Computing Envy-Freeable Allocations with Limited Subsidies

Ioannis Caragiannis[1(✉)] and Stavros D. Ioannidis[2]

[1] Department of Computer Science, Aarhus University, Aarhus, Denmark
iannis@cs.au.dk

[2] Department of Informatics, King's College London, London, UK
stavros.ioannidis@kcl.ac.uk

Abstract. Fair division has emerged as a very hot topic in EconCS research, and envy-freeness is among the most compelling fairness concepts. An allocation of indivisible items to agents is envy-free if no agent prefers the bundle of any other agent to his own in terms of value. As envy-freeness is rarely a feasible goal, there is a recent focus on relaxations of its definition. An approach in this direction is to complement allocations with payments (or subsidies) to the agents. A feasible goal then is to achieve envy-freeness in terms of the total value an agent gets from the allocation and the subsidies.

We consider the natural optimization problem of computing allocations that are *envy-freeable* using the minimum amount of subsidies. As the problem is NP-hard, we focus on the design of approximation algorithms. On the positive side, we present an algorithm which, for a constant number of agents, approximates the minimum amount of subsidies within any required accuracy, at the expense of a graceful increase in the running time. On the negative side, we show that, for a superconstant number of agents, the problem of minimizing subsidies for envy-freeness is not only hard to compute exactly (as a folklore argument shows) but also, more importantly, hard to approximate.

Keywords: Fair division · Indivisible goods · Subsidy minimization · Approximation algorithms

1 Introduction

Fairly dividing goods among people is an extremely important quest since antiquity. Today, fair division is a flourishing area of research in computer science, economics, and political science and *envy-freeness* is considered as the ultimate fairness concept [26]. Following a research trend that is very popular recently, we consider allocation problems with indivisible items. An allocation of items to agents is envy-free if no agent prefers the bundle of items allocated to some other agent to her own. Traditionally, agents' preferences are based on cardinal valuations they have for the items.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, pp. 522–539, 2022.
https://doi.org/10.1007/978-3-030-94676-0_29

Unfortunately, with indivisible items, envy-freeness is rarely a feasible goal. For example, no such allocation exists in the embarrassingly simple case with a single item and two agents with some value for it. Recently proposed relaxations of envy-freeness aim to serve as useful alternative fairness notions. In a line of research that emerged very recently, allocations are complemented with payments (or *subsidies*) to the agents [7,18]. Now, envy-freeness dictates that no agent prefers the allocation and payment of another agent to hers, and becomes a feasible goal. However, important questions arise related to the sparing use of money.

In this paper, we follow an optimization approach. We define and study the optimization problem SMEF (standing for Subsidy Minimization for Envy-Freeness). Given an allocation problem consisting of items and agents with valuations for the items, SMEF asks for an allocation that is envy-freeable using the minimum total amount of subsidies.

SMEF is NP-hard; this follows by the NP-hardness of deciding whether a given allocation problem has an envy-free allocation or not. Thus, we resort to *approximation algorithms* for SMEF. As multiplicative approximation guarantees are hopeless, our aim is to design algorithms that run in polynomial-time and compute an allocation that is envy-freeable with an amount of subsidies that does not exceed the minimum possible amount of subsidies (denoted χ) by much. In particular, we use the total valuation of all agents for all goods (denoted by sum v) as a benchmark and seek allocations that are envy-freeable with an amount of at most $\chi + \rho \cdot \text{sum } v$ as subsidies. The goal for the approximation guarantee ρ of an algorithm is to be as small as possible.

We initiate the study of SMEF and present two results. On the positive side, we design an algorithm that achieves an arbitrary low approximation guarantee of $\epsilon > 0$. When applied to allocation instances with a constant number of agents, the algorithm uses *dynamic programming* and runs in time that is polynomial in the number of items and $1/\epsilon$. On the negative side, we show that, in general, SMEF is not only hard to solve exactly, but also hard to approximate within a small constant. Unlike the folklore reduction[1] for proving hardness of envy-freeness, our proof uses a novel *approximation-preserving* reduction. Besides separating the general case from that with constantly many agents, our negative result indicates that achieving good approximation guarantees will be a challenging goal.

1.1 Related Work

The concept of envy-freeness was formally introduced by Foley [15] and Varian [29]. As envy-freeness may not be achievable when goods are indivisible, recent research has focused on defining approximations of envy-freeness. These include envy-freeness up to one good [8], envy-freeness up to any good [10], epistemic envy-freeness [4], and more. Still, achieving even them in polynomial time

[1] Notice that deciding whether an envy-free allocation exists for two agents with identical item valuations requires solving Partition, a well-known NP-hard problem [16].

can be challenging, and recent work has focused on approximation algorithms; see, e.g., [2,6,9,11,12,21,25].

The approach of mixing allocations with payments either from or to the agents has been extensively considered in the economics literature. A typical example is the rent division problem, where n items (rooms) and a fixed rent have to be divided among n agents in an envy-free manner [27,28]. Compensations to the agents were first considered by Maskin [22]. Subsequent papers consider unit-demand allocation problems, where each agent can get at most one item; see, e.g., [1]. Aragones [3] and Klijn [20] give polynomial-time algorithms that compute allocations and payments. More general models are studied by Haake et al. [17] and Meertens et al. [23].

In the AI literature, Chevaleyre et al. [13] consider allocation problems and monetary transfers between the agents. In a model that is the closest to ours, Halpern and Shah [18] aim to bound the amount of external subsidies assuming that all agent valuations for goods are in $[0, 1]$. Among several results, they conjectured that subsidies of $n-1$ suffice; an even stronger version of the conjecture was proved very recently by Brustle et al. [7].

1.2 Roadmap

The rest of the paper is structured as follows. We begin with preliminary definitions in Sect. 2. Our approximation algorithm is presented in Sect. 3 and our result on the hardness of approximation for SMEF is presented in Sect. 4. We conclude in Sect. 5.

2 Preliminaries

We consider allocation instances with a set M of m items and a set N of n agents. Each agent $i \in N$ has a valuation function $v_i : M \rightarrow \mathbb{R}_{\geq 0}$ over the items.[2] With some abuse of notation, we use $v_i(B)$ to denote the valuation of agent i for the set (or *bundle*) of items B. Valuations are additive, i.e., $v_i(B) = \sum_{g \in B} v_i(g)$. An allocation is simply a partition $X = (X_1, X_2, ..., X_n)$ of the items of M into n disjoint bundles, where agent $i \in N$ is supposed to get the bundle X_i. We use the abbreviations sum $v = \sum_{i \in N} v_i(M)$ and $\max v = \max_{i \in N} v_i(M)$.

As usual, we define the *social welfare* of an allocation $X = (X_1, ..., X_n)$ to be $\mathrm{SW}(X, v) = \sum_{i \in N} v_i(X_i)$. An allocation $X = (X_1, X_2, ..., X_n)$ is *envy-free* if $v_i(X_i) \geq v_i(X_j)$ for every pair of agents i and j. Informally, envy-freeness requires that no agent envies the bundle allocated to any other agent compared to her own.

For an allocation $X = (X_1, ..., X_n)$ in an instance with agent valuations v, the *envy graph* $\mathrm{EG}(X, v)$, introduced by Lipton et al. [21], is an edge-weighted complete directed graph that has a node for each agent and the weight of the

[2] In our exposition, we assume that valuations are non-negative, even though our positive result can be extended to work without this assumption, in the model of [5] where items can be goods or chores.

directed edge (i, j) represents the "envy" of agent i for agent j. Using $G = \mathrm{EG}(X, v)$ and $\mathrm{wgt}_G(i, j)$ for the weight of the directed edge from node i to node j in the envy graph $\mathrm{EG}(X, v)$, we define $\mathrm{wgt}_G(i, j) = v_i(X_j) - v_i(X_i)$.

Following the modelling assumptions of [18], we also consider *payments* (or *subsidies*) to the agents, represented by a payment vector $\pi = \langle \pi_1, ..., \pi_n \rangle$ with non-negative entries, i.e., $\pi_i \geq 0$ for every agent $i \in N$. Below, we use the terms "payment" and "subsidy" interchangeably. Now, we say that the pair (X, π) of the allocation X and payment vector π is envy-free if $v_i(X_i)+\pi_i \geq v_i(X_j)+\pi_j$ for every pair of agents $i, j \in N$. Informally, this extended version of envy-freeness requires that no agent envies the bundle and the payment of any other agent compared to the bundle and payment she gets.

We say that allocation X is *envy-freeable* if there is a payment vector π so that the pair (X, π) is envy-free. Although the use of payments makes envy-freeness a feasible goal, not all allocations are envy-freeable. The following theorem, due to Halpern and Shah [18], gives sufficient and necessary conditions so that an allocation is envy-freeable.

Theorem 1 (Halpern and Shah [18]). *The following statements are equivalent:*

- *The allocation $X = (X_1, X_2, ..., X_n)$ is envy-freeable.*
- *The allocation X maximizes social welfare among all redistributions of its bundles to the agents.*
- *The envy graph $EG(X, v)$ contains no directed cycles of positive total weight.*

Detecting whether a given allocation X is envy-freeable can be done using the following linear program $\mathrm{LP}(X, v)$:

$$\begin{aligned} \text{minimize} \quad & \sum_{i \in N} \pi_i \qquad (1) \\ \text{subject to:} \quad & \pi_i - \pi_j \geq v_i(X_j) - v_i(X_i), \forall i, j \in N \\ & \pi \geq 0 \end{aligned}$$

$\mathrm{LP}(X, v)$ aims to find a payment vector π so that the envy-freeness constraints between pairs of agents are satisfied. In addition, it minimizes the total amount of payments. As it is observed by Halpern and Shah [18], the payment π_i of agent i obtained in this way is equal to the maximum total weight in any simple path that originates from node i in the envy graph $\mathrm{EG}(X, v)$.

We study the optimization problem SMEF (standing for Subsidy Minimization for Envy-Freeness). Given an allocation instance, SMEF aims to compute an allocation that is envy-freeable with the minimum amount of subsidies. Since the problem of computing an envy-free allocation is NP-hard, SMEF is NP-hard as well.

We are interested in the design of approximation algorithms for SMEF. As algorithms with finite multiplicative approximation ratio are hopeless (since it is NP-hard to decide whether the minimum amount of subsidies is zero or not), we seek polynomial-time algorithms that compute an allocation that is envy-freeable with subsidies $\chi + \rho \cdot \text{sum } v$, with the approximation guarantee ρ being as low as possible.

As a warmup, consider the algorithm that allocates all items to the agent i^* who has maximum value for M and paying a subsidy of $v_{i^*}(M)$ to every other agent i. Clearly, this is a polynomial-time algorithm. The allocation obtained is envy-freeable since no redistribution of the bundles (i.e., giving all items to another agent) results in higher social welfare. And the particular payments are right: agent i^* is indifferent between the bundle M and the payment to any other agent, while the other agents are indifferent between the (equal) payments, and prefer their payment to getting the whole bundle M. It can be easily verified that the algorithm guarantees an amount of at most $\chi+(n-1)\max v \leq \chi+(n-1)\text{sum } v$ as subsidies; this is the best guarantee of this form for this algorithm in the worst-case.

3 An Approximation Algorithm

We now present an algorithm that does much better. The algorithm exploits ideas that have led to polynomial-time approximation schemes for combinatorial optimization problems like KNAPSACK; e.g., see [30]. It first discretizes all valuations to multiples of a discretization parameter. In this way, the different discretized valuations an agent can have for bundles of items in the new instance is small. This allows to classify all allocations into a relatively small number of classes, each defined by specific discretized valuation levels of each agent for all bundles. Dynamic programming is used to decide the classes that are non-empty and to select a representative allocation from each class. The final allocation is selected among all representative allocations, possibly after redistributing the bundles so that social welfare (with respect to the original valuations) is maximized (in order to get envy-freeability). This requires a call to linear program (1) to compute the minimum amount of subsidies for each representative allocation.

The classification of allocations guarantees that the algorithm will consider a representative allocation from the class that also contains the optimal one (i.e., the allocation that is envy-freeable with the minimum amount of subsidies overall). Our analysis shows that the amount of subsidies for making the representative allocation envy-free is close to optimal. Polynomial running time for the case of a constant number of agents follows by setting the discretization parameter appropriately.

We now present our algorithm in detail. It uses an accuracy parameter $\epsilon > 0$ and initially decides the value of the discretization parameter δ as follows:

$$\delta = \frac{\epsilon \max v}{4mn^2}.$$

First, the algorithm implicitly discretizes all agent valuations by defining new valuations $\tilde{v}$ as follows: for an agent i with valuation $v_i(g)$ for item g, the discretized valuation $\tilde{v}_i(g)$ is equal to $\lfloor v_i(g)/\delta \rfloor \delta$.

The algorithm uses an arbitrary ordering of the items in M; let $M = \{g_1, g_2, ..., g_m\}$, where the item indices are those in this ordering. The algorithm builds a table $\mathbf{T}$ which classifies all possible allocations of subsets of M. Consider an (n^2+1)-dimensional tuple $\tau = (t, P_{ij}, 1 \le i, j \le n)$, where t is an integer from 1 to m and P_{ij} is an integer from 0 to $\lfloor \max v/\delta \rfloor$, for every pair of agents i and j. The entry $\mathbf{T}(\tau)$ of the table indicates whether an allocation $A^t = (A^t_1, A^t_2, ..., A^t_n)$ of the first t items $g_1, ..., g_t$ of M to the n agents, satisfying $\tilde{v}_i(A^t_j) = P_{ij}\delta$ for every pair of agents i and j, exists ($\mathbf{T}(\tau) = 1$) or not ($\mathbf{T}(\tau) = 0$).

The entries of $\mathbf{T}$ are computed using the following recursive relation:

- For a tuple $\tau = (t, P_{ij}, 1 \le i, j \le n)$ with $t = 1$, the algorithm sets $\mathbf{T}(\tau) = 1$ if there exists $k \in [n]$ such that, for every $i \in [n]$, $\tilde{v}_i(g_1) = P_{ik}\delta$ and $P_{ij} = 0$ for every $j \neq k$. Otherwise, the algorithm sets $\mathbf{T}(\tau) = 0$.
- For a tuple $\tau = (t, P_{ij}, 1 \le i, j \le n)$ with $t > 1$, the algorithm sets $\mathbf{T}(\tau) = 1$ if there exists $k \in [n]$ and tuple $\tau' = (t-1, P'_{ij}, 1 \le i, j \le n)$ such that, for every $i \in [n]$, $P_{ik} = P'_{ik} + \tilde{v}_i(g_t)/\delta$ and $P_{ij} = P'_{ij}$ for every $j \neq k$. Otherwise, the algorithm sets $\mathbf{T}(\tau) = 0$.

Essentially, each non-zero entry of $\mathbf{T}$ (i.e., $\mathbf{T}(\tau) = 1$) indicates a non-empty class $\mathcal{A}_\tau$ of (possibly partial, when the first argument of τ is an integer smaller than m) allocations. To compute a representative complete allocation $A_\tau \in \mathcal{A}_\tau$ among those implied by the non-zero entry corresponding to the tuple $(m, P^m_{ij}, 1 \le i, j \le n)$, the algorithm does the following for $t = m$ downto 2. Let $k \in [n]$ be such that $\mathbf{T}(\tau') = 1$ for a tuple $\tau' = (t-1, P^{t-1}_{ij}, 1 \le i, j \le n)$ with $P^{t-1}_{ik} = P^t_{ik} - \tilde{v}_i(g_t)/\delta$ and $P^{t-1}_{ij} = P^t_{ij}$ for every pair of agents i and $j \neq k$. The algorithm assigns item g_t to agent k and proceeds to considering the next item. The first item g_1 is assigned to agent k such that $\mathbf{T}(\tau') = 1$ for a tuple $\tau' = (1, P^1_{ij}, 1 \le i, j \le n)$ with $P^1_{ik} = \tilde{v}_i(g_1)/\delta$ and $P^1_{ij} = 0$ for every pair of agents i and $j \neq k$.

Next, the algorithm redistributes the bundles of each allocation A_τ that represents a non-empty class $\mathcal{A}_\tau$ so that an allocation A'_τ of maximum social welfare (among those that distribute the particular bundles to the agents) is obtained (in terms of the original valuations). It solves $\text{LP}(A'_\tau, v)$ (for the original valuations) to compute the minimum amount of subsidies that make A'_τ envy-free. Among all allocations A'_τ, it outputs the one with the minimum amount of subsidies. The approximation guarantee of the algorithm is given by the next lemma.

Lemma 1. *Given an instance of SMEF that has an allocation that is envy-freeable with an amount of χ as total subsidies, the algorithm computes an allocation that is envy-freeable with total subsidies of at most $\chi + 4mn^2\delta$.*

Proof. Let τ be a full tuple such that $\mathcal{A}_\tau$ contains an allocation $O = (O_1, ..., O_n)$ that is envy-freeable with subsidies of χ. Since $\mathcal{A}_\tau$ is non-empty, it is $\mathbf{T}(\tau) = 1$. Let A be the allocation computed by the algorithm as representative of $\mathcal{A}_\tau$ and A' the allocation that is obtained after redistributing the bundles of A. By Theorem 1, A' is clearly envy-freeable; we will show that the corresponding subsidies are at most $\chi + 4mn^2\delta$. Clearly, the output of the algorithm will be envy-freeable with at most this amount of subsidies.

Let $\sigma \in \mathcal{L}(n)$ be the permutation over $[n]$ such that $A'_j = A_{\sigma(j)}$ for every $j \in [n]$. Let G and H be the envy graphs $\mathrm{EG}(O, v)$ and $\mathrm{EG}(A', v)$, respectively.

We now present the most crucial component of our analysis. It exploits the fact that both O and A belong to class $\mathcal{A}_\tau$ and uses the third statement of Theorem 1.

Lemma 2. *For every pair of agents i and j, there exists a (not necessarily simple) path $p(i,j)$ from node $\sigma(i)$ to node $\sigma(j)$ such that*

$$wgt_H(i,j) \leq \sum_{e \in p(i,j)} wgt_G(e) + 4m\delta.$$

Proof. In the proof, we will use the following simple lemma.

Lemma 3. *For every agent i and every two bundles B_1 and B_2 such that $\tilde{v}_i(B_1) = \tilde{v}_i(B_2)$, it holds that*

$$-|B_2|\delta \leq v_i(B_1) - v_i(B_2) \leq |B_1|\delta. \tag{2}$$

Proof. First observe that, by the definition of $\tilde{v}$ and its relation to v, for every agent i and item $g \in M$, it holds that $\tilde{v}_i(g) \leq v_i(g) \leq \tilde{v}_i(g) + \delta$. Hence, for every bundle B,

$$\tilde{v}_i(B) \leq v_i(B) \leq \tilde{v}_i(B) + |B|\delta.$$

The lemma follows by applying this inequality for bundles B_1 and B_2 and using the fact that $\tilde{v}_i(B_1) = \tilde{v}_i(B_2)$. □

We use the notation σ^{-1} to refer to the inverse permutation of σ, i.e., $\sigma^{-1}(k) = j$ when $k = \sigma(j)$. Consider the set C that contains edge $(k, \sigma^{-1}(k))$ for every agent k such that $k \neq \sigma^{-1}(k)$. C is either empty (if $k = \sigma^{-1}(k)$ for every agent k) or consists of disjoint directed cycles. For an agent i, if $\sigma^{-1}(i) \neq i$, we denote by C_i the set of nodes that are spanned by the cycle of C that includes node i. Otherwise, we define C_i to contain only node i.

Define the (not necessarily simple) path $p(i,j)$ from node $\sigma(i)$ to node $\sigma(j)$ to contain edge $(k, \sigma(k))$ for every node k in the set C_i besides node i and, if $i \neq \sigma(j)$, the directed edge $(i, \sigma(j))$.

For every pair of agents i and j, we have that the weight of the directed edge (i,j) in H is

$$\begin{aligned}
\mathrm{wgt}_H(i,j) &\leq \mathrm{wgt}_H(i,j) - \sum_{k \in C_i} \mathrm{wgt}_H(k, \sigma^{-1}(k)) \\
&= v_i(A'_j) - v_i(A'_i) - \sum_{k \in C_i} \left(v_k(A'_{\sigma^{-1}(k)}) - v_k(A'_k) \right) \\
&= v_i(A_{\sigma(j)}) - v_i(A_{\sigma(i)}) - \sum_{k \in C_i} \left(v_k(A_k) - v_k(A_{\sigma(k)}) \right)
\end{aligned}$$

$$\begin{aligned}
&\leq v_i(O_{\sigma(j)}) - v_i(O_{\sigma(i)}) - \sum_{k\in C_i} \big(v_k(O_k) - v_k(O_{\sigma(k)})\big)\\
&\quad + \left(|A_{\sigma(j)}| + |O_{\sigma(i)}| + \sum_{k\in C_i} |O_k| + \sum_{k\in C_i} |A_{\sigma(k)}|\right)\delta\\
&\leq v_i(O_{\sigma(j)}) - v_i(O_{\sigma(i)}) - \sum_{k\in C_i} \big(v_k(O_k) - v_k(O_{\sigma(k)})\big) + 4m\delta\\
&= v_i(O_{\sigma(j)}) - v_i(O_i) + \sum_{k\in C_i\setminus\{i\}} \big(v_k(O_{\sigma(k)}) - v_k(O_k)\big) + 4m\delta\\
&= \mathrm{wgt}_G(i,\sigma(j)) + \sum_{k\in C_i\setminus\{i\}} \mathrm{wgt}_G(k,\sigma(k)) + 4m\delta\\
&= \sum_{e\in p(i,j)} \mathrm{wgt}_G(e) + 4m\delta.
\end{aligned}$$

The first inequality follows since C_i consists of node i only (when $i = \sigma(i)$) or the edges $(k, \sigma^{-1}(k))$ for $k \in C_i$ form a directed cycle of non-positive total weight in H. The second inequality follows by applying Lemma 3 (recall that both allocations A and O belong to the class $\mathcal{A}_\tau$ and, hence, $\tilde{v}_\ell(A_q) = \tilde{v}_\ell(O_q)$ for every pair of agents ℓ and q). The third inequality follows since the bundles $A_{\sigma(k)}$ (respectively, O_k) for $k \in C_i$ are disjoint. The equalities are obvious or follow by the definition of the weights. □

Now, let π' and π be the solutions of LP(A', v) and LP(O, v), respectively. Hence, $\chi = \mathrm{Sub}(O, v) = \sum_{i=1}^{n} \pi_i$. We will use Lemma 2 to argue that

$$\pi'_i \leq \pi_{\sigma(i)} + 4mn\delta. \tag{3}$$

This will yield total subsidies of

$$\mathrm{Sub}(A', v) = \sum_{i=1}^{n} \pi'_i \leq \sum_{i=1}^{n} \big(\pi_{\sigma(i)} + 4mn\delta\big) = \chi + 4mn^2\delta,$$

completing the proof.

Recall from Theorem 1 that the payment π'_ℓ (respectively, π_ℓ) is equal to the maximum path weight over all simple paths that originate from node ℓ in graph H (respectively, graph G). Let Q_ℓ be the corresponding simple path that is destined for some node s (and originates from node ℓ), i.e., $\pi'_\ell = \sum_{e\in Q_\ell} \mathrm{wgt}_H(e)$. We construct the (not necessarily simple) path P_ℓ from node $\sigma(\ell)$ to node $\sigma(s)$ of G that consists of path $p(i, j)$ for every directed edge (i, j) in the path Q_ℓ. Using Lemma 2, we get

$$\begin{aligned}
\pi'_\ell = \sum_{e\in Q_\ell} \mathrm{wgt}_H(e) &\leq \sum_{e\in Q_\ell} \left(\sum_{e'\in p(e)} \mathrm{wgt}_G(e') + 4m\delta\right) \\
&\leq \sum_{e\in Q_\ell} \sum_{e'\in p(e)} \mathrm{wgt}_G(e') + 4mn\delta = \sum_{e\in P_\ell} \mathrm{wgt}_G(e) + 4mn\delta.
\end{aligned} \tag{4}$$

The second inequality follows since path Q_ℓ is simple (and, hence, contains at most $n-1$ edges). Now, create the simple path P'_ℓ from node $\sigma(\ell)$ to node $\sigma(s)$ by removing the cycles in P_ℓ. Since graph G does not have any directed cycles of positive total weight (by Theorem 1), we have $\text{wgt}_G(P_\ell) \leq \text{wgt}_G(P'_\ell)$. Now, (4) yields

$$\pi'_\ell \leq \sum_{e \in P'_\ell} \text{wgt}_G(e) + 4mn\delta,$$

which implies (3) since P'_ℓ is a simple path that originates from node $\sigma(\ell)$. □

The running time of the algorithm depends on the number of table entries, the number of steps required for computing each table entry using the recursive relation, the number of steps required to compute a representative allocation for a non-empty allocation class, the redistribution time, and the time required to solve the linear programs.

The dimensions of the table $\mathbf{T}$ are m for the first one that enumerates over all items, and at most $1 + \lfloor \max v/\delta \rfloor = 1 + \frac{4mn^2}{\epsilon}$ for each of the other dimensions. Overall, the size of the table is $\mathcal{O}\left(\left(\frac{m}{\epsilon}\right)^{n^2+1}\right)$. The computation of each table entry using the recursive relation needs the values in n^2 table entries that have been previously computed. In a representative allocation, the agent to which each of the m items is allocated requires time n^2 as well, i.e., time $\mathcal{O}(m)$ in total. The redistribution of the bundles can be implemented using a matching computation in a complete edge-weighted bipartite graph that has a node for each agent and for each bundle and the weight of an edge indicates the valuation of an agent for a bundle. As n is constant, this takes constant time. Also, the linear programs have constant size. In general, since n is a constant, it is ignored in the $\mathcal{O}$ notation unless it appears in the exponent. The above discussion is summarized in the next statement.

Theorem 2. *Let $\epsilon > 0$ be the accuracy parameter used by the algorithm. Given an instance of SMEF consisting of a constant number n of agents with valuations v over a set M of m items that has an envy-freeable allocation using an amount χ of subsidies, the algorithm runs in time $\mathcal{O}\left((m/\epsilon)^{n^2+2}\right)$ and computes an allocation that is envy-freeable using a total subsidy of at most $\chi + \epsilon \max v$.*

4 Hardness of Approximating SMEF

In this section, we show that approximation guarantees like the one in the statement of Theorem 2 are not possible when the number of agents is part of the input.

Theorem 3. *Approximating SMEF within an additive term of $3 \cdot 10^{-4}$sum v is NP-hard.*

We prove Theorem 3 by presenting a reduction from Maximum 3-Dimensional Matching (MAX-3DM). An instance of MAX-3DM consists of three disjoint sets of elements $A = \{a_1, a_2, ..., a_n\}$, $B = \{b_1, b_2, ..., b_n\}$, and $C = \{c_1, c_2, ..., c_n\}$, each of size n, and a set T of m triplets of the form (a_i, b_j, c_k) with $a_i \in A$, $b_j \in B$, and $c_k \in C$. The objective is to compute a disjoint subset of T (or, simply, a 3D matching) of maximum size. The problem is well-known to be NP-hard not only to solve exactly [16] but also to approximate [19].

We will use the inapproximability result of Chlebík and Chlebíková [14], which applies to bounded instances of MAX-3DM in which each element appears in exactly two triplets (i.e., $m = 2n$); we will refer to this restriction of MAX-3DM as MAX-3DM-2. In particular, Chlebík and Chlebíková [14] show that it is NP-hard to distinguish between instances of MAX-3DM-2 with a 3D matching of size at least K and instances of MAX-3DM-2 in which any 3D matching has size at most $K - 0.01n$.[3]

4.1 The Reduction

We present our reduction and full proof for the case $\chi > 0$. We omit the case $\chi = 0$, which requires a minor modification of the reduction. On input an instance of MAX-3DM-2, our reduction constructs in polynomial time an instance of SMEF, in which the minimum amount of subsidies that can make some allocation envy-free is exactly $\chi(1 + \max\{K - L, 0\})$, where L is the size of the maximum 3D matching in the MAX-3DM-2 instance. Using the result of [14], we will get that it is NP-hard to distinguish between SMEF instances in which the minimum amount of subsidies is at most χ and instances in which it is at least $\chi(1 + 0.01n)$. Hence, SMEF will be proved to be NP-hard to approximate within $0.01n\chi$. Our construction will be such that sum $v < 30n\chi$. In this way, we will obtain a hardness of approximating SMEF within an additive term of (at least) $3 \cdot 10^{-4}$sum v, as desired.

Our reduction is as follows. Given an instance of MAX-3DM-2 consisting of sets of elements A, B, and C, each of size n, and a set of $2n$ triplets T, the instance of SMEF has

- three agents 1, 2, and 3,
- three agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ for every triplet $t \in T$,
- an item A_i for every element $a_i \in A$,
- an item B_i for every element $b_i \in B$,
- an item Γ_i for every element $c_i \in C$,
- three items Δ_t, Z_t, and Θ_t for every triplet $t \in T$, and
- an additional item Λ.

The agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ that correspond to the triplet $t = (a_i, b_j, c_k)$ have valuations 0 for all items besides the items A_i, B_j, Γ_k, Δ_t, Z_t, and Θ_t.

[3] This statement is actually weaker than the one proved in [14]. However, it suffices for our purpose to prove hardness of approximation. Note that we have made no particular attempt to optimize our inapproximability threshold.

Agents $1, 2$ have valuation 0 for all items besides item Λ and agent 3 has valuation zero for all items besides item Λ and items Θ_t for $t \in T$. Their remaining valuations are as follows:

	A_i	B_j	Γ_k	Δ_t	Z_t	Θ_t	Λ
1	0	0	0	0	0	0	χ
2	0	0	0	0	0	0	χK
3	0	0	0	0	0	χ	χK
$J_1(t)$	χ	χ	χ	3χ	3χ	0	0
$J_2(t)$	0	0	0	χ	χ	χ	0
$J_3(t)$	0	0	0	0	χ	0	0

Recall that each element belongs to exactly two triplets. Hence, two agents have positive value for item A_i (similarly for items B_j and Γ_k): agents $J_1(t_1)$ and $J_1(t_2)$ such that the triplets t_1 and t_2 contain element a_i (similarly for elements b_j and c_k). It is easy to see that either two or three agents have positive value for each item. For every triplet t, the agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ have total valuation 9χ, 3χ, and χ, respectively. Taking into account that $K \leq n$, we obtain that sum $v < 30n\chi$.

4.2 Lower Bound on Subsidies

Consider an instance of SMEF constructed by our reduction and let X be an envy-freeable allocation in it. We will first lower-bound the minimum amount of subsidies that make X envy-free. First observe that X cannot give item Λ to agent 1; in that case, exchanging the bundles of agents 1 and 2 would result to an increase of the social welfare and, hence, X would not be envy-freeable. If X gives item Λ to agent 3, agents 1 and 2 would need subsidies of at least χ and χK, respectively, so that they do not envy agent 3. Hence, $\mathrm{Sub}(A, v) \geq \chi(1+K)$ in this case.

In the following, we will lower-bound the minimum total subsidies that make X envy-free assuming that item Λ is given to agent 2. Let θ be the number of items Θ_t for $t \in [2n]$ agent 3 gets. Then, agent 3 should be given a subsidy of at least $\chi \max\{K - \theta, 0\}$ so that she does not envy agent 2. Agent 1 needs a subsidy of $\chi \max\{K - \theta, 1\}$ so that she does not envy agents 1 and 2.

For a triplet $t = (a_i, b_j, c_k)$ in the original instance of MAX-3DM-2, we call it *full* if all items A_i, B_j, and Γ_k (which correspond to the elements of the triplet) have been allocated to the agents $J_1(t)$, $J_2(t)$, or $J_3(t)$. Otherwise, we call it *partial*. We call t *supported* if item Θ_t has been allocated to agent $J_2(t)$; otherwise, we call t *unsupported*.

In the next four lemmas, we lower-bound the total amount of subsidies the agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ of a triplet t need, depending on the type of t.

Lemma 4. *The agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ of a full and supported triplet t need subsidies of at least $\chi \max\{K - \theta - 2, 0\}$.*

Proof. Consider a full and supported triplet t. If agent $J_2(t)$ has value at most 2χ (i.e., getting Θ_t and at most one of the items Δ_t and Z_t), then she needs a subsidy of at least $\chi \max\{K - \theta - 2, 0\}$ so that she does not envy agent 3. If agent $J_2(t)$ has value 3χ by getting both items Δ_t and Z_t in addition to Θ_t, she needs a subsidy of at least $\chi \max\{K-\theta-3, 0\}$, while then agents $J_1(t)$ and $J_3(t)$ need subsidies of at least $3\chi + \chi \max\{K - \theta - 3, 0\}$ and $\chi + \chi \max\{K - \theta - 3, 0\}$, respectively, so that they do not envy agent $J_2(t)$. In both cases, the total amount of subsidies of the agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ is at least $\chi \max\{K - \theta - 2, 0\}$. □

Lemma 5. *The agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ of a full and unsupported triplet t need subsidies of at least $\chi \max\{K - \theta - 1, 0\}$.*

Proof. Consider a full and unsupported triplet t. If agent $J_2(t)$ has value at most χ (i.e., getting at most one of the items Δ_t and Z_t), then she needs a subsidy of at least $\chi \max\{K - \theta - 1, 0\}$ so that she does not envy agent 3. If agent $J_2(t)$ has value 2χ by getting both items Δ_t and Z_t, she needs a subsidy of at least $\chi \max\{K - \theta - 2, 0\}$, while then agents $J_1(t)$ and $J_3(t)$ need subsidies of at least $3\chi + \chi \max\{K - \theta - 2, 0\}$ and $\chi + \chi \max\{K - \theta - 2, 0\}$, respectively, so that they do not envy agent $J_2(t)$. In both cases, the total amount of subsidies of agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ is at least $\chi \max\{K - \theta - 1, 0\}$. □

Lemma 6. *The agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ of a partial and supported triplet t need subsidies of at least $\chi \max\{K - \theta - 1, 0\}$.*

Proof. Let t be a partial and supported triplet. If agent $J_2(t)$ does not get items Δ_t and Z_t, then she gets only a value of χ from item Θ_t and needs a subsidy of at least $\chi \max\{K - \theta - 1, 0\}$ so that she does not envy agent 3.

If agent $J_2(t)$ gets item Δ_t but not item Z_t, she needs a subsidy of $\chi \max\{K - \theta - 2, 0\}$ so that she does not envy agent 3. Then, if agent $J_1(t)$ does not get item Z_t, her value is at most 2χ (from at most two of the items A_i, B_j, and Γ_k) and needs a subsidy of $\chi + \chi \max\{K - \theta - 2, 0\}$ so that she does not envy agent $J_2(t)$. If agent $J_3(t)$ does not get item Z_t, she needs a subsidy of at least $\chi + \chi \max\{K - \theta - 2, 0\}$ so that she does not envy agent $J_2(t)$.

If agent $J_2(t)$ gets item Z_t but not Δ_t, she needs a subsidy of $\chi \max\{K - \theta - 2, 0\}$ so that she does not envy agent 3 and agent $J_3(t)$ needs a subsidy of at least $\chi + \chi \max\{K - \theta - 2, 0\}$ so that she does not envy agent $J_2(t)$.

Finally, if agent $J_2(t)$ gets items Δ_t and Z_t, her value is 3χ and needs a subsidy of at least $\chi \max\{K - \theta - 3, 0\}$ so that she does not envy agent 3. Then, each of agents $J_1(t)$ and $J_3(t)$ need a subsidy of at least $\chi + \chi \max\{K - \theta - 3, 0\}$ so that they do not envy agent $J_2(t)$.

In all cases, the total amount of subsidies the agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ need is at least $\chi \max\{K - \theta - 1, 0\}$. □

Lemma 7. *The agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ of a partial and unsupported triplet t need subsidies of at least $\chi \max\{K - \theta, 1\}$.*

Proof. Let t be a partial and unsupported triplet. If agent $J_2(t)$ gets both items $\varDelta_t$ and Z_t, she needs a subsidy of $\chi \max\{K-\theta-2,0\}$ so that she does not envy agent 3, while agents $J_1(t)$ and $J_3(t)$ would then need subsidies of at least $4\chi+\chi \max\{K-\theta-2,0\}$ and $\chi+\chi \max\{K-\theta-2,0\}$, respectively, so that they do not envy agent $J_2(t)$.

If agent $J_2(t)$ gets only item $\varDelta_t$, she needs a subsidy of $\chi \max\{K-\theta-1,0\}$ so that she does not envy agent 3. Then, the agent who does not get item Z_t among $J_1(t)$ and $J_3(t)$ would need a subsidy of at least $\chi+\chi \max\{K-\theta-1,0\}$ so that she does not envy agent $J_2(t)$.

If agent $J_2(t)$ gets only item Z_t, she needs a subsidy of $\chi \max\{K-\theta-1,0\}$ so that she does not envy agent 3, while agent $J_3(t)$ needs a subsidy of at least $\chi+\chi \max\{K-\theta-1,0\}$ so that she does not envy agent $J_2(t)$.

Finally, if agent $J_2(t)$ gets no item (among $\varDelta_t$ and Z_t), she needs a subsidy of at least χ so that she does not envy the agents who get items $\varDelta_t$ and Z_t and a subsidy of at least $\chi \max\{K-\theta,0\}$ so that she does not envy agent 3.

In all cases, the total amount of subsidies the agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ need is at least $\chi \max\{K-\theta,1\}$. □

We now denote by L_1, L_2, P_1, and P_2, the number of full and supported, full and unsupported, partial and supported, and partial and unsupported triplets defined by X, respectively. Notice that the full triplets form a 3D matching. Denoting by L the maximum size over all 3D matchings of the MAX-3DM-2 instance, we have $L \geq L_1+L_2$. Using Lemmas 4–7, and our observations for agents 1 and 3, we have that the total amount of subsidies X needs to become envy-free is

$$\begin{aligned}\mathrm{Sub}(X,v) \geq \chi\,(&L_1 \max\{K-\theta-2,0\}+L_2 \max\{K-\theta-1,0\}\\ &+P_1 \max\{K-\theta-1,0\}+P_2 \max\{K-\theta,1\}\\ &+\max\{K-\theta,0\}+\max\{K-\theta,1\})\,.\end{aligned} \tag{5}$$

We will distinguish between two cases for $K-\theta$. If $K-\theta \geq 2$, (5) yields

$$\begin{aligned}\mathrm{Sub}(X,v) &\geq \chi\,(L_2+P_1+2P_2+4)=\chi\,(2n-L_1+P_2+4)\\ &\geq \chi(1+\max\{K-L,0\}).\end{aligned}$$

Now, notice that θ, the number of items $\varTheta_t$ agent 3 gets in X is upper-bounded by the number of unsupported triplets, i.e., $\theta \leq L_2+P_2$. Thus, if $K-\theta \leq 1$, (5) yields

$$\mathrm{Sub}(X,v) \geq \chi\,(P_2+K-\theta+1) \geq \chi\,(K-L_2+1) \geq \chi(1+\max\{K-L,0\}).$$

We conclude that the minimum amount of subsidies necessary to make X envy-free is at least $\chi(1+\max\{K-L,0\})$.

4.3 Upper Bound on Minimum Subsidies

We now present our upper bound on the minimum amount of subsidies for envy-freeness. Given a 3D matching $\mathcal{M}$ of maximum size L in the MAX-3DM-2 instance, we will construct an allocation for the SMEF instance and will show that it is envy-freeable with an amount of subsidies equal to $\chi(1+\max\{K-L,0\})$.

For defining the allocation, we partition $T\setminus\mathcal{M}$ in two disjoint sets of triplets T_1 and T_2 of size $2n-\max\{K,L\}$ and $\max\{K-L,0\}$, respectively.

- For every triplet $t=(a_i,b_j,c_k)\in\mathcal{M}$, agent $J_1(t)$ gets items A_i, B_j, and Γ_k, agent $J_2(t)$ gets item Δ_t and agent $J_3(t)$ gets item Z_t.
- For every triplet $t=(a_i,b_j,c_k)\notin\mathcal{M}$, let $F(t)$ be the set of items that correspond to the elements of t that have not been included in triplets of $\mathcal{M}$. Note that, due to the maximality of $\mathcal{M}$, $F(t)$ has zero, one, or two elements among A_i, B_j, and Γ_k. For every triplet $t=(a_i,b_j,c_k)\in T_1$, agent $J_1(t)$ gets item Δ_t, agent $J_2(t)$ gets the items in $F(t)$, if any, and item Θ_t, and agent $J_3(t)$ gets item Z_t.
- For every triplet $t=(a_i,b_j,c_k)\in T_2$, agent $J_1(t)$ gets item Δ_t, agent $J_2(t)$ gets the items in $F(t)$, if any, and agent $J_3(t)$ gets item Z_t.
- Agent 3 gets item Θ_t for every triplet $t\in\mathcal{M}\cup T_2$.
- Agent 2 gets item Λ.
- Agent 1 gets no items.

We claim that the allocation above is envy-freeable by assigning a subsidy of χ to agent 1 and a subsidy of χ to agent $J_2(t)$ for every triplet $t\in T_2$ (if any).

Indeed, agent 1 has positive value only for item Λ, which is given to agent 2, who gets no subsidy. Also, no other agent gets a subsidy more than the subsidy χ that is given to agent 1. Hence, agent 1 is not envious. Agent 2 gets item Λ, which is the only item she values positively and much higher than the subsidy given to any other agent. Hence, agent 2 is not envious either. Agent 3 gets exactly $\max\{K,L\}$ items of total value $\chi\max\{K,L\}$. She does not envy agent 2 who gets item Λ (which agent 3 values for χK) since no subsidy is given to agent 2. Clearly, the value of agent 3 is much higher than the subsidy given to any other agent.

Consider a triplet $t=(a_i,b_j,c_k)\in\mathcal{M}$. Agent $J_1(t)$ has a value of 3χ for the items A_i, B_j, and Γ_k she gets. The remaining items for which she has positive valuation of 3χ have been given to agents $J_2(t)$ and $J_3(t)$, respectively. Since these agents do not get subsidies, agent $J_1(t)$ is not envious of them. Clearly, agent $J_1(t)$ is not envious of any other agent since she has zero value for all other items and no agent gets a subsidy more than χ. Agent $J_3(t)$ gets item Z_t, the only item for which she has positive value and does not envy any other agent since no one gets a subsidy higher than χ. Agent $J_2(t)$ gets a value of χ from item Δ_t and does not envy agent $J_3(t)$, who gets item Z_t, or agent 3, who gets item Θ_t, as these agents receive no subsidy. Clearly, agent $J_2(t)$ envies no other agent.

Now consider a triplet $t=(a_i,b_j,c_k)\notin\mathcal{M}$. Agent $J_1(t)$ has a value of 3χ for the item Δ_t she gets. The remaining items for which she has positive valuation have been allocated as follows. Item Z_t has been given to agent $J_3(t)$; clearly, agent $J_1(t)$ is not envious of $J_3(t)$ since the latter gets no subsidies. The items in $F(t)$ have been given to agent $J_2(t)$. Again, agent $J_1(t)$ is not envious of $J_2(t)$

since $F(t)$ contains at most two items (which agent $J_1(t)$ values for χ each) and agent $J_2(t)$ gets a subsidy of zero (if $t \in T_1$) or χ (if $i \in T_2$). Clearly, $J_1(t)$ does not envy any other agent. Agent $J_3(t)$ gets item Z_t, the only item for which she has positive value and does not envy any other agent since no one gets a subsidy higher than χ. Agent $J_2(t)$ gets a value of χ, either from item Θ_t (if $t \in T_1$) or as subsidy (if $t \in T_2$), and does not envy agent $J_1(t)$ who gets item Δ_t or agent 3 who gets item Θ_t only when $t \in T_2$; recall that these two agents never get subsidies. Again, agent $J_2(t)$ envies no other agent.

4.4 Adapting the Proof for the Case $\chi = 0$

The modification required in our reduction so that it covers the case $\chi = 0$ as well is to remove agent 1 and replace χ with 1 in the definition of valuations. In particular, the agents $J_1(t)$, $J_2(t)$, and $J_3(t)$ that correspond to the triplet $t = (a_i, b_j, c_k)$ have valuations 0 for all items besides the items A_i, B_j, Γ_k, Δ_t, Z_t, and Θ_t. Agent 2 has valuation 0 for all items besides item Λ and agent 3 has valuation zero for all items besides item Λ and items Θ_t for $t \in T$. The remaining valuations are now as follows:

	A_i	B_j	Γ_k	Δ_t	Z_t	Θ_t	Λ
2	0	0	0	0	0	0	K
3	0	0	0	0	0	1	K
$J_1(t)$	1	1	1	3	3	0	0
$J_2(t)$	0	0	0	1	1	1	0
$J_3(t)$	0	0	0	0	1	0	0

The same reasoning as in our proof for the case $\chi \neq 0$ gives a minimum amount of subsidies for the SMEF instance of exactly $\max\{K - L, 0\}$, where L is the maximum 3D matching size in the MAX-3DM-2 instance. In this way, we get that SMEF is NP-hard to approximate within $0.01n$ (i.e., it is NP-hard to distinguish between envy-free instances and instances that need subsidies of $0.01n$) and the construction satisfies sum $v < 30n$. This yields the desired inapproximability result in the statement of Theorem 3 for the case $\chi = 0$ as well.

5 Concluding Remarks

We have initiated the study of the optimization problem SMEF. The challenging open problem that deserves investigation is to close the gap between the trivial approximation guarantee of $n - 1$ in Sect. 2 and our negative result for super-constant numbers of agents in Sect. 4. Unfortunately, more sophisticated existing algorithms, such as the recent one by Brustle et al. [7], do not lead to better approximations.

We remark that $\max v$ could be used alternatively to sum in the definition of the approximation guarantees of SMEF. Actually, the guarantee for our dynamic programming algorithm is stated in terms of $\max v$. We can express the rest of our results using $\max v$ as well. First, the trivial algorithm presented at the end of Sect. 2 uses an amount of $\chi + (n-1)\max v$ as subsidies. Second, an adaptation of the current proof of the inapproximability result can easily give that approximating SMEF within an additive term of $c \cdot \max v$ for a constant c is NP-hard. The important observation is that $\max v < n\chi$ (or $\max v < n$ when $\chi = 0$) in our construction. Then, distinguishing between SMEF instances in which the minimum amount is at most χ and at least $\chi(1 + 0.01n)$ (or at least $0.01n$ when $\chi = 0$) requires to distinguish between SMEF instances in which the minimum amount is at most χ and at least $\chi + 0.01 \max v$. So, the inapproximability constant is a bit higher in this case. The main advantage of adopting sum is that it makes the problem of computing the tight approximation factor more challenging.

Interestingly, an advantage of the trivial algorithm is that the particular payments incentivize the agents to report their valuations truthfully. What is the best possible approximation guarantee that can be obtained for SMEF by truthful algorithms? Unfortunately, a simple application of Myerson's characterization in single-item settings [24] indicates that no approximation guarantee better than $n-1$ is possible. Indeed, consider instances with a single item. By the characterization of envy-freeable allocations by Halpern and Shah [18] (i.e., the second statement in Theorem 1), we know that the agent with the highest valuation should get the item. Then, Myerson's characterization for truthful mechanisms in single parameters environments and our requirement for non-negative payments give us the specific form payments should have so that truthful reporting is a dominant strategy for all agents when this algorithm is used: if the agent i who gets the item receives payment of $p \geq 0$, agent t should get a payment of exactly $p + v_i - v_t$, where v_i and v_t are the payments of agents i and t. Now, consider specifically the instance in which one agent has value 1 for the item, and all other agents have value 0. Truthfulness requires (at least) a unit of subsidy to each agent that does not get the item (i.e., total subsidies of $n-1$ while sum $= 1$), even though there is clearly an allocation that is envy-free without any payments. This yields the claimed lower bound of $n-1$ in the approximation guarantee.

References

1. Alkan, A., Demange, G., Gale, D.: Fair allocation of indivisible goods and criteria of justice. Econometrica **59**(4), 1023–1039 (1991)
2. Amanatidis, G., Markakis, E., Ntokos, A.: Multiple birds with one stone: beating 1/2 for EFX and GMMS via envy cycle elimination. Theoret. Comput. Sci. **841**, 94–109 (2020)
3. Aragones, E.: A derivation of the money Rawlsian solution. Soc. Choice Welf. **12**(3), 267–276 (1995). https://doi.org/10.1007/BF00179981
4. Aziz, H., Bouveret, S., Caragiannis, I., Giagkousi, I., Lang, J.: Knowledge, fairness, and social constraints. In: Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI), pp. 4638–4645 (2018)

5. Aziz, H., Caragiannis, I., Igarashi, A., Walsh, T.: Fair allocation of indivisible goods and chores. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), pp. 53–59 (2019)
6. Barman, S., Krishnamurthy, S.K., Vaish, R.: Finding fair and efficient allocations. In: Proceedings of the 19th ACM Conference on Economics and Computation (EC), pp. 557–574 (2018)
7. Brustle, J., Dippel, J., Narayan, V.V., Suzuki, M., Vetta, A.: One dollar each eliminates envy. In: Proceedings of the 21st ACM Conference on Economics and Computation (EC), pp. 23–39 (2020)
8. Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. J. Polit. Econ. **119**(6), 1061–1103 (2011)
9. Caragiannis, I., Gravin, N., Huang, X.: Envy-freeness up to any item with high Nash welfare: the virtue of donating items. In: Proceedings of the 20th ACM Conference on Economics and Computation (EC), pp. 527–545 (2019)
10. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum nash welfare. ACM Trans. Econ. Comput. **7**(3), 12:1–12:32 (2019)
11. Chaudhury, B.R., Kavitha, T., Mehlhorn, K., Sgouritsa, A.: A little charity guarantees almost envy-freeness. SIAM J. Comput. **50**(4), 1336–1358 (2021)
12. Chevaleyre, Y., Endriss, U., Estivie, S., Maudet, N.: Reaching envy-free states in distributed negotiation settings. In: Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI), pp. 1239–1244 (2007)
13. Chevaleyre, Y., Endriss, U., Maudet, N.: Distributed fair allocation of indivisible goods. Artif. Intell. **242**, 1–22 (2017)
14. Chlebík, M., Chlebíková, J.: Complexity of approximating bounded variants of optimization problems. Theoret. Comput. Sci. **354**(3), 320–338 (2006)
15. Foley, D.K.: Resource allocation and the public sector. Yale Econ. Essays **7**(1), 45–98 (1967)
16. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman & Co., New York (1979)
17. Haake, C.J., Raith, M.G., Su, F.E.: Bidding for envy-freeness: a procedural approach to n-player fair-division problems. Soc. Choice Welf. **19**(4), 723–749 (2002)
18. Halpern, D., Shah, N.: Fair division with subsidy. In: Fotakis, D., Markakis, E. (eds.) SAGT 2019. LNCS, vol. 11801, pp. 374–389. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30473-7_25
19. Kann, V.: Maximum bounded 3-dimensional matching is MAXSNP-complete. Inf. Process. Lett. **37**(1), 27–35 (1991)
20. Klijn, F.: An algorithm for envy-free allocations in an economy with indivisible objects and money. Soc. Choice Welf. **17**(2), 201–215 (2000)
21. Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: Proceedings of the 5th ACM Conference on Electronic Commerce (EC), pp. 125–131 (2004)
22. Maskin, E.: On the fair allocation of indivisible goods. In: Arrow and the Foundations of the Theory of Economic Policy (Essays in Honor of Kenneth Arrow), pp. 341–349. MacMillan (1987). https://doi.org/10.1007/978-1-349-07357-3_12
23. Meertens, M., Potters, J., Reijnierse, H.: Envy-free and pareto efficient allocations in economies with indivisible goods and money. Math. Soc. Sci. **44**(3), 223–233 (2002)
24. Myerson, R.B.: Optimal auction design. Math. Oper. Res. **6**(1), 58–73 (1981)
25. Plaut, B., Roughgarden, T.: Almost envy-freeness with general valuations. SIAM J. Discret. Math. **34**(2), 1039–1068 (2020)

26. Procaccia, A.D.: An answer to fair division's most enigmatic question: technical perspective. Commun. ACM **63**(4), 118 (2020)
27. Su, F.E.: Rental harmony: Sperner's lemma in fair division. Am. Math. Mon. **106**(10), 930–942 (1999)
28. Svensson, L.G.: Large indivisibles: an analysis with respect to price equilibrium and fairness. Econometrica **51**(4), 939–954 (1983)
29. Varian, H.R.: Equity, envy, and efficiency. J. Econ. Theory **9**(1), 63–91 (1974)
30. Vazirani, V.V.: Approximation Algorithms. Springer, Heidelberg (2001). https://doi.org/10.1007/978-3-662-04565-7