Computing Efficient Envy-Free Partial Allocations of Indivisible Goods

Robert Bredereck [a,*], Andrzej Kaczmarczyk [b], Junjie Luo [c] and Bin Sun [a]

[a]Institut für Informatik, TU Clausthal, Clausthal-Zellerfeld, Germany
[b]AGH University, Kraków, Poland
[c]School of Mathematics and Statistics, Beijing Jiaotong University, Beijing, China
ORCID (Robert Bredereck): https://orcid.org/0000-0002-6303-6276, ORCID (Andrzej Kaczmarczyk): https://orcid.org/0000-0003-1401-0157, ORCID (Junjie Luo): https://orcid.org/0000-0001-8892-8863, ORCID (Bin Sun): https://orcid.org/0009-0001-9127-2452

Abstract. Envy-freeness is one of the most prominent fairness concepts of indivisible goods allocation. Even though trivial envy-free allocations always exist, rich literature shows that this is not true when one additionally requires some efficiency concept (like, e.g., completeness, Pareto-efficiency, maximization of social welfare). In fact, in such case even deciding the existence of an efficient envy-free allocation is notoriously computationally hard. On our way to deeply study the limits of efficient computability of such allocations, we relax the standard efficiency concepts and analyze how it impacts the computational complexity of the respective problems. Specifically, we allow for partial allocations (where not all goods are allocated) and impose only very mild efficiency constraints (e.g., we require each agent to have a positive utility from their bundle). Surprisingly, even such seemingly weak efficiency requirements lead to a diverse computational complexity landscape. We identify several polynomial-time solvable cases, yet we also find NP-hardness for very restricted scenarios of ternary (cardinal) utilities.

1 Introduction

Computing fair allocations of indivisible resources is an important issue with many applications in all kinds of disciplines [8, 9, 22]. Envy-freeness, which ensures that no agent strictly prefers the resources allocated to a different agent over their own, is one of the most prominent fairness concepts [9]. Unfortunately, envy-free allocations do not always exist, and computing them is often associated to computationally very difficult problems [8]. In consequence, researchers have developed several ways to relax that fairness notion, such as envy-free up to one good (EF1) [11] and envy-free up to any good (EFX) [13].

If one has a close look, however, then one quickly realizes that envy-freeness alone does not enforce any computational or existence issues: allocating no resource to anyone is envy-free. Only when adding an *efficiency* component, such as requiring each resource to be allocated to someone (completeness), the picture changes. A folklore example is an instance with n agents (say employees) and $n+1$ identical resources (say laptops): in every possible complete allocation there is at least one agent a who gets at most one resource and another agent a' that gets at least two resources, so that (for reasonable preferences) a envies a'. While there are certainly applications where this is indeed a problem, there is likely a trivial solution in most applications: allocating only n of the $n+1$ resources (one to each agent). Such observations lead to the main question of our paper: which (weaker) efficiency concepts can help to identify additional (cf. completeness) envy-free allocations and what is the consequence on the computational complexity of finding such allocations?

We come up with two basic ideas: What if the goal is not to allocate *all resources*, but to either just allocate *some resources* to the agents or just provide *some utility* for the agents? In each case, we can focus on either the whole society or individual agents. More concretely, we ask for an envy-free (partial) allocation that (i) allocates at least t resources in total, or (ii) allocates at least t resources to each agent, or (iii) has utilitarian welfare of at least t, or (iv) has egalitarian welfare of at least t.

Note that even variants for $t = 1$ have meaningful (potential) applications. They allow us to ask if there is an envy-free allocation of (some of) the resources such that (i) at least one resource is allocated, (ii) each agent gets at least one resource, (iii) at least one agent has a positive value for the allocated resources, or (iv) each agent has a positive value for the allocated resources. The first two cases (i,ii) model natural formal requirements while the other two cases (iii,iv) model basic (individual) quality requirements.

The efficiency requirements are also relevant from the computational complexity perspective. To see this, assume—as we do in our paper—that the resources are goods, that is, agents report nonnegative utilities for them. In this case, all our efficiency concepts for $t = 1$ are significantly less demanding than multiple other prominent efficiency concepts; for example, than completeness, as we demonstrated using the folklore example earlier in this section. Hence, analyzing computational complexity, especially showing hardness, of these very special cases allows us to identify borders of efficient computability more accurately than before. On the other hand, if we find efficient algorithms for these relaxed cases, their results can be practically interpreted as the minimum efficiency levels that can be achieved. Indeed, given an instance of an allocation problem, by computing the result with such an algorithm, one can argue that any fair allocation that is less efficient is unjustified. Before we describe our findings, we briefly review the related literature to present the context helpful to interpret our results.

* Corresponding Author. Email: robert.bredereck@tu-clausthal.de

1.1 Related Work

Computing fair and efficient allocations has recently emerged as a very prominent stream of research in the area of fair allocation of indivisible resources. Allocations with maximum Nash welfare are both Pareto optimal and EF1, but computing such allocations is NP-hard [13]. Likewise, computing an allocation with the highest utilitarian social welfare among all EF1 allocations is NP-hard even for two agents [4]. As discussed before, the main difference of our model is that we allow partial allocations, and consequently we consider envy-freeness instead of its relaxation EF1.

Allowing partial allocations is an important approach to guarantee the existence of EFX allocations (the existence of EFX complete allocations is still an open question). Caragiannis et al. [12] showed that there always exists an EFX partial allocation with at least half of the maximum Nash welfare. Chaudhury et al. [14] showed that donating at most $n-1$ resources can guarantee the existence of EFX allocation such that no agent prefers the donated resources to its own bundle, where n denotes the number of agents. This bound was latter improved to $n-2$ in general and to 1 for the case with four agents [5]. Besides existence, Bu et al. [10] studied the problem of computing partial allocations with the maximum utilitarian welfare among all EFX allocations. Our work differs from this stream of research in that we focus on envy-freeness instead of EFX.

Aziz et al. [3] studied the problem of deleting (or adding) a minimum number of resources such that the resulting instance admits an envy-free allocation; which is equivalent to finding an envy-free allocation with the maximum size. However, they consider ordinal preferences whereas we consider cardinal preferences. Moreover, Aziz et al. [3] considered the number of deleted resources, where the problem is NP-hard even if no resource can be deleted. In contrast, we consider the dual parameter the lower bound on the allocated resources, which allows us to identify polynomial-time solvable cases.

Boehmer et al. [7] studied the problem of transforming a given unfair allocation into an EF or EF1 allocation by donating few resources. In addition to upper bounds on the number of donated resources and the decrease on the utilitarian welfare, they also consider the lower bounds on the remaining allocated resources and the remaining utilitarian welfare. Dorn et al. [15] studied the same problem but focused on a different fairness notion. The most prominent difference to our work is that in our model there is no given allocation.

Hosseini et al. [18] introduced a fairness notion where agents can hide some of the resources in their own bundles such that no agent is envious assuming that the agents do not know the existence of the hidden resources in other agents' bundles. Then the goal is to find a complete allocation and a minimum number of hidden resources such that no agent is envious. While the idea is similar to find an envy-free partial allocation with the maximum size, note that the hidden resources are not deleted; their owners get utility from them just like normal resources.

A series of works [16, 6, 19] studied the computational complexity of finding an envy-free house allocation when the number of houses is larger than the number of agents. This is equivalent to finding an envy-free (partial) allocation that allocates exactly one resource to each agent. Our model does not have this kind of upper bound on the number of resources allocated to each agent. Aigner-Horev and Segal-Halevi [1] studied the problem of finding an envy-free matching of maximum cardinality in a bipartite graph. Taking the bipartite graph as the representation of binary utilities of agents on one side towards resources on the other side, the problem studied by Aigner-Horev and Segal-Halevi [1] is equivalent to finding an envy-free (partial) allocation with the maximum size such that each agent gets at most one resource liked by it. Our model differs from it in that we do not add an upper bound for agents' bundles and we allow agents to receive resources with utility 0. Nevertheless, many of our algorithms for binary utilities use the structural properties of envy-free matching established by Aigner-Horev and Segal-Halevi [1].

Table 1. Summary of results. Columns denote different utility constraints and efficiency threshold t values. Rows represent different efficiency concepts $\mathcal{E}$. The hardness results for $t = 1$ apply to every positive t as well.

	Identical	Binary		Ternary
	$t = 1$	$t = 1$	t	$t = 1$
utilitarian social welfare (usw)	NP-h	P	NP-h (FPT)	NP-h
egalitarian social welfare (esw)		P	P	NP-h
#resources allocated (size)		P	NP-h (FPT)	NP-h
min-cardinality (mcar)		P	NP-h	NP-h

1.2 Contributions and Outline

We study the computational complexity of finding envy-free partial allocations with mild efficiency requirements. To this end, we consider a lower bound t on utilitarian welfare, egalitarian welfare, the number of allocated resources, or the minimum bundle size among all agents. Formal definitions can be found in Section 2. In Section 3, we show that finding such allocations is strongly NP-hard, even if all agents have identical preferences. In Section 4, we show that for binary preferences all variants are polynomial-time solvable when $t = 1$, but that they are in general strongly NP-hard. A surprising exception is the egalitarian welfare variant (which is typically harder than the utilitarian welfare): we show a polynomial-time algorithm that finds an envy-free partial allocation where each agent obtains a bundle with value at least t (for arbitrary t). To overcome NP-hardness, we show that the utilitarian welfare variant and the number of allocated resources variant are both fixed-parameter tractable (FPT) with respect to t[1], implying that the problems can be efficiently solved for small t. In Section 5, we go beyond binary preferences and allow for three different utility values. We show that all variants become strongly NP-hard already when $t = 1$ for any ternary utility values $\{0, v, u\}$ with $0 < v < u$. See Table 1 for an overview of our results. Some details (e.g., of proofs of theorems marked with (★) will be presented in the full version of this paper in the future.

2 Preliminaries

We fix a collection $\mathcal{R}$ of m *resources* and a set $\mathcal{A}$ of n *agents*. Each agent $a \in \mathcal{A}$ reports its cardinal *utility* from each resource via the *utility function* $u_a : \mathcal{R} \to \mathbb{N}_0$.[2] We assume additive utilities, hence, with a slight abuse of notation, for some set $B \subseteq \mathcal{R}$ of resources, the utility $u_a(B)$ of agent $a \in \mathcal{A}$ from B is the sum of the agent's utilities for each resource in B, i.e., $u_a(B) \coloneqq \sum_{r \in B} u_a(r)$.

We often use specific classes of cardinal utilities reported by agents. *Identical utilities* denote a family of utilities in which every agent's utility functions are the same. The utilities are *binary* if

[1] A problem is fixed-parameter tractable with respect to some parameter k if it can be solved in $f(k)|I|^{O(1)}$ time, where $|I|$ denotes the input size.

[2] $\mathbb{N}_0$ denotes the set of all nonnegative integers.

agents' utilities use only values 0 or 1. We also define *ternary* utilities, where there are always three possible values of utility that agents can report.

An *allocation* $\boldsymbol{\pi} : \mathcal{A} \to 2^{\mathcal{R}}$ assigns each agent $a \in \mathcal{A}$ its *bundle* $\boldsymbol{\pi}(a)$ for their private use, i.e. $\boldsymbol{\pi}(a) \cap \boldsymbol{\pi}(a') = \emptyset$ for each distinct $a, a' \in \mathcal{A}$. If $\boldsymbol{\pi}(i) = \emptyset$, it is an *empty bundle*. If $\boldsymbol{\pi}$ is a partition of $\mathcal{R}$, we say that $\boldsymbol{\pi}$ is *complete*, otherwise we call it *partial*. We call the smallest number $\mathrm{mcar}(\boldsymbol{\pi}) := \min_{a\in\mathcal{A}} |\boldsymbol{\pi}(a)|$ of resources allocated to some agent the *min-cardinality* of $\boldsymbol{\pi}$, whereas by $\mathrm{size}(\boldsymbol{\pi}) := \sum_{a\in\mathcal{A}} |\boldsymbol{\pi}(a)|$ we denote the total number of resources allocated by $\boldsymbol{\pi}$.

Given an allocation $\boldsymbol{\pi} : \mathcal{A} \to 2^{\mathcal{R}}$ and some collection $(u_a)_{a\in\mathcal{A}}$ of utility functions, we say that agent $a \in \mathcal{A}$ is *envious* regarding (u_a) under $\boldsymbol{\pi}$ if there is another agent $a' \in \mathcal{A}$ whose bundle $\boldsymbol{\pi}(a')$ is preferred by a over their own bundle $\boldsymbol{\pi}(a)$; formally $u_a(\boldsymbol{\pi}(a')) < u_a(\boldsymbol{\pi}(a))$. An allocation $\boldsymbol{\pi}$ is *envy-free* regarding (u_a) if no agent is envious under $\boldsymbol{\pi}$. The *utilitarian social welfare* $\mathrm{usw}(\boldsymbol{\pi})$ of $\boldsymbol{\pi}$ regarding (u_a) is the sum of the utilities of agents for their bundles, i.e., $\mathrm{usw}(\boldsymbol{\pi}) := \sum_{a\in\mathcal{A}} u_a(\boldsymbol{\pi}(a))$. Analogously, *egalitarian social welfare* $\mathrm{esw}(\boldsymbol{\pi})$ is the minimum of the agent's utilities, i.e., $\mathrm{esw}(\boldsymbol{\pi}) := \min_{a\in\mathcal{A}} u_a(\boldsymbol{\pi}(a))$. (We omit "regarding (u_a)" and "under $\boldsymbol{\pi}$", respectively, when the context is clear.)

Our problem of interest is a computational problem of deciding if, for a given input, one can find allocations that are envy-free and efficient. Following the introduction, we define our problem generally, using an efficiency measure placeholder $\mathcal{E}$ to be substituted by any of the efficiency measures of our interest: utilitarian and egalitarian social welfare, size, and min-cardinality.

$\mathcal{E}$-ENVY-FREE PARTIAL ALLOCATION ($\mathcal{E}$-EF-PA)
Input: A set $\mathcal{R}$ of resources, a set $\mathcal{A}$ of agents, a collection $(u_a)_{a\in\mathcal{A}}$ of utility functions $u_a\colon \mathcal{R} \to \mathbb{N}_0$ and an efficiency threshold t.
Question: Is there an envy-free allocation $\boldsymbol{\pi}$ such that $\mathcal{E}(\boldsymbol{\pi}) \geq t$?

3 Identical valuations

The case in which all agents have identical preferences is potentially simpler to solve than the general case when finding our desired allocations. However, we show even in this scenario, our problem is NP-hard for each efficiency notion we consider.

We show hardness via a reduction from the 3-PARTITION problem [17]. The main idea is to have one resource for each number of the 3-PARTITION instance as well as some well-designed dummy resources and extra agents, ensuring that each agent receives either one dummy resource or three non-dummy resources such that the utility for them adds up to the same value as the agents have for a dummy resource.

Theorem 1 (★). *For each $\mathcal{E} \in \{\mathrm{usw}, \mathrm{esw}, \mathrm{size}, \mathrm{mcar}\}$ it holds that $\mathcal{E}$-EF-PA is strongly NP-hard, even if $t = 1$ and each agent has the same utility function.*

The presented result categorically sets the limits of our expectations, as the hardness holds for the weakest variants of efficiency concepts, that is, when the threshold $t = 1$. Hence, we need to focus on another way of constraining the agent preferences to find polynomial-time tractable cases.

In the remaining sections, we will focus on restrictions on the set of utilities (resp. the images of the utility functions), since it seems essential that they are unrestricted in the above hardness reduction for identical preferences.

4 Binary utilities

The case of binary utilities, where agents express preferences by pointing out which resources they desire and which not, is another natural constraint to our problems. Given that identifying exact utility values imposes a high cognitive burden for human agents, in practice, binary utilities are sometimes even preferred over more complicated variants. It is then easier to elicit correct preference data and to avoid excessive fatigue of the agents.

The good news is that for binary preferences, our problem with $t = 1$ is solvable in polynomial time for all the four efficiency notions. On the negative side, $t = 1$ is where the good news ends. For arbitrary t, except for esw-EF-PA, the other three efficiency concepts yield NP-hardness. For some of these cases, however, we could find efficient (FPT) algorithms for bounded values of the threshold t.

4.1 Egalitarian social welfare

Beginning with esw-EF-PA, we show that it is polynomial-time solvable by providing a reduction to computing a maximum cardinality matching in bipartite graphs.

Theorem 2. *For $0/1$-utilities* esw*-EF-PA is solvable in $O(m^{2.5})$ time.*

Proof. If $t > \frac{m}{n}$, then no allocation can get $\mathrm{esw}(\boldsymbol{\pi}) \geq t$. So, in the following we assume $t \leq \frac{m}{n}$. Given an envy-free allocation $\boldsymbol{\pi}$ with $\mathrm{esw}(\boldsymbol{\pi}) \geq t$, we construct a new allocation $\boldsymbol{\pi}'$ by keeping t arbitrary resources from each agent's bundle that are liked by the agent and deleting the other resources. Note that $\boldsymbol{\pi}'$ also satisfies envy-freeness and $\mathrm{esw}(\boldsymbol{\pi}') \geq t$. Therefore, it suffices to check whether there exists an allocation such that every agent gets exactly t resources liked by it. To this end, we create a bipartite graph where one side consists of t copies of each agent and the other side consists of all resources, and there is an edge between an agent and a resource if the agent likes the resource. Then there exists an envy-free allocation with $\mathrm{esw}(\boldsymbol{\pi}) \geq t$ if and only if a maximum cardinality matching of this bipartite graph, which can be computed in $O((tn)^{1.5}m) = O(m^{2.5})$ time [20], saturates the agent side. □

4.2 Utilizing envy-free matchings

For the other three efficiency measures, we create a bipartite graph $G = (X\dot{\cup}Y, E)$, where $X = \mathcal{A}$, $Y = \mathcal{R}$, and there is an edge between $x_i \in X$ and $y_j \in Y$ if $u_i(r_j) = 1$. We use the concept of *envy-free matchings* (EFM) for bipartite graphs introduced by Aigner-Horev and Segal-Halevi [1]. A matching M in a bipartite graph $G = (X\dot{\cup}Y, E)$ is envy-free with respect to X if no vertex in $X \setminus X_M$ is adjacent to any vertex in Y_M, where X_M (resp. Y_M) represents the set of vertices from X (resp. Y) saturated by M. Note that each envy-free matching M in $G = (X\dot{\cup}Y, E)$ induces an envy-free allocation $\boldsymbol{\pi}^M$, where every agent gets at most one resource. Slightly abusing the notation, we sometimes use subsets of X (resp. Y) to denote the corresponding subsets of agents (resp. resources).

Aigner-Horev and Segal-Halevi [1] show that finding an envy-free matching of maximum cardinality is solvable in polynomial time. The idea is to first compute an arbitrary matching M of maximum cardinality. Then, starting with each vertex from X that are not saturated by M, we find M-alternating paths, which partition the vertex set into two parts according to whether they are covered by these paths or not. It is shown that this partition is independent of the initial matching M and that all envy-free matchings are contained in the part not covered by the above M-alternating paths. In the following theorem, we summarize the findings of Aigner-Horev and Segal-Halevi [1] related to envy-free matchings that are relevant to our results.

Theorem 3 ([1]). *Every bipartite graph $G = (X\dot{\cup}Y, E)$ admits a unique partition $X = X_S\dot{\cup}X_L$ and $Y = Y_S\dot{\cup}Y_L$, called the* EFM partition *of G, satisfying the following conditions:*

1. *Every X_L-saturating matching in $G[X_L; Y_L]$ is an envy-free matching in G;*
2. *Every envy-free matching in G is contained in $G[X_L; Y_L]$;*
3. *There are no edges between X_S and Y_L;*
4. *Each vertex in Y_S is connected to at least one vertex in X_S.*

Moreover, the unique EFM partition and a maximum envy-free matching (X_L-saturating matching in $G[X_L; Y_L]$) can be computed in $O(|E|\sqrt{\min\{|X|, |Y|\}})$ time.

Based on Theorem 3, we derive the following lemma, which will be useful for designing algorithms in the remainder of this section.

Lemma 1. *For any envy-free allocation, all agents from X_S receive a bundle of utility 0 and all the allocated resources are from Y_L.*

Proof. Given any envy-free allocation $\boldsymbol{\pi}$, denote by $\mathcal{A}_z$ the set of agents receiving a bundle of utility 0 and by $\mathcal{A}_p$ the set of remaining agents (receiving a bundle of utility larger than 0). We construct a new allocation $\boldsymbol{\pi}'$ as follows. For each agent from $\mathcal{A}_z$, delete all resources from its bundle. For each agent from $\mathcal{A}_p$ keep an arbitrary resource in its bundle with utility 1 for the agent and delete the other resources. We show that $\boldsymbol{\pi}'$ is still envy-free. Since the original allocation $\boldsymbol{\pi}$ is envy-free and all agents from $\mathcal{A}_z$ receive a bundle of utility 0 under $\boldsymbol{\pi}$, it must be that every agent from $\mathcal{A}_z$ values every resource allocated under $\boldsymbol{\pi}$ as 0, and hence no agent from $\mathcal{A}_z$ will envy other agents under $\boldsymbol{\pi}'$. Moreover, under $\boldsymbol{\pi}'$, every agent from $\mathcal{A}_p$ receives a bundle of utility 1 and every agent gets exactly one resource, so no agent from $\mathcal{A}_p$ will be envious. Therefore, $\boldsymbol{\pi}'$ is envy-free. Since each agent either gets nothing or gets one resource liked by it under $\boldsymbol{\pi}'$, it induces an envy-free matching M in $G = (X\dot{\cup}Y, E)$. According to Theorem 3, we have $\mathcal{A}_p \subseteq X_L$. Since $\mathcal{A} = X_S \cup X_L = \mathcal{A}_z \cup \mathcal{A}_p$, we have $X_S \subseteq \mathcal{A}_z$, which means that all agents from X_S receive a bundle of utility 0. Since $\boldsymbol{\pi}$ is envy-free, it follows that all the allocated resources under $\boldsymbol{\pi}$ have utility 0 for agents from X_S. According to Theorem 3, each resource in Y_S is liked by at least one agent from X_S, so all the allocated resources are from Y_L. □

4.3 Social welfare and allocation size

Based on Lemma 1, we can design an FPT algorithm for usw-EF-PA. The idea is that according to Lemma 1, we just need to consider allocations restricted to X_L and Y_L. If $|X_L| \geq t$, then there is a trivial solution following from the envy-free matching. Otherwise, we can bound the size of the instance by a function depending only on t.

Theorem 4. *For 0/1-utilities* usw*-EF-PA is NP-hard and fixed-parameter tractable with respect to t. In particular, if $t = 1$, then* usw*-EF-PA is solvable in $O(n^{1.5}m)$ time for 0/1-utilities.*

Proof. The NP-hardness follows from the fact that usw-EF-PA for 0/1-utilities with t setting as the maximum utilitarian social welfare among all allocations is equivalent to the problem of deciding the existence of a Pareto efficient and envy-free allocation, which is NP-hard [8].

Next we show that usw-EF-PA for 0/1-utilities is fixed-parameter tractable with respect to t. According to Lemma 1, it suffices to check allocations that only allocate resources from Y_L. In addition, since in any desired allocation agents from X_S receive a bundle of utility 0, it suffices to check allocations that only allocate resources from Y_L to agents from X_L. If $X_L = \emptyset$, then no such allocations exists. In the following analysis we assume $X_L \neq \emptyset$. According to Theorem 3, there exists an envy-free matching M of cardinality $|X_L|$ in $G[X_L; Y_L]$. If $|X_L| \geq t$, then M induces an envy-free allocation with social welfare at least t and we are done. Otherwise, we have $|X_L| < t$. Since agents have binary utilities, we can partition all resources from Y_L into at most $2^{|X_L|} < 2^t$ groups according to the subset of agents from X_L who like the resource. If there is a group with more than t^2 resources, then allocating each agent from X_L a different set of t resources from this group is an envy-free allocation with social welfare $t|X_L| \geq t$ and we are done. Otherwise, we have $|Y_L| < 2^t t^2$ and then we can bound the number of all possible allocations restricted to X_L and Y_L by $O(2^{t^2} t^{2t})$. Thus, the problem is fixed-parameter tractable with respect to t.

When $t = 1$, it suffices to compute the EFM partition of G and check whether $|X_L| \geq 1$, so the running time is $O(n^{1.5}m)$ according to Theorem 3. □

Next, we provide an FPT algorithm for size-EF-PA using similar ideas. Here we just need to consider allocations restricted to Y_L and we will compare the size of X (instead of X_L) and t.

Theorem 5. *For 0/1-utilities* size*-EF-PA is NP-hard and is fixed-parameter tractable with respect to t. In particular, if $t = 1$, then* size*-EF-PA is solvable in $O(n^{1.5}m)$ time for 0/1-utilities.*

Proof. The NP-hardness follows from the fact that size-EF-PA for 0/1-utilities with $t = |\mathcal{R}|$ is equivalent to the problem of deciding the existence of a complete and envy-free allocation, which is NP-hard [18, 2].

Next we show that size-EF-PA for 0/1-utilities is fixed-parameter tractable with respect to t. According to Lemma 1, it suffices to check allocations that only allocate resources from Y_L. If $|Y_L| < t$, then there is no such allocation with size at least t. In the following analysis we assume $|Y_L| > t$. If $|X| \leq t$, then similar to the case for usw, we can bound the number of all possible allocations restricted to Y_L by $O(2^{t^2} t^{2t})$, and hence the problem is fixed-parameter tractable with respect to t. If $|X| > t$, then we can find an envy-free allocation with size at least t as follows. According to Theorem 3, there exists an envy-free matching M of cardinality $|X_L|$ in $G[X_L; Y_L]$, which induces an envy-free allocation $\boldsymbol{\pi}^M$. We extend $\boldsymbol{\pi}^M$ by letting each agent from X_S select a different resource from $Y_L \setminus Y_M$ until there is no remaining resource or each agent from X_S gets one resource. Denote the resulting allocation by $\boldsymbol{\pi}$. We have $\text{size}(\boldsymbol{\pi}) \geq \min\{|X|, |Y_L|\} \geq t$. According to Theorem 3, no resource from Y_L is liked by any agent from X_S, so $\boldsymbol{\pi}$ is still envy-free.

When $t = 1$, it suffices to compute the EFM partition of G and check whether $|Y_L| \geq 1$, so the running time is $O(n^{1.5}m)$ according to Theorem 3. □

4.4 Min-cardinality

Finally, we consider mcar-EF-PA. The following lemma reduces mcar-EF-PA with $t = 1$ to comparing the cardinality of X and Y_L

in the EFM partition of G.

Lemma 2. *The following three statements are equivalent:*

1. *There exists an envy-free allocation* $\boldsymbol{\pi}$ *where every agent gets a non-empty bundle, i.e.,* $\mathrm{mcar}(\boldsymbol{\pi}) \geq 1$*;*
2. *There exists an envy-free allocation* $\boldsymbol{\pi}$ *where every agent gets exactly one resource, i.e.,* $|\boldsymbol{\pi}(a)| = 1$ *for each* $a \in \mathcal{A}$*;*
3. $|X| \leq |Y_L|$.

Proof. (1) $\Leftarrow$ (2): If there exists an envy-free allocation $\boldsymbol{\pi}$ with $|\boldsymbol{\pi}(a)| = 1$ for each $a \in \mathcal{A}$, then clearly $\mathrm{mcar}(\boldsymbol{\pi}) \geq 1$.

(1) $\Rightarrow$ (2): Given an envy-free allocation $\boldsymbol{\pi}$ with $\mathrm{mcar}(\boldsymbol{\pi}) \geq 1$, denote by $\mathcal{A}_z$ the set of agents receiving a bundle of utility 0 and by $\mathcal{A}_p$ the set of remaining agents (receiving a bundle of utility larger than 0). For each agent from $\mathcal{A}_z$, keep an arbitrary resource in its bundle and delete the other resources. For each agent from $\mathcal{A}_p$, keep an arbitrary resource in its bundle with utility 1 for the agent and delete the other resources. Denote by $\boldsymbol{\pi}'$ the resulting allocation, where every agent gets exactly one resource. It remains to show that $\boldsymbol{\pi}'$ is envy-free. Since the original allocation $\boldsymbol{\pi}$ satisfies envy-freeness and all agents from $\mathcal{A}_z$ have utility 0 under $\boldsymbol{\pi}$, it must be that every agent from $\mathcal{A}_z$ values every resource allocated under $\boldsymbol{\pi}$ as 0, and hence no agent from $\mathcal{A}_z$ will envy other agents under $\boldsymbol{\pi}'$. Moreover, under $\boldsymbol{\pi}'$, since every agent from $\mathcal{A}_p$ has utility 1 and every agent gets exactly one resource, no agent from $\mathcal{A}_p$ will be envious. Thus, $\boldsymbol{\pi}'$ satisfies envy-freeness.

(2) $\Leftarrow$ (3): Suppose that $|X| \leq |Y_L|$. According to Theorem 3 we can find a X_L-saturating envy-free matching M in $G[X_L; Y_L]$, which induces an envy-free allocation $\boldsymbol{\pi}^M$, where every agent gets at most one resource. To get an envy-free allocation where every agent gets exactly one resource, we let each remaining agent corresponding to X_S select a different resource from $Y_L \setminus Y_M$. Since $|Y_L| \geq |X|$, there are enough remaining resources from $Y_L \setminus Y_M$. Denote the resulting allocation by $\boldsymbol{\pi}$, where every agent now gets exactly one resource. Since there are no edges between X_S and Y_L, all agents corresponding to X_S are non-envious. For agents corresponding to X_L, since they all have utility 1 and every agent gets exactly one resource, all of them are non-envious. Therefore, $\boldsymbol{\pi}$ is envy-free.

(2) $\Rightarrow$ (3): Let $\boldsymbol{\pi}$ be an envy-free allocation where every agent gets exactly one resource. According to Lemma 1, all the allocated resources are from Y_L. Thus, $|X| \leq |Y_L|$. □

It immediately follows that mcar-EF-PA with $t = 1$ is solvable in polynomial time. We prove the NP-hardness for the general case with arbitrary t in the following theorem. Whether the problem is fixed-parameter tractable with respect to t is left as an open question.

Theorem 6. *For* $0/1$*-utilities* mcar*-EF-PA is NP-hard. If* $t = 1$ *then it is solvable in* $O(n^{1.5}m)$ *time.*

Proof. We show the NP-hardness of mcar-EF-PA by providing a simple many-one reduction from size-EF-PA with $t = |\mathcal{R}|$, which is shown to be NP-hard in Theorem 5. Given an instance $(\mathcal{A}, \mathcal{R}, t = |\mathcal{R}|)$ of size-EF-PA, we create an instance $(\mathcal{A}, \mathcal{R}', t')$ of mcar-EF-PA, where $\mathcal{R}'$ contains all resources in $\mathcal{R}$ and also $t(|\mathcal{R}|-1)$ dummy resources that are not liked by any agent, and $t' = t$. It is easy to verify that there exists an envy-free and complete allocation for the former instance if and only if there exists an envy-free allocation such that every agent gets exactly t resources for the latter instance.

When $t = 1$, according to Lemma 2 and Theorem 3, it suffices to compute the EFM partition for G and check whether $|X| \leq |Y_L|$, so the running time is $O(n^{1.5}m)$. □

5 Ternary Valuations

We have seen that our problems are tractable for binary preferences and $t = 1$, which already has quite clear practical relevance as discussed in the introduction. A very natural question is whether these positive results transfer to three different utility values. In this section we answer this question negatively by showing strong NP-hardness for all the four goals under any three different utility values $\{0, v, u\}$ with $0 < v < u$. Before the proofs, let's introduce a known NP-Hard problem: Exact Cover by 3-Sets. [17]

EXACT COVER BY 3-SETS (X3C)
Input: A set X, with $|X| = 3n$ and a collection C of 3-element subsets of X.
Question: Is there a subset C' of C where every element of X occurs in exactly one member of C'?

We start by providing a reduction from esw to the other three problems for the above ternary utilities and $t = 1$.

Lemma 3. *Let* v *and* u *be two positive integers with* $0 < v < u$*. Let* $\mathcal{R}$ *be a set of resources,* $\mathcal{A}$ *be a set of agents, and* $(u_a)_{a\in\mathcal{A}}$ *be a collection of utility functions with* $u_a : \mathcal{R} \to \{0, u, v\}$*. Then, there exist extended sets of resources* $\mathcal{R}^* = \mathcal{R} \cup \mathcal{R}_{shadow}$ *and agents* $\mathcal{A}^* = \mathcal{A} \cup \mathcal{A}_{shadow}$*, and a collection of extended utility functions* $(u^*_a)_{a\in A^*}$ *(with* $u^*_a(r) = u_a(r)$ *for each* $a \in \mathcal{A}$ *and each* $r \in \mathcal{R}$*) such that:*

Regarding (u_a) *there exists an envy-free allocation* $\pi^{\mathrm{esw}} : \mathcal{A} \to 2^{\mathcal{R}}$ *with* $\mathrm{esw}(\pi^{\mathrm{esw}}) \geq 1$*, if and only if regarding* (u^*_a) *there exists an envy-free allocation* $\pi^* : \mathcal{A}^* \to 2^{\mathcal{R}^*}$ *with* $\mathcal{E}(\pi^*) \geq 1$ *for each* $\mathcal{E} \in \{\mathrm{mcar}, \mathrm{usw}, \mathrm{size}\}$[3]*. Moreover,* $(\mathcal{R}^*, \mathcal{A}^*, (u^*_a)_{a\in\mathcal{A}^*})$ *can be computed in linear time.*

Proof. Given $(\mathcal{R}, \mathcal{A}, (u_a)_{a\in\mathcal{A}})$, we construct $(\mathcal{R}^* = \mathcal{R} \cup \mathcal{R}_{\text{shadow}}, \mathcal{A}^* = \mathcal{A} \cup \mathcal{A}_{\text{shadow}}, (u^*_a)_{a\in A^*})$ as follows. For each resource, we create two corresponding shadow agents and two corresponding shadow resources. That is, $\mathcal{A}_{\text{shadow}} := \{a'_r, a''_r \mid r \in R\}$ and $\mathcal{R}_{\text{shadow}} := \{r', r'' \mid r \in \mathcal{R}\}$. We distinguish between *original agents* $\mathcal{A}$ and *shadow agents* $\mathcal{A}_{\text{shadow}}$, as well as between *original resources* $\mathcal{R}$ and *shadow resources* $\mathcal{R}_{\text{shadow}}$. The idea is to define utilities functions $(u^*_a)_{a\in A^*}$ such that whenever any agent gets a resource, each shadow agent will also require a shadow resource, which in turn ensures that every agent gets a resource of positive value. Formally, $(u^*_a)_{a\in A^*}$ is defined as follows (see also Table 2).

- For each original agent a and each original resource r, u^* is identical to u, i,e., $u^*_a(r) = u_a(r)$.
- Each original agent is interested in all the shadow resources and values each of them as v.
- Each shadow agent is interested in all the shadow resources and value each of them as u.
- Each shadow agent a'_r or $a''_r \in \mathcal{A}^*_{\text{shadow}}$ is also interested in its unique corresponding original resource $r \in \mathcal{R}$, i.e., $u^*_{a'_r}(r) = u^*_{a''_r}(r) = v$, and values all other original resources as 0.

	$\bar{r} \in \mathcal{R} \setminus \{r\}$	$r \in \mathcal{R}$	$r^* \in \mathcal{R}_{\text{shadow}}$
$a \in \mathcal{A}$	$u_a(\bar{r})$	$u_a(r)$	v
$a'_r, a''_r \in \mathcal{A}_{\text{shadow}}$	0	v	u

Table 2. Utility functions of the agents in the proof of the Lemma 3

Next, we show that for $(\mathcal{R}^*, \mathcal{A}^*, (u^*_a)_{a\in\mathcal{A}^*})$ and any $\mathcal{E}, \mathcal{E}' \in \{\mathrm{mcar}, \mathrm{usw}, \mathrm{size}\}$ it holds that for every envy-free allocation π with

[3] Note that given any $\pi^{\mathcal{E}}$ for $\mathcal{E} \in \{\mathrm{esw}, \mathrm{mcar}, \mathrm{usw}, \mathrm{size}\}$, we can compute each of the respective other allocations in polynomial time.

$\mathcal{E}(\pi) \geq 1$ we also have $\mathcal{E}'(\pi) \geq 1$. By definition, it is obvious that an envy-free allocation π with $\mathrm{mcar}(\pi) \geq 1$ or $\mathrm{usw}(\pi) \geq 1$ must in both cases have $\mathrm{size}(\pi) \geq 1$. Let us conversely assume that there exists some envy-free allocation π with $\mathrm{size}(\pi) \geq 1$ for $(\mathcal{R}^*, \mathcal{A}^*, (u_a^*)_{a\in\mathcal{A}^*})$. We want to show that $\mathrm{mcar}(\pi) \geq 1$ and $\mathrm{usw}(\pi) \geq 1$ also hold for $(\mathcal{R}^*, \mathcal{A}^*, (u_a^*)_{a\in\mathcal{A}^*})$. Since $\mathrm{size}(\pi) \geq 1$, at least one resource r is allocated. If r is not a shadow resource, then at least one of the two corresponding shadow agents a_r' or a_r'' gets a shadow resource. Thus, at least one shadow resource is allocated under π. Considering that each shadow agent can only gain a maximum value of v from the shadow resources, and $u > v$, the fact that at least one shadow resource is allocated under π makes every shadow agent require at least one shadow resource with value at least u. Since $|\mathcal{A}_{\text{shadow}}| = |\mathcal{R}_{\text{shadow}}| = 2|\mathcal{R}|$, each shadow agent should receive exactly one shadow resource. Since each original agent values each shadow resource as v, this enforces that each original agent gets a bundle with value at least v. Therefore, we have $\mathrm{mcar}(\pi) \geq 1$ and $\mathrm{usw}(\pi) \geq 1$.

To prove the lemma, it remains to show that there exists an envy-free allocation π^{esw} with $\mathrm{esw}(\pi) \geq 1$ for $(\mathcal{R}, \mathcal{A}, (u_a)_{a\in\mathcal{A}})$ if and only if there exists an envy-free allocation π^{size} with $\mathcal{E}(\pi^{\mathrm{size}}) \geq 1$ for $(\mathcal{R}^*, \mathcal{A}^*, (u_a^*)_{a\in\mathcal{A}^*})$.

($\Longrightarrow$) Assume there exists an envy-free allocation π^{esw} with $\mathrm{esw}(\pi^{\mathrm{esw}}) \geq 1$ for $(\mathcal{R}, \mathcal{A}, (u_a)_{a\in\mathcal{A}})$. A desired allocation π^{size} for $(\mathcal{R}^*, \mathcal{A}^*, (u_a^*)_{a\in\mathcal{A}^*})$ can be constructed in the following way. Analogously to π^{esw}, we let $\pi_a^{\mathrm{size}} = \pi_a^{\mathrm{esw}}$ for each original agent $a \in \mathcal{A}$. Aside from that, each shadow agent is assigned an arbitrary shadow resource. Clearly, original agents will not envy each other, and each of them receives a bundle with positive value v or u. Consequently, original agents will not envy shadow agents either, since they perceive the value of each shadow agent's bundle to be exactly v. Meanwhile, shadow agents will not envy original agents because, in their views, the value of each shadow agent's bundle is u, whereas the value of any original agent's bundle does not exceed v.

($\Longleftarrow$) Assume there exists some envy-free allocation π^{size} with $\mathrm{size}(\pi^{\mathrm{size}}) \geq 1$ for $(\mathcal{R}^*, \mathcal{A}^*, (u_a^*)_{a\in\mathcal{A}^*})$. Recall that in π^{size}, each shadow agent must get exactly one shadow resource, and each original agent must get a bundle with positive value. Thus, we have $\mathrm{esw}(\pi^{\mathrm{size}}) \geq 1$. We create an allocation π^{esw} for $(\mathcal{R}, \mathcal{A}, (u_a)_{a\in\mathcal{A}})$ in a straight-forward way by setting $\pi_a^{\mathrm{esw}} := \pi_a^{\mathrm{size}}$ for each original agent $a \in A$. Note that this is indeed a well-defined allocation for $(\mathcal{R}, \mathcal{A}, (u_a)_{a\in\mathcal{A}})$ since π^{size} allocates shadow resources only to shadow agents. Since the original agents do not envy each another in π^{size} for $(\mathcal{R}^*, \mathcal{A}^*, (u_a^*)_{a\in\mathcal{A}^*})$, and the utility functions of the original agents for original resources are identical for $(\mathcal{R}^*, \mathcal{A}^*, (u_a^*)_{a\in\mathcal{A}^*})$ and $(\mathcal{R}, \mathcal{A}, (u_a)_{a\in\mathcal{A}})$, it follows that π^{esw} is envy-free for $(\mathcal{R}, \mathcal{A}, (u_a)_{a\in\mathcal{A}})$. □

According to Lemma 3, if we show that esw-EF-PA is strongly NP-hard for ternary utility values $0 < v < u$, then we will get the strong NP-hardness of $\mathcal{E}$-EF-PA for each $\mathcal{E} \in \{\mathrm{mcar}, \mathrm{usw}, \mathrm{size}\}$ for free. Our main result in this section is that all the four goals are strongly NP-hard for ternary utility values $0 < v < u$ even if $t = 1$, stated as follows.

Theorem 7. *For each $\mathcal{E} \in \{\mathrm{esw}, \mathrm{mcar}, \mathrm{usw}, \mathrm{size}\}$, $\mathcal{E}$-EF-PA is strongly NP-hard, even if each agent assigns to each resource only values from $\{0, u, v\}$ with $v, u \in \mathbb{N}, 0 < v < u$, and $t = 1$.*

By Lemma 3, it suffices to show the strong NP-hardness for esw-EF-PA. To this end, we make a case distinction over the values of u and v and show all cases via different reductions from the NP-hard EXACT COVER BY 3-SETS (X3C) problem [17]. Given a multiset $X = \{x_1, x_2, \ldots, x_{3n}\}$, with $|X| = 3n$, and a collection $C = \{S_1, S_2, \ldots, S_m\}$ of 3-element subsets of X, X3C asks whether there is some $C' \subseteq C$ where every element of X occurs in exactly one member of C'. We assume without loss of generality that $m > 3n$. If this requirement is not fulfilled then we can easily obtain an equivalent instance which fulfills it by adding some dummy 3-sets. Let's start with $u = kv, k \geq 3$.

Lemma 4 (★). esw*-EF-PA with ternary utility values* $\{0, u, v\}$, $u = kv > 0$, $k \geq 3$, *and* $t = 1$ *is strongly NP-hard.*

Next we consider the case with $u = 2v$. The distinctive feature of the following proof, compared to the previous one, lies in our creation of standard agents and special resources as benchmarks, ensuring that the value of the bundle desired by each agent exceeds a certain constant value. Additionally, we introduce a large number of observers and corresponding blank resources to monitor potential combinations of resources that may interfere with the reduction.

Lemma 5. esw*-EF-PA with ternary utility values* $\{0, u, v\}$, $u = 2v > 0$, *and* $t = 1$ *is strongly NP-hard.*

Proof. The hardness proof also proceeds by a reduction from X3C. Given an instance (X, C) of the X3C, we construct an instance of $\mathcal{I} = (\mathcal{R}, \mathcal{A}, (u_a)_{a\in\mathcal{A}}, t = 1)$ of the esw-EF-PA problem as follows.

- There are m *cover agents* $\mathcal{A}_C = \{a_1, a_2, \ldots, a_m\}$, 3 *standard agents* b, c, d, and a set $\mathcal{W}$ of *observers* (of finite size to be specified later), i.e., $\mathcal{A} = \mathcal{A}_C \cup \{b, c, d\} \cup \mathcal{W}$.
- There $3n$ *normal resources* $\mathcal{R}_N = \{r_1, r_2, \ldots, r_{3n}\}$, n *small resources* $\mathcal{R}_S = \{s_1, s_2, \ldots, s_n\}$, $2(m-n)$ *dummy resources* $\mathcal{R}_D = \{d_1, \ldots, d_{2(m-n)}\}$ and a finite number of *blank resources* $\mathcal{R}_B$ (where $|\mathcal{R}_B| = 2|\mathcal{W}|$) and 4 *special resources* $r_1^*, r_2^*, r_3^*, r_4^*$, i.e, $\mathcal{R} = \mathcal{R}_N \cup \mathcal{R}_S \cup \mathcal{R}_D \cup \mathcal{R}_B \cup \{r_1^*, r_2^*, r_3^*, r_4^*\}$.
- For each cover agent a_j and each normal resource r_i, the utility function is defined such that $u_j(r_i) = v$ if $x_i \in S_j$ and $u_j(r_i) = 0$ otherwise. Besides, each cover agent values each small resource as v. In addition, each cover agent values each dummy resource and each special resource as $2v$. The cover agents are not interested in blank resources.
- For each standard agent and each special resource, the utility function is defined as described in the table below and the standard agents are not interested in any of the other resources:

	r_1^*	r_2^*	r_3^*	r_4^*
b	$2v$	0	$2v$	0
c	$2v$	v	$2v$	v
d	0	0	$2v$	0

Table 3. Utility functions of the agents in the proof of the Lemma 5

- Each observer assigns value 2 to each blank resource and each special resource. In particular, there are three different kinds of observers where we only describe which other resources have nonzero value: (1) Observers $w_{i,j;k}$ of type 1: An observer $w_{i,j;k}$ values the two normal resources r_i, r_j, and one dummy resource d_k at $2v$, respectively. (2) Observers $w'_{i;j;k}$ of type 2: An observer $w'_{i;j;k}$ values the normal resource r_i and the dummy resource d_j and the small resource s_k at $2v$, respectively. (3) An observer w^* values every small resource and every dummy resource at $2v$.

Overall, we create $\binom{3n}{2} * m + 3m * n * 2(m-n) + 1$ observers. Thus, there are $\mathcal{O}(n^2 m)$ numbers of observers and blank resources.

Assuming that there is solution for the constructed instance $\mathcal{I}$ of esw-EF-PA, we have the following observations.

Ob. 1. We can take a look at standard agents at first. Since each agent has to get positive value, the standard agent d will get r_3^*. Then, the standard agent b will get r_1^* and the standard agent c will get r_2^* and r_4^*. Since c gets r_2^* and r_4^*, each of the cover agents and the observers will require a value of at least $4v$.

Ob. 2. The normal resources, dummy resources and small resources can only be allocated to the m cover agents. This is because cover agents are not interested in blank resources and the sum of the value that these three kinds of resources can provide is at most $4mv$.

Ob. 3. It follows that each cover agent gets utility exactly $4v$. Thus, if a cover agent gets 2 dummy resources, he cannot get any other resources.

Ob. 4. Since the observers can only get blank resources. Each observer will get exactly two blank resources.

Ob. 5. From the previous observations 1–4, we can claim that each resource is allocated to one agent in this allocation.

Ob. 6. No cover agent can get three different kinds of resources. Otherwise, some observer $w'_{i;j;k}$ of type 2 will envy.

Ob. 7. No cover agent can get one dummy resource and one small resource. Assume this was the case, then this agent needs another resource to ensure the bundle is of value at least $4v$. However, this resource cannot be a normal resource according to observation 6, and it cannot be a small or dummy resource since otherwise observer w^* would envy this agent.

Ob. 8. No cover agent receives one dummy resource and one normal resource. Again, another resource would be needed, which cannot be a dummy resource because of observers of types 2 and cannot be a normal resource because of observers of type 1.

Ob. 9. From the previous observations 3,6–8, we can claim that if some cover agent gets a dummy resource then it will get exactly two dummy resources and nothing else. Thus, there are $m-n$ cover agents who get $2(m-n)$ two dummy resources.

Ob. 10. The remaining n cover agents get normal resources and small resources and each of them get exactly 1 small resource. This is because the value that each of them can get from normal resources is at most $3v$. According to the pigeonhole principle, there is and can only be one small resource for each cover agent.

We show that (X, C) is a YES-instance if and only if $\mathcal{I}$ is a YES-instance.

($\Longrightarrow$) Since (X, C) is a YES-instance, there exists a subset $C' \subseteq C$ with $|C'| = n$ such that every element of X occurs in exactly one member of C'. If $S_j \in C'$, we allocate the 3 corresponding normal resources to a_j such that the value that a_j can get is exactly $3v$ for now. In addition, he will also get 1 small resource and finally get the value $4v$. If $S_j \notin C'$, we allocate 2 dummy resources and the value is also $4v$. In addition, each observer gets 2 blank resources and the value is also $4v$. Finally, b gets r_1^*, c gets r_2^* and r_4^*, d gets r_3^*. In this case, no one envies the other. Thus, I is also a YES-instance.

($\Longleftarrow$) Since I is a YES-instance, combining the observations above, note that there are n agents a_j who only get three normal resources $I_j = \{i_{ja}, i_{jb}, i_{jc}\}$ and one small resource such that we can find n corresponding set $S_j = \{x_{ja}, x_{jb}, x_{jc}\}$. We can find exactly n such disjoint sets. This induces a feasible solution C'. Thus, (X, C) is a YES-instance. □

Finally, we consider the case when u is not divisible by v. The following proof, while similar to the previous one, involves additional considerations. These arise primarily because u may be significantly greater than v. For some agents, in order to achieve a value exceeding u or even $2u$ solely through resources valued at v, they would need to acquire a multiple of these resources.

Lemma 6 (★). esw-*EF-PA with ternary utility values* $\{0, v, u = kv + c\}$ *for* $v > c > 0, k > 0$ *and* $t = 1$ *is strongly NP-hard.*

Now combining Lemma 3 to 6 we get the claim in the Theorem 7.

6 Conclusions

We studied how to allocate indivisible resources to agents in an envy-free manner by relaxing the common requirement that all resources should be allocated. We considered envy-free partial allocations that can provide at least some utilities or allocate some resources from both systematic or individual perspectives, and we obtained comprehensive results under various classes of utilities. While all problems we considered are NP-hard under identical utilities, we identified many tractable results for binary utilities and showed interesting connections to matching problems on bipartite graphs. For ternary utilities, it is somewhat surprising that all problems we studied are NP-hard even if we only require the bare minimum of efficiency or cardinality. Our results provide a more fine-grained view of the computational complexity of finding efficient envy-free allocations.

Our work can be extended in several directions. First, our results showing that some cases are tractable under binary utilities but all scenarios become NP-hard under ternary utilities form a stark contrast worth of being further investigated. In particular, it is interesting to study bivalued utilities, with utilities taking one of two values but not necessarily 0 and 1, that lie between binary and ternary utilities. In the full version of the paper, we provide some initial results for $1/2$ utilities, where we show that when $t = 1$ all the four efficiency measures are equivalent, and we can reduce the problem to the case where each agent can get at most two resources. Second, so far we assumed all resources are goods. A natural extension is to study chores or mixed resources. Note that for chores it is perhaps even more meaningful to consider the efficiency measures size and mcar since the planner might want to distribute as many tasks to agents as possible. Here a relevant result is that for chores and binary values (or even binary marginals), there always exists an envy-free allocation with at most $n - 1$ unallocated resources [21]. Finally, one can also study equitability instead of envy-freeness in our setting. We note that for identical utilities, these two fairness notions are equivalent. For binary utilities, the problems seem to be easier for equitability since all agents have to get identical utilities. For example, it is easy to show that maximizing the utilitarian or egalitarian welfare, for any t, can be reduced to the maximum cardinality matching problem (similarly as done in Theorem 2).

References

[1] E. Aigner-Horev and E. Segal-Halevi. Envy-free matchings in bipartite graphs and their applications to fair division. *Information Sciences*, 587: 164–187, 2022.

[2] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227: 71–92, 2015.

[3] H. Aziz, I. Schlotter, and T. Walsh. Control of fair division. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, pages 67–73, 2016.

[4] H. Aziz, X. Huang, N. Mattei, and E. Segal-Halevi. Computing welfare-maximizing fair allocations of indivisible goods. *European Journal of Operational Research*, 307(2):773–784, 2023.

[5] B. Berger, A. Cohen, M. Feldman, and A. Fiat. Almost full EFX exists for four agents. In *Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI-22)*, pages 4826–4833, 2022.

[6] A. Beynier, Y. Chevaleyre, L. Gourvès, A. Harutyunyan, J. Lesca, N. Maudet, and A. Wilczynski. Local envy-freeness in house allocation problems. *Autonomous Agents and Multi-Agent Systems*, 33(5): 591–627, 2019.

[7] N. Boehmer, R. Bredereck, K. Heeger, D. Knop, and J. Luo. Multivariate algorithmics for eliminating envy by donating goods. In *Proceedings of the Twenty-First International Conference on Autonomous Agents and Multiagent Systems (AAMAS-22)*, pages 127–135, 2022.

[8] S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *Journal of Artificial Intelligence Research*, 32(1):525–564, 2008.

[9] S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair allocation of indivisible goods. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12. Cambridge University Press, 2016.

[10] X. Bu, Z. Li, S. Liu, J. Song, and B. Tao. On the complexity of maximizing social welfare within fair allocations of indivisible goods. *CoRR*, abs/2205.14296, 2022.

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

[12] I. Caragiannis, N. Gravin, and X. Huang. Envy-freeness up to any item with high Nash welfare: The virtue of donating items. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC-19)*, pages 527–545, 2019.

[13] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3):12:1–12:32, 2019.

[14] B. R. Chaudhury, T. Kavitha, K. Mehlhorn, and A. Sgouritsa. A little charity guarantees almost envy-freeness. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, (SODA-20)*, pages 2658–2672, 2020.

[15] B. Dorn, R. de Haan, and I. Schlotter. Obtaining a proportional allocation by deleting items. *Algorithmica*, 83(5):1559–1603, 2021.

[16] J. Gan, W. Suksompong, and A. A. Voudouris. Envy-freeness in house allocation problems. *Mathematical Social Sciences*, 101:104–106, 2019.

[17] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[18] H. Hosseini, S. Sikdar, R. Vaish, H. Wang, and L. Xia. Fair division through information withholding. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*, pages 2014–2021, 2020.

[19] N. Kamiyama, P. Manurangsi, and W. Suksompong. On the complexity of fair house allocation. *Operations Research Letters*, 49(4):572–577, 2021.

[20] L. Ramshaw and R. E. Tarjan. On minimum-cost assignments in unbalanced bipartite graphs. *technical report*, 2012.

[21] B. Tao, X. Wu, Z. Yu, and S. Zhou. On the existence of efx (and pareto-optimal) allocations for binary chores. *arXiv preprint arXiv:2308.12177*, 2023.

[22] T. Walsh. Challenges in resource and cost allocation. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, pages 4073–4077, 2015.