



Fair Division Is Hard Even for Amicable Agents

Neeldhara Misra and Aditi Sethia^(✉)

Indian Institute of Technology, Gandhinagar, India
{neeldhara.m,aditi.sethia}@iitgn.ac.in

Abstract. We consider the problem of distributing a collection of indivisible objects among agents in a manner that satisfies some desirable notions of fairness and efficiency. We allow agents to “share” goods in order to achieve efficiency and fairness goals which may be otherwise impossible to attain. In this context, our goal is to find allocations that minimize the “amount of sharing”. We follow up on recent work demonstrating that finding fair allocations with minimum sharing is tractable when valuations are non-degenerate, a notion which captures scenarios that are “far from identical”. This result holds for any fixed number of agents. We show that the usefulness of non-degeneracy does not scale to the setting of many agents. In particular, we demonstrate that the problem of finding fractionally Pareto optimal and envy-free allocations is NP-complete even for instances with constant degeneracy and no sharing. We also demonstrate an alternate approach to enumerating distinct consumption graphs for allocations with a small number of sharings.

Keywords: Fair division · Indivisible items · NP-completeness

1 Introduction

The task of fairly distributing indivisible goods among interested agents is challenging already for the simplest possible scenario: one object valued by two or more people. We are typically dealing with m objects and n agents. All agents specify their *value* for each of the objects, and the *utility* they derive from a set of objects is the sum of the values for the individual objects in the collection¹. An *allocation* of m goods to n agents is a partition of the goods into n bundles.

A natural and well-studied notion of fairness is *envy-freeness*, which demands that every agent finds themselves no worse than any other in that they value their own bundle at least as much as any of the other bundles. Note that by itself, envy-freeness can be achieved by trivial allocations where all bundles are empty. This motivates the pursuit of some notion of efficiency—for instance, *completeness*

¹ This is the setting of additive utilities, which will be the focus of our discussion. In a more general setting agents might have different and unrelated utility associated with every possible subset of items.

Supported by the Indian Institute of Technology, Gandhinagar and DST-SERB.

© Springer Nature Switzerland AG 2021

T. Bureš et al. (Eds.): SOFSEM 2021, LNCS 12607, pp. 421–430, 2021.

https://doi.org/10.1007/978-3-030-67731-2_31

requires all items to be allocated, and *fractionally Pareto Optimal* (fPO), where no agent can be made better off without making another worse off. The opening example involving one good valued equally by two agents already shows that there are instances where no allocation is simultaneously complete and envy-free (EF). This has led to several notions of “workarounds”: approximate envy-freeness up to one good [2, 7], or any good [3], or using hidden goods [6], subsidy [1], donating items [4]), and sharing [8, 9]).

Our focus is on settings sharing goods appears to be the most reasonable of all workarounds, and the question of interest is to find allocations that meet our goals of fairness and efficiency with minimum sharing. In a recent development, Sandomirskiy and Segal-Halevi [8] propose a notion of degeneracy which captures the degree of similarity across agent valuations and argue that the intractable cases are those that have a rather high degree of similarity. In retrospect, one might argue that similar valuations signal high conflict, and this possibly contributes to making this a hard scenario. We say that a set of goods are valued similarly by two agents if the ratios of their values for all goods are the same. The degree of similarity between two agents is one less than the largest number of goods that are valued similarly by them. The degeneracy of an instance with n agents is the highest degree of similarity across all pairs of agents. In particular, the degeneracy of an instance with identical valuations is $m - 1$ and it can be as small as zero, when all agents view all goods differently.

Informally speaking, we refer to the setting of low degeneracy, the ones where agent valuations over goods are generally dissimilar, as a scenario involving *amicable agents*. Unlike the case of identical valuations, we expect such valuations to invoke relatively “less conflict”. One of the key results in [8] is that while finding EF allocations remains hard even with two amicable agents, finding allocations that are both fPO and EF is tractable for a constant number of amicable agents. In particular, the time to compute such allocations was shown to be $O(3^{\frac{n(n-1)}{2}d} m^{\frac{n(n-1)}{2}+2})$, where d is the degeneracy of the valuation matrix. In contrast, it was shown that the problem remains NP-hard for instances that have high degeneracy.

Our Contributions. The two results above nicely illustrate the influence of degeneracy on the complexity of finding fPO and EF allocations. We investigate the complexity from the perspective of the number of agents. For example, can this running time be improved to $(n + m)^{O(d)}$, which would increase the realm of tractability to scenarios with any number of agents and constant degeneracy, or more ambitiously, $O(2^{O(d)} \cdot (m + n)^{O(1)})$, which would make the problem tractable for instances with any number of agents and degeneracy logarithmic in $(n + m)$? Our main contribution here is to show that even the former goal is unlikely to be achievable: when the number of agents is unbounded, the problem of finding allocations that are fPO and EF remains *strongly* NP-complete for instances with degeneracy one, even for the specific question of allocations with no sharings.

Our result also has consequences for the problem of finding EF allocations. We recall that the problem of finding EF allocations is weakly NP-complete by

a reduction from PARTITION [8]. It turns out that the arguments in the reverse direction of our reduction do not require the allocation in question to be fPO. Since the valuation matrix of our reduced instance happens to only have values that are bounded by a polynomial function of n and m , we obtain a stronger hardness result for the problem of finding complete EF allocations for instances with constant degeneracy.

We also revisit the algorithm for finding fPO+EF allocations from [8]. The algorithm relies on enumerating certain *consumption graphs* corresponding to fPO allocations that fix the sharing structure of a potential solution, after which the task of determining the exact proportions of sharing while respecting fairness constraints is outsourced to an ILP formulation. It is shown [8, Lemma 2.5] that there always exists a fPO allocation with at most $(n - 1)$ sharings. We propose an alternate method for generating the relevant consumption graphs that takes advantage of the upper bound on the number of sharings upfront. This leads to a slightly different bound that leads to a better exponential term at the cost of a worse polynomial factor. Although the difference in the bound is not significant, we believe our approach lends additional understanding to the structure of class of graphs based on fPO allocations. The arguments regarding this are deferred to the full version of the paper.

2 Preliminaries

We use $\mathcal{A} = \{a_1, \dots, a_n\}$ to denote a set of agents and $\mathcal{G} = \{g_1, \dots, g_m\}$ to denote a collection of objects.

Allocations and Sharing. A *bundle* of objects is a vector $\mathbf{b} = (b_j)_{j \in [m]} \in [0, 1]^m$, where the component b_j represents the portion of g_j in the bundle. The total amount of each object is normalized to one. An *allocation* \mathbf{z} is a collection of bundles $(\mathbf{z}_i)_{i \in [n]}$, one for each agent, with the condition that all the objects are fully allocated. Note that an allocation can be identified with the matrix $\mathbf{z} := (z_{i,j})_{i \in [n], j \in [m]}$ such that all $z_{i,j} \geq 0$ and $\sum_{i \in [n]} z_{i,j} = 1$ for each $j \in [m]$.

Let $j \in [m]$ be arbitrary but fixed. If for some $i \in [n]$, $z_{i,j} = 1$, then the object g_j is not shared—it is fully allocated to agent a_i . Else, object g_j is *shared* between two or more agents.

- The *number of shared objects* is given by the number of items that are shared:

$$\#s^\dagger(\mathbf{z}) = |\{j \in [m] : z_{i,j} \in (0, 1) \text{ for some } i \in [n]\}|.$$

- The *total number of sharings* accounts for the number of times that an object is shared, i.e:

$$\#s^*(\mathbf{z}) = \sum_{j \in [m]} (|\{i \in [n] : z_{i,j} > 0\}| - 1).$$

For allocations with no shared objects, both measures are zero, but they can differ by as much as $m(n - 2)$ in general. Note that the number of sharings is at least the number of shared objects, since each shared object is shared at least once by definition. Unless mentioned otherwise, our measure for “extent of sharing” in the computational questions that we will shortly define will be the notion of the total number of sharings.

Value and Utility. For every $i \in [n], j \in [m]$, $v_{i,j}$ denotes agent a_i ’s *value* for the entire object g_j . In the setting of *additive* utilities, the valuations naturally lead us to a utility function over bundles defined as $u_i(\mathbf{b}) = \sum_{j \in [m]} v_{i,j} \cdot b_j$.

The matrix $\mathbf{v} = (v_{i,j})_{i \in [n], j \in [m]}$ is called the *valuation matrix*; it encodes the information about the preferences of agents and is used as the input of fair division algorithms. We use v^* to denote the largest value in a valuation matrix v . We say that a class of inputs \mathcal{C} has *bounded valuations* if there exists a polynomial $p(n, m)$ such that $v^* \leq p(n, m)$ for all instances in \mathcal{C} .

We recall the notion of degeneracy that was proposed in [8, 9]. To this end, we say that two goods g_p, g_q are valued *similarly* by a pair of agents i, j if there exists a constant r such that $v_{i,p} \cdot v_{j,q} = v_{i,q} \cdot v_{j,p} = r$.

Any collection of goods valued identically by a pair of agents would be pairwise similar with respect to the agents in question, but this definition generalizes the notion of “identical” to, roughly speaking, “identical up to a scaling factor”. Now, let us define the *similarity* between a pair of agents i and j as:

$$s_{\mathbf{v}}(i, j) = \max_{r > 0} \left| \left\{ k \in [m] : v_{i,k} = r \cdot v_{j,k} \right\} \right| - 1.$$

Note that the similarity of a pair of agents captures the notion of the largest number of goods that the agents value similarly when considered pairwise. This finally leads us to the notion of *degeneracy*, which is defined as:

$$d(\mathbf{v}) = \max_{i, j \in [n], i \neq j} s_{\mathbf{v}}(i, j).$$

Valuations for which $d(\mathbf{v}) = 0$ are called *non-degenerate*. Also, note that if any two agents have the same valuations for all goods, then $d(\mathbf{v}) = m - 1$.

Fairness and Efficiency. An allocation $\mathbf{z} = (z_i)_{i \in [n]}$ is called *envy-free (EF)* if every agent prefers her bundle to the bundles of others. Formally, for all $i, j \in [n]$: $u_i(\mathbf{z}_i) \geq u_i(\mathbf{z}_j)$. An allocation \mathbf{z} is *proportional (Prop)* if each agent prefers her bundle to the equal division: $\forall i \in [n] \ u_i(\mathbf{z}_i) \geq \frac{1}{n} \sum_{o \in [m]} v_{i,o}$. An allocation \mathbf{z} is *equitable (EQ)* if any pair of agents derive equal utility from their respective bundles. Formally, for all $i, j \in [n]$: $u_i(\mathbf{z}_i) = u_j(\mathbf{z}_j)$.

An allocation \mathbf{z} is Pareto-dominated by an allocation \mathbf{y} if \mathbf{y} gives at least the same utility to all agents and strictly more to at least one of them. An allocation \mathbf{z} is *fractionally Pareto-optimal (fPO)* if no feasible \mathbf{y} dominates it. If \mathbf{y} is such that $y_{i,o} \in \{0, 1\}$, then \mathbf{z} is called *discrete Pareto-Optimal (dPO)*. The following lemma provides a complete characterisation of fPO allocations.

Lemma 1 ([8], **Lemma 2.3**). *An allocation \mathbf{z} is fractionally Pareto Optimal if and only if there exists a vector of weights $\lambda = (\lambda_i)_{i \in [n]}$ with $\lambda_i > 0$, such that for all agents $i \in [n]$ and goods $p \in [m]$, if $z_{i,p} > 0$ then for any agent $j \in [n]$,*

$$\lambda_i \cdot v_{i,p} \geq \lambda_j \cdot v_{j,p}$$

Computational Questions. Formally, for a fairness concept $\alpha \in \{\text{EF}, \text{EQ}, \text{Prop}\}$ and an efficiency concept $\beta \in \{\text{fPO}, \text{dPO}\}$, the (α, β) -MINIMAL SHARING problem is the following. Given $(\mathcal{A}, \mathcal{G}, \mathbf{v}, t \in \mathbb{N})$ as input, the question is if there exists an α, β allocation where the total number of sharings is at most t . In this paper, we focus on $\{\text{EF}, \text{fPO}\}$ -minimal sharing problem.

3 Hardness for Instances of Constant Degeneracy

To prove the hardness of the minimal sharing problem, we will show a reduction from a structured version of SATISFIABILITY problem called LINEAR NEAR-EXACT SATISFIABILITY (LNES) which is known to be NP-complete [5]. An instance of LNES consists of $5p$ clauses (where $p \in \mathbb{N}$) denoted as follows:

$$\mathcal{C} = \{U_1, V_1, U'_1, V'_1, \dots, U_p, V_p, U'_p, V'_p\} \cup \{C_1, \dots, C_p\}.$$

We will refer to the first $4p$ clauses as the *core* clauses, and the remaining clauses as the *auxiliary* clauses. The set of variables consists of p *main variables* x_1, \dots, x_p and $4p$ *shadow variables* y_1, \dots, y_{4p} . Each core clause consists of two literals $\forall i \in [p], U_i \cap V_i = \{x_i\}$ and $U'_i \cap V'_i = \{\bar{x}_i\}$. Each main variable x_i occurs exactly twice as a positive literal and exactly twice as a negative literal. The main variables only occur in the core clauses. Each shadow variable makes two appearances: as a positive literal in an auxiliary clause and as a negative literal in a core clause. Each auxiliary clause consists of four literals, each corresponding to a positive occurrence of a shadow variable. We will use u_i, v_i, u'_i , and v'_i to refer to the shadow variables in the main clauses U_i, V_i, U'_i , and V'_i , respectively.

The LNES problem asks whether, given a set of clauses with the aforementioned structure, there exists an assignment τ of truth values to the variables such that *exactly one* literal in every core clause and *exactly two* literals in every auxiliary clause evaluate to TRUE under τ . The main result of this section is the following, and is established by a reduction from LNES.

Theorem 1. *(EF,fPO)-MINIMAL SHARING is NP-hard even when restricted to inputs with bounded valuations, degeneracy one, and no sharing.*

Proof. We reduce from LNES. Let $\mathcal{C} = \{U_1, V_1, U'_1, V'_1, \dots, U_p, V_p, U'_p, V'_p\} \cup \{C_1, \dots, C_p\}$ be an instance of LNES as described above.

We begin with a description of the construction of the reduced instance. We refer the reader to Fig. 1 for a high-level schematic of this construction. For each main variable x_i we introduce three agents: $\{a_i, \bar{a}_i, d_i\}$, and the goods $\{g_i, \bar{g}_i, h_i\}$.

We refer to d_i as the *dummy* agent associated with x_i and a_i and \bar{a}_i as the *key* agents associated with x_i . Also, we refer to h_i as the *trigger* good and g_i and \bar{g}_i as *consolation* goods. For the shadow variables u_i, v_i, u'_i, v'_i , we introduce four agents: b_i, c_i, b'_i, c'_i which we simply refer to as *shadow agents* and four goods: r_i, s_i, r'_i, s'_i , which we refer to as the *essential* goods. Finally, for each auxiliary clause C_j , we introduce the goods f_j^1 and f_j^2 . These goods are called *backup* goods.

Note that our instance consists of $n = 7p$ agents and $m = 9p$ goods. Thus the size of the valuation matrix is $N := 63 \cdot p^2$. We let $L = 4000 \cdot p^5$. We will use \mathcal{A} and \mathcal{G} to refer to the set of agents and goods that we have defined here.

Let $\mathbf{w} = (w_{i,j})_{i \in [n], j \in [m]}$ denote the $(7p \times 9p)$ matrix whose entries are given by $w_{i,j} = (i - 1) \cdot m + j$. Intuitively, we can think of these values as being small enough to be negligible, and we will obtain our final valuation matrix by starting from \mathbf{w} and “overwriting” some entries to reflect the fact that certain goods are valued highly by certain agents. This is done to ensure that the final valuation matrix has low degeneracy. We now describe the specific modifications that we have to make to \mathbf{w} .

To this end, let us define another set of values given by $\mathbf{w}^* = (w_{i,j}^*)_{i \in [n], j \in [m]}$. Let $\pi : \mathcal{A} \rightarrow [n]$ and $\sigma : \mathcal{G} \rightarrow [m]$ be arbitrary but fixed orderings of the agents and goods, respectively.

- For $i \in [p]$, we have that the dummy agent corresponding to the main variable x_i has a high value for the consolation goods g_i and \bar{g}_i .

$$w_{\pi(d_i),j}^* = \begin{cases} L & \text{if } \sigma^{-1}(j) \in \{g_i, \bar{g}_i\}, \\ 0 & \text{otherwise.} \end{cases}$$

- For $i \in [p]$, we have that the first key agent corresponding to the main variable x_i has a somewhat high value for the consolation good g_i and the essential goods r_i and s_i , and a high value for the trigger good h_i .

$$w_{\pi(a_i),j}^* = \begin{cases} L/3 & \text{if } \sigma^{-1}(j) \in \{g_i, r_i, s_i\}, \\ L & \text{if } \sigma^{-1}(j) = h_i, \\ 0 & \text{otherwise.} \end{cases}$$

- For $i \in [p]$, we have that the second key agent corresponding to the main variable x_i has a somewhat high value for the consolation good \bar{g}_i and the essential goods r'_i and s'_i , and also has a high value for the trigger good h_i .

$$w_{\pi(\bar{a}_i),j}^* = \begin{cases} L/3 & \text{if } \sigma^{-1}(j) \in \{\bar{g}_i, r'_i, s'_i\}, \\ L & \text{if } \sigma^{-1}(j) = h_i, \\ 0 & \text{otherwise.} \end{cases}$$

- For $i \in [p]$ the shadow agents have a high value for their associated essential goods and the backup good which represents an auxiliary clause that contains the shadow variable associated with the shadow agent. Formally, we have:

$$w_{\pi(b_i),j}^* = \begin{cases} L & \text{if } \sigma^{-1}(j) \in \{r_i, f_\ell^1, f_\ell^2\}, \\ 0 & \text{otherwise.} \end{cases}$$

where ℓ is such that C_ℓ is the unique clause that contains the shadow variable u_i . The valuations for $w_{\pi(c_i),j}$, $w_{\pi(b'_i),j}$ and $w_{\pi(c'_i),j}$ are analogously defined, with r_i being replaced by s_i , r'_i , and s'_i , respectively, and ℓ would be such that C_ℓ is the unique clause that contains v_i , u'_i , and v'_i , respectively.

The final valuations that we will work with are obtained by taking a point-wise max of the two valuation matrices defined above with the following exceptions:

- Dummy agents value the four essential goods associated with them at zero.
- The shadow agent b_i (respectively, c_i) values the consolation good g_i and the essential good s_i (respectively, r_i) at zero.
- The shadow agent b'_i (respectively, c'_i) values the consolation good \bar{g}_i and the essential good s'_i (respectively, r'_i) at zero.

In particular, we propose the final valuation matrix $\mathbf{v} = (v_{i,j})_{i \in [n], j \in [m]}$ as follows:

$$v_{i,j} = \begin{cases} \min(w_{i,j}, w_{i,j}^*) & \text{if } \pi^{-1}(i) = d_k \text{ and } \sigma^{-1}(j) \in \{r_k, s_k, r'_k, s'_k\}, \\ & \text{or } \pi^{-1}(i) = b_k \text{ and } \sigma^{-1}(j) \in \{g_k, s_k\}, \\ & \text{or } \pi^{-1}(i) = c_k \text{ and } \sigma^{-1}(j) \in \{g_k, r_k\}, \\ & \text{or } \pi^{-1}(i) = b'_k \text{ and } \sigma^{-1}(j) \in \{\bar{g}_k, s'_k\}, \\ & \text{or } \pi^{-1}(i) = c'_k \text{ and } \sigma^{-1}(j) \in \{\bar{g}_k, r'_k\}, \\ & \text{for any } k \in [p] \\ \max(w_{i,j}, w_{i,j}^*) & \text{otherwise.} \end{cases}$$

For convenience, we say an entry of \mathbf{v} is *large* if it is at least $L/3$ and is *small* otherwise. For (i, j) which are such that $v_{i,j}$ is small, we introduce the notation $\varepsilon_{i,j}$ to denote $v_{i,j}$. We ask if this instance admits an allocation with zero sharing. We now argue the equivalence of the instances.

Forward Direction. Let τ be a boolean assignment for the variables of the LNES instance that we start with. We now propose an allocation:

- If $\tau(x_i) = 1$, then the first key agent a_i gets $\{g_i, r_i, s_i\}$, the second key agent \bar{a}_i gets the trigger good h_i and the dummy agent d_i gets the consolation good $\{\bar{g}_i\}$.
- If $\tau(x_i) = 0$, then the first key agent a_i gets the trigger good $\{h_i\}$, the second key agent \bar{a}_i gets $\{\bar{g}_i, r'_i, s'_i\}$ and the dummy agent gets the consolation good $\{g_i\}$.
- If $\tau(x_i) = 1$, then the shadow agents b'_i and c'_i get the essential goods that they value highly, i.e. r'_i and s'_i .

	g_i	\bar{g}_i	h_i	r_i	s_i	r'_i	s'_i	f_j^1	f_j^2
a_i	L/3	*	L	L/3	L/3	*	*	*	*
\bar{a}_i	*	L/3	L	*	*	L/3	L/3	*	*
d_i	L	L	*	*	*	*	*	*	*
b_i	*	*	*	L	*	*	*	L	L
c_i	*	*	*	*	L	*	*	*	*
b'_i	*	*	*	*	*	L	*	*	*
c'_i	*	*	*	*	*	*	L	*	*

Fig. 1. The overall schematic of the construction in the proof of Theorem 1. The entries depicted by a \star indicate small values

- If $\tau(x_i) = 0$, then the shadow agents b_i and c_i get the essential goods that they value highly, i.e. r_i and s_i .

Note that there are $2p$ shadow agents who have not been allocated any goods so far. It is easy to check that these shadow agents correspond exactly to shadow variables x for which $\tau(x) = 1$. Since τ is a satisfying assignment for the LNES instance, we know that each auxiliary clause C_ℓ contains exactly two shadow variables which evaluate to true under τ . Let $\mu(C_\ell)$ denote the shadow agents corresponding to these shadow variables. Then, for each $j \in [p]$, the goods f_j^1 and f_j^2 are allocated arbitrarily, one each, to the two shadow agents in $\mu(C_j)$.

We claim that this allocation is fPO and EF, and we defer the proofs of these properties to the full version.

Reverse Direction. For the discussion in the reverse direction, we say that an allocation is *valid* if it is EF and fPO and involves no sharing. Let $\mathbf{z} := (z_{i,j})_{i \in [n], j \in [m]}$ be a valid allocation. First, we argue that \mathbf{z} must have a certain structure in a series of claims whose proofs are deferred to the full version.

Claim. In the allocation \mathbf{z} , any trigger good h_i must be allocated to one of the corresponding key agents $\{a_i, \bar{a}_i\}$.

Claim. Every consolation good g_i is allocated to either to the key agent a_i or to the dummy agent d_i . Likewise, the good \bar{g}_i is allocated to either to the key agent \bar{a}_i or to the dummy agent d_i .

Claim. If the consolation good g_i is allocated to a key agent a_i , then the shadow agents b_i and c_i must be allocated the backup goods.

Claim. If the consolation good \bar{g}_i is allocated to a key agent \bar{a}_i , then the shadow agents b'_i and c'_i must be allocated the backup goods.

Now observe that the two claims above account for the allocation of $2p$ backup goods among $2p$ distinct shadow agents. Let us call these shadow agents *happy* and the remaining shadow agents *unhappy*. We claim that the bundle of every unhappy shadow agent must contain an essential good—this is because these are the only highly valued goods left in the pool and are the only way to eliminate the envy that the unhappy agents feel for the happy ones. Note that every unhappy agent values the bundle of exactly two happy shadow agents.

Based on this, we propose the following assignment of truth values:

$$\tau(x_i) = \begin{cases} 1 & z_{\pi(a_i), \sigma(g_i)} = 1, \\ 0 & \text{otherwise.} \end{cases}$$

We extend this assignment to shadow variables in the natural way: if $\tau(x_i) = 1$, then $\tau(u_i) = \tau(v_i) = 1$ and $\tau(u'_i) = \tau(v'_i) = 0$, while if $\tau(x_i) = 0$, then $\tau(u_i) = \tau(v_i) = 0$ and $\tau(u'_i) = \tau(v'_i) = 1$. We now argue that τ is a satisfying assignment for the original LNES instance.

Suppose g_i is allocated to a_i . We set $\tau(x_i) = 1$. This satisfies all the clauses containing the literal x_i , namely, U_i and V_i . Further, note that these clauses are satisfied exactly once, since we also set $\tau(u_i) = \tau(v_i) = 1$ (recall that u_i and v_i appear in these clauses with negative polarity). The other main clauses U'_i and V'_i are satisfied since we set $\tau(u'_i) = \tau(v'_i) = 0$, and these clauses are satisfied exactly once as well, since x_i appears in them with a negative polarity and we are in the case when $\tau(x_i) = 1$. The case when $\tau(x_i) = 0$ is analogous, and we see that all core clauses are satisfied exactly once by τ , as desired.

We now turn to the auxiliary clauses. Observe that $\tau(x_i) = 1$ if and only if $z_{\pi(a_i), \sigma(g_i)} = 1$, that is, the key agent a_i gets the consolation good g_i . This implies that b_i and c_i are happy agents. On the other hand, recall that we also set $\tau(u_i)$ and $\tau(v_i)$ to one. Similarly, it can be argued that if $\tau(x_i) = 0$, then b'_i and c'_i are happy agents, and in this case, we had also set $\tau(u'_i)$ and $\tau(v'_i)$ to one. So we conclude that all happy agents correspond to variables that evaluate to one under τ . Along similar lines, it is easy to check that all unhappy agents who receive essential goods as explained in the last claim correspond to variables that are set to zero under τ .

Now consider an auxiliary clause C_ℓ . Notice that f_ℓ^1 and f_ℓ^2 have been allocated to happy agents that value these goods highly, so we know that C_ℓ contains at least two variables that evaluate to true. Now suppose there is some auxiliary clause that contains more than two variables that evaluate to true. This would imply the existence of more than $2p$ happy agents, which is a contradiction. The argument for the reduced instance having constant degeneracy is deferred to the full version. □

4 Concluding Remarks

We demonstrated the hardness of finding fPO+EF and EF allocations even for instances with constant degeneracy for instances with an unbounded number of agents. We note that running times of the form $d^{O(n)} \cdot \text{poly}(m, n)$ are “weakly ruled out” because of the hardness result in [8] which is based on a reduction from PARTITION. However, all the hardness results combined so far do not rule out the possibility of an algorithm with a running time of $c^{O(d+n)} \cdot m^{O(1)}$, which would imply strongly polynomial running times for instances where $(d + n)$ is bounded by $O(\log m)$. One framework to rule out such a possibility would be parameterized complexity, where one might attempt demonstrating W-hardness in the combined parameter (n, d) . On a related note, we show that instances that have bounded degeneracy and a bounded number of values in the valuation matrix are essentially bounded—we refer the reader to full version of the paper for a more detailed discussion.

References

1. Brustle, J., Dippel, J., Narayan, V.V., Suzuki, M., Vetta, A.: One dollar each eliminates envy. In: Proceedings of the 2020 ACM Conference on Economics and Computation (2019)
2. Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. In: Proceedings of the Behavioral and Quantitative Game Theory BQGT, p. 74:1. ACM (2010)
3. 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)
4. Chaudhury, B.R., Kavitha, T., Mehlhorn, K., Sgouritsa, A.: A little charity guarantees almost envy-freeness. In: Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, pp. 2658–2672. SIAM (2020)
5. Dayal, P., Misra, N.: Deleting to structured trees. In: Du, D.-Z., Duan, Z., Tian, C. (eds.) COCOON 2019. LNCS, vol. 11653, pp. 128–139. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26176-4_11
6. Hosseini, H., Sikdar, S., Vaish, R., Wang, J., Xia, L.: Fair division through information withholding. In: Proceedings of the Thirty-Fourth AAI Conference on Artificial Intelligence (AAAI) (2020)
7. 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. ACM (2004)
8. Sandomirskiy, F., Segal-Halevi, E.: Fair division with minimal sharing. CoRR abs/1908.01669 (2019). <http://arxiv.org/abs/1908.01669>
9. Segal-Halevi, E.: Fair division with bounded sharing. CoRR abs/1912.00459 (2019). <http://arxiv.org/abs/1912.00459>