 31–37, 2016.

[Amanatidis *et al.*, 2017] Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 545–562, 2017.

[Amanatidis *et al.*, 2022] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. Fair division of indivisible goods: A survey. *CoRR*, 2202.08713v1, 2022. Preliminary versions in IJCAI'22 and SIGecom Exchanges.

[Aziz and Mackenzie, 2016] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 416–427, 2016.

[Aziz and Rey, 2020] Haris Aziz and Simon Rey. Almost group envy-free allocation of indivisible goods and chores. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 39–45, 2020.

[Aziz and Ye, 2014] Haris Aziz and Chun Ye. Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In *Proceedings of the 10th International Conference on Web and Internet Economics (WINE)*, pages 1–14, 2014.

[Babaioff *et al.*, 2021] Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5119–5126, 2021.

[Bei *et al.*, 2017] Xiaohui Bei, Ning Chen, Guangda Huzhang, Biaoshuai Tao, and Jiajun Wu. Cake cutting: Envy and truth. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3625–3631, 2017.

[Bei *et al.*, 2020] Xiaohui Bei, Guangda Huzhang, and Warut Suksompong. Truthful fair division without free disposal. *Social Choice and Welfare*, 55:523–545, 2020.

[Bei *et al.*, 2021a] Xiaohui Bei, Zihao Li, Jinyan Liu, Shengxin Liu, and Xinhang Lu. Fair division of mixed divisible and indivisible goods. *Artificial Intelligence*, 293:103436, 2021.

[Bei *et al.*, 2021b] Xiaohui Bei, Shengxin Liu, Xinhang Lu, and Hongao Wang. Maximin fairness with mixed divisible and indivisible goods. *Autonomous Agents and Multi-Agent Systems*, 35(2):34:1–34:21, 2021.

[Bhaskar *et al.*, 2021] Umang Bhaskar, A. R. Sricharan, and Rohit Vaish. 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)*, pages 1:1–1:23, 2021.

[Bogomolnaia and Moulin, 2004] Anna Bogomolnaia and Hervé Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.

[Brams and Taylor, 1995] Steven J. Brams and Alan D. Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18, 1995.

[Brams *et al.*, 2006] Steven J. Brams, Michael A. Jones, and Christian Klamler. Better ways to cut a cake. *Notices of the AMS*, 53(11):1314–1321, 2006.

[Brânzei and Miltersen, 2015] Simina Brânzei and Peter B. Miltersen. A dictatorship theorem for cake cutting. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 482–488, 2015.

[Bu *et al.*, 2023] Xiaolin Bu, Jiaxin Song, and Biaoshuai Tao. On existence of truthful fair cake cutting mechanisms. *Artificial Intelligence*, 319:103904, 2023.

[Budish and Cantillon, 2012] Eric Budish and Estelle Cantillon. The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *American Economic Review*, 102(5):2237–71, 2012.

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

[Caragiannis *et al.*, 2009] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On low-envy truthful allocations. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT)*, pages 111–119, 2009.

[Caragiannis *et al.*, 2019] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg

Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3):12:1–12:32, 2019.

[Chen *et al.*, 2013] Yiling Chen, John K. Lai, David C. Parkes, and Ariel D. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1):284–297, 2013.

[Foley, 1967] Duncan Karl Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7(1):45–98, 1967.

[Ghodsi *et al.*, 2011] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2011.

[Goko *et al.*, 2022] Hiromichi Goko, Ayumi Igarashi, Yasushi Kawase, Kazuhisa Makino, Hanna Sumita, Akihisa Tamura, Yu Yokoi, and Makoto Yokoo. Fair and truthful mechanism with limited subsidy. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 534–542, 2022.

[Halpern *et al.*, 2020] Daniel Halpern, Ariel D. Procaccia, Alexandros Psomas, and Nisarg Shah. Fair division with binary valuations: One rule to rule them all. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*, pages 370–383, 2020.

[Kurokawa *et al.*, 2013] David Kurokawa, John Lai, and Ariel D. Procaccia. How to cut a cake before the party ends. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 2013.

[Li *et al.*, 2015] Minming Li, Jialin Zhang, and Qiang Zhang. Truthful cake cutting mechanisms with externalities: Do not make them care for others too much! In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 589–595, 2015.

[Li *et al.*, 2023] Zihao Li, Shengxin Liu, Xinhang Lu, and Biaoshuai Tao. Truthful fair mechanisms for allocating mixed divisible and indivisible goods. *CoRR*, 2305.09206v1, 2023.

[Lipton *et al.*, 2004] Richard Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 125–131, 2004.

[Lu *et al.*, 2023] Xinhang Lu, Jannik Peters, Haris Aziz, Xiaohui Bei, and Warut Suksompong. Approval-based voting with mixed goods. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023. Forthcoming.

[Maya and Nisan, 2012] Avishay Maya and Noam Nisan. Incentive compatible two player cake cutting. In *Proceedings of the 8th International Workshop on Internet and Network Economics (WINE)*, pages 170–183, 2012.

[Menon and Larson, 2017] Vijay Menon and Kate Larson. Deterministic, strategyproof, and fair cake cutting. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 352–358, 2017.

[Mossel and Tamuz, 2010] Elchanan Mossel and Omer Tamuz. Truthful fair division. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*, pages 288–299, 2010.

[Moulin, 2019] Hervé Moulin. Fair division in the internet age. *Annual Review of Economics*, 11(1):407–441, 2019.

[Nishimura and Sumita, 2023] Koichi Nishimura and Hanna Sumita. Envy-freeness and maximum Nash welfare for mixed divisible and indivisible goods. *CoRR*, abs/2302.13342v1, 2023.

[Ortega and Segal-Halevi, 2022] Josué Ortega and Erel Segal-Halevi. Obvious manipulations in cake-cutting. *Social Choice and Welfare*, pages 1–20, 2022.

[Robertson and Webb, 1998] Jack Robertson and William Webb. *Cake-Cutting Algorithm: Be Fair If You Can*. A K Peters/CRC Press, 1998.

[Suksompong, 2021] Warut Suksompong. Constraints in fair division. *ACM SIGecom Exchanges*, 19(2):46–61, 2021.

[Troyan and Morrill, 2020] Peter Troyan and Thayer Morrill. Obvious manipulations. *Journal of Economic Theory*, 185:104970, 2020.

[Varian, 1974] Hal R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.