Computing Efficient Envy-Free Partial Allocations of Indivisible Goods

Robert Bredereck, Andrzej Kaczmarczyk, Junjie Luo, Bin Sun

Abstract

Envy-freeness is one of the most prominent fairness concepts in the allocation of indivisible goods. Even though trivial envy-free allocations always exist, rich literature shows this is not true when one additionally requires some efficiency concept (e.g., completeness, Pareto-efficiency, or social welfare maximization). In fact, in such case even deciding the existence of an efficient envy-free allocation is notoriously computationally hard. In this paper, we explore the limits of efficient computability by relaxing standard efficiency concepts and analyzing how this impacts the computational complexity of the respective problems. Specifically, we allow partial allocations (where not all goods are allocated) and impose only very mild efficiency constraints, such as ensuring each agent receives a bundle with positive utility. Surprisingly, even such seemingly weak efficiency requirements lead to a diverse computational complexity landscape. We identify several polynomial-time solvable or fixed-parameter tractable cases for binary utilities, yet we also find NP-hardness in very restricted scenarios involving ternary utilities.

1 Introduction

Computing fair allocations of indivisible resources is an important issue with many applications in all kinds of disciplines [11, 12, 27]. 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 [12]. Unfortunately, non-trivial envy-free allocations do not always exist, and computing them is often associated to computationally very difficult problems [11]. In consequence, researchers have developed several ways to relax that fairness notion, such as envy-free up to one good (EF1) [14] and envy-free up to any good (EFX) [16].

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. 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 (in comparison to 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 non-negative utilities for them. In this case, all our efficiency concepts for $t = 1$ are significantly less demanding than multiple other prominent efficiency concepts, such as completeness, as demonstrated by the earlier folklore example. Hence, analyzing computational complexity 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.

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 [16]. 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. [15] showed that there always exists an EFX partial allocation with at least half of the maximum Nash welfare. Chaudhury et al. [17] showed that donating at most $n - 1$ resources can guarantee the existence of EFX allocation such that no agent prefers the donated resources to their 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 [7]. Besides existence, Bu et al. [13] 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 to identify polynomial-time solvable cases.

Boehmer et al. [10] 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. [18, Chap 5] 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. [23] 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

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. For both FPT entries in the table, the parameterization is by t.

	Identical	Binary		Ternary
	$t = 1$	$t = 1$	t	$t = 1$
utilitarian social welfare (usw)	NP-h (Th. 1)	P (Th. 4)	NP-h, FPT (Th. 4)	NP-h (Th. 7)
egalitarian social welfare (esw)		P (Th. 2)	P (Th. 2)	
#resources allocated (size)		P (Th. 5)	NP-h, FPT (Th. 5)	
min-cardinality (mcar)		P (Th. 6)	NP-h (Th. 6)	

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 [20, 8, 24] 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 matchings by Aigner-Horev and Segal-Halevi [1].

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. An overview of our results is provided in Table 1. In Section 3, we show that finding such allocations is strongly NP-hard, even if all agents have identical preferences. In Section 4, we focus on the case with binary utilities, where each agent values a resource as either 0 or 1. We show that all the four variants are polynomial-time solvable when $t = 1$, indicating that determining the existence of envy-free allocations with minimal efficiency requirements can be done efficiently. For arbitrary t, while most problem variants become strongly NP-hard, we show that the utilitarian welfare variant and the number of allocated resources variant are both fixed-parameter tractable (FPT)[1] with respect to t, implying that the problems can still be efficiently solved for small t. 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). In Section 5, we go beyond binary preferences and allow for three different utility values. We show a reduction from the egalitarian welfare variant to the other three variants for ternary utilities and $t = 1$, which reveals an interesting connection between the four efficiency requirements and might be of independent interest. Based on this reduction, we show that

[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.

all variants become strongly NP-hard already when $t = 1$ for any ternary utility values $\{0, v, u\}$ with $0 < v < u$. Furthermore, all the problems shown to be NP-hard in this paper are contained in NP, since verifying that an allocation (guessed non-deterministically) is envy-free and meets the respective efficiency criterion is possible in polynomial time.

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 their 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 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 agents' utilities use only values 0 or 1 and *ternary* when there are three possible values of utility that agents can report.

An *allocation* $\boldsymbol{\pi} : \mathcal{A} \to 2^{\mathcal{R}}$ assigns each agent $a \in \mathcal{A}$ their private *bundle* $\boldsymbol{\pi}(a)$, 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 $\mathrm{mcar}(\boldsymbol{\pi}) \coloneqq \min_{a \in \mathcal{A}} |\boldsymbol{\pi}(a)|$ of resources allocated to some agent the *min-cardinality* of $\boldsymbol{\pi}$, whereas by $\mathrm{size}(\boldsymbol{\pi}) \coloneqq \sum_{a \in \mathcal{A}} |\boldsymbol{\pi}(a)|$ 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., $\mathrm{usw}(\boldsymbol{\pi}) \coloneqq \sum_{a \in \mathcal{A}} u_a(\boldsymbol{\pi}(a))$. Analogously, *egalitarian social welfare* $\mathrm{esw}(\boldsymbol{\pi})$ is the minimum of the agent's utilities, i.e., $\mathrm{esw}(\boldsymbol{\pi}) \coloneqq \min_{a \in \mathcal{A}} u_a(\boldsymbol{\pi}(a))$. (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 for all studied efficiency concepts with $t = 1$ via a reduction from the 3-partition problem [21]. 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.

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

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.*

Proof. The hardness proof proceeds by a reduction from the 3-PARTITION problem [21]. Given a multiset of positive numbers $N = \{e_1, e_2, \ldots, e_{3n}\}$, in 3-PARTITION we ask whether N can be partitioned into n three-element parts such that each of them sums up to $b := \sum_{e_i \in N} e_i/n$. Given an instance $\mathcal{I} = (N)$ of 3-PARTITION, we build an instance $\mathcal{I}' = (\mathcal{R}, \mathcal{A}, (u_a), t)$ of $\mathcal{E}$-EF-PA, with $t = 1$ and with identical utilities, as follows:

- We put $3n + 1$ *agents* a_0, a_1, a_2, …, a_{3n} in $\mathcal{A}$;
- We construct $3n$ *normal resources* $R_N = \{r_1, r_2, \ldots, r_{3n}\}$ and a set $\mathcal{R}_S$ of *special resources* with $|\mathcal{R}_S| = 2n + 1$, i.e., $\mathcal{R} = \mathcal{R}_N \cup \mathcal{R}_S$;
- We let each agent $a \in \mathcal{A}$ have the same *utility function* u such that $u(r_i) = e_i + b$ for each resource $r_i \in \mathcal{R}_N$ and $u(r^*) = 4b$ for each special resource $r^* \in \mathcal{R}_S$.

First we describe a structural property of every solution $\boldsymbol{\pi}$ to $(\mathcal{R}, \mathcal{A}, (u_a))$. Note that it is clear that $\boldsymbol{\pi}$ cannot be an empty allocation, as this would violate the threshold value $t = 1$ for each of efficiency concepts $\mathcal{E} \in \{\mathrm{usw}, \mathrm{esw}, \mathrm{size}, \mathrm{mcar}\}$. Further, since there are $3n + 1$ agents and only $3n$ normal resources, one agent has to get some special resource $r^* \in \mathcal{R}_S$ consequently obtaining utility at least $4b$. Otherwise, this agent would be envious. Hence, due to identical utilities, each agent has to get utility at least $4b$ from their bundles. Finally, since the sum of the values of all resources is exactly $(3n + 1) \cdot 4b$, it follows that $\boldsymbol{\pi}$ assigns to each agent a bundle of value exactly $4b$. Given that, it is easy to see that each agent that gets a special resource, does not have any other resource assigned by $\boldsymbol{\pi}$. Thus, in $\boldsymbol{\pi}$ there are exactly $2n + 1$ agents that get one special resource. The remaining n agents get exactly three normal resources out of $3n$ of them. In any other case, because $e_i < b$ for each $e_i \in N$, there would be at least one agent getting a utility smaller than $4b$ from their bundle.

We now show that the original instance $\mathcal{I}$ is a YES-instance if and only if the constructed one, $\mathcal{I}'$, is a YES-instance. Our argument is solely based on the above-analyzed structure of solutions to $\mathcal{I}'$. Since the structure applies to each studied efficiency concept when $t = 1$, the argument works for each choice of the efficiency concept.

($\Longrightarrow$) Suppose $\mathcal{I}$ is a YES-instance of 3-PARTITION. Then, there are n disjoint subsets, each summing up to b. We denote each subset $N_j = \{e_{x(j)}, e_{y(j)}, e_{z(j)}\}$ for $j \in [n]$. We build allocation $\boldsymbol{\pi}$ which certifies that $\mathcal{I}'$ is a YES-instance. First, $\boldsymbol{\pi}$ allocates the three resources $r_{x(j)}, r_{y(j)}, r_{z(j)}$ to each agent a_j, $j \in [n]$. Then, it gives the remaining $2n + 1$ special resources to the remaining $2n + 1$ agents, one resource per agent. It is easy to verify that each agents gets utility exactly $4b$ from their bundle, which shows that $\mathcal{I}'$ is indeed a YES-instance.

($\Longleftarrow$) Let us assume that $\mathcal{I}'$ is a YES-instance. Due to the structure of solutions to $\mathcal{I}'$ discussed earlier, we know that there is an allocation $\boldsymbol{\pi}$ letting each agent have utility exactly $4b$. The structure also requires that $\boldsymbol{\pi}$ assigns one special resource to $2n + 1$ agents (one per agent) and that each of the remaining n agents gets exactly three resources whose utility is $4b$. Without loss of generality, we label the latter agents a_1, a_2, …, a_n and denote the resources of agent a_j, $j \in [n]$, by $r_{x(j)}$, $r_{y(j)}$, and $r_{z(j)}$ for $j \in [n]$. It follows that we can split N into n subsets constructing a subset N_j of numbers $\{e_{x(j)}, e_{y(j)}, e_{z(j)}\}$ for each $j \in [n]$. By our construction, the sum $e_{x(j)} + e_{y(j)} + e_{z(j)}$ of numbers in each such subset N_j is b, which proves that $\mathcal{I}$ is a YES-instance.

Clearly, the reduction is computable in polynomial time.

$\square$

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 focus on other aspects to identify 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. Specifically, we study binary and ternary utilities, which are commonly studied utility restrictions in the literature [6, 22, 5, 19].

4 Binary utilities

Given that identifying exact utility values imposes a high cognitive burden for human agents, in practice binary utilities, where agents express preferences by pointing out which resources they desire and which not, 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, 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 [25], 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 is 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:*

- *An X_L-saturating matching in $G[X_L; Y_L]$ always exists, and every X_L-saturating matching in $G[X_L; Y_L]$ is an envy-free matching in G;*
- *Every envy-free matching in G is contained in $G[X_L; Y_L]$;*
- *There are no edges between X_S and Y_L;*
- *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 Thm. 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 in 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 their bundle. For each agent from $\mathcal{A}_p$ keep an arbitrary resource in their 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 Thm. 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 Thm. 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. $\square$

Based on the properties of envy-free matchings, we prove in the following that both $0/1$-utilities usw-EF-PA and $0/1$-utilities size-EF-PA admit FPT algorithms parameterized by t, and that $0/1$-mcar-EF-PA is in P regardless of the choice of t.

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, it suffices to consider allocations restricted to X_L and Y_L. If $|X_L| \geq t$, 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. Hardness follows from the equivalence of usw-EF-PA for $0/1$-utilities with t setting as the maximum utilitarian social welfare among all allocations and the NP-hard problem of deciding the existence of a Pareto efficient and envy-free allocation [11], since Bliem et al. [9, Ob.1] shows that, in case of $0/1$-utilities, an allocation is Pareto-efficient if and only if it is complete and every resource is allocated to an agent that assigns 1 to it.

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 Thm. 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. This is because resources with zero utility for all agents are irrelevant and can be removed during preprocessing. Thus, every resource has a positive value for at least one agent. 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 for 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 Thm. 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. Hardness follows from size-EF-PA for $0/1$-utilities with $t = |\mathcal{R}|$ being equivalent to the problem of deciding the existence of a complete and envy-free allocation, which is NP-hard [23, 2].

Next we show that size-EF-PA for $0/1$-utilities is fixed-parameter tractable with respect to t. By 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 Thm. 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 $\mathrm{size}(\boldsymbol{\pi}) \geq \min\{|X|, |Y_L|\} \geq t$. According to Thm. 3, no resource from Y_L is liked by any agent from X_S, so $\boldsymbol{\pi}$ is still envy-free.

For $t = 1$, computing the EFM partition of G and checking whether $|Y_L| \geq 1$ suffices; so Thm. 3 yields running time $O(n^{1.5}m)$. □

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 their bundle and delete the other resources. For each agent from $\mathcal{A}_p$, keep an arbitrary resource in their 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 Thm. 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 subsequently prove the NP-hardness for the general case with arbitrary t. However, whether the problem is fixed-parameter tractable with respect to t remains open.

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 Thm. 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 $t(|\mathcal{A}| - 1)$ 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 Thm. 3, it suffices to compute the EFM partition for G and check whether $|X| \leq |Y_L|$, so the running time is $O(n^{1.5}m)$. □

Table 2: Agent's utility functions in the proof of Lemma 3.

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

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 $0 < v < u$.

We start by providing a very general reduction from esw to the other three problems for any ternary utilities which include utility zero 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 $u_a : \mathcal{R} \to \{0, v, u\}$. 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 $\pi^{\text{esw}} : \mathcal{A} \to 2^{\mathcal{R}}$ with $\text{esw}(\pi^{\text{esw}}) \geq 1$, 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 $\mathcal{E} \in \{\text{mcar}, \text{usw}, \text{size}\}$[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 $(\mathcal{R}^* = \mathcal{R} \cup \mathcal{R}_{\text{shadow}}, \mathcal{A}^* = \mathcal{A} \cup \mathcal{A}_{\text{shadow}}, (u^*_a)_{a\in\mathcal{A}^*})$ 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 values each of them as u.
- Each shadow agent a'_r or $a''_r \in \mathcal{A}^*_{\text{shadow}}$ is also interested in their 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.

Next, we show that for $(\mathcal{R}^*, \mathcal{A}^*, (u^*_a)_{a\in\mathcal{A}^*})$ and any $\mathcal{E}, \mathcal{E}' \in \{\text{mcar}, \text{usw}, \text{size}\}$ it holds that for every envy-free allocation π with $\mathcal{E}(\pi) \geq 1$ we also have $\mathcal{E}'(\pi) \geq 1$. By definition, it is obvious that an envy-free allocation π with $\text{mcar}(\pi) \geq 1$ or $\text{usw}(\pi) \geq 1$ must in both cases have $\text{size}(\pi) \geq 1$. Let us conversely assume that there exists some envy-free allocation π with $\text{size}(\pi) \geq 1$ for $(\mathcal{R}^*, \mathcal{A}^*, (u^*_a)_{a\in\mathcal{A}^*})$. We want to show that $\text{mcar}(\pi) \geq 1$ and $\text{usw}(\pi) \geq 1$ also hold for $(\mathcal{R}^*, \mathcal{A}^*, (u^*_a)_{a\in\mathcal{A}^*})$. Since $\text{size}(\pi) \geq 1$, at least one resource r is allocated. If r is not a shadow resource, then at least one

[3] Note that given any $\pi^{\mathcal{E}}$ for $\mathcal{E} \in \{\text{esw}, \text{mcar}, \text{usw}, \text{size}\}$, we can compute each of the respective other allocations in polynomial time. Here, the condition $\mathcal{E}(\pi^*) \geq 1$ corresponds to the setting $t = 1$ in $\mathcal{E}$-EF-PA.

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 original 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 $\text{mcar}(\pi) \geq 1$ and $\text{usw}(\pi) \geq 1$.

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

($\Longrightarrow$) Assume there exists an envy-free allocation π^{esw} with $\text{esw}(\pi^{\text{esw}}) \geq 1$ for $(\mathcal{R}, \mathcal{A}, (u_a)_{a\in\mathcal{A}})$. We construct a desired allocation π^{size} for $(\mathcal{R}^*, \mathcal{A}^*, (u^*_a)_{a\in\mathcal{A}^*})$ as follows. Analogously to π^{esw}, we let $\pi^{\text{size}}_a = \pi^{\text{esw}}_a$ 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 of at least v. 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 $\text{size}(\pi^{\text{size}}) \geq 1$ 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 a positive value. Thus, we have $\text{esw}(\pi^{\text{size}}) \geq 1$. We create an allocation π^{esw} for $(\mathcal{R}, \mathcal{A}, (u_a)_{a\in\mathcal{A}})$ in a straight-forward way by setting $\pi^{\text{esw}}_a := \pi^{\text{size}}_a$ 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 $(\mathcal{R}, \mathcal{A}, (u_a)_{a\in\mathcal{A}})$. $\square$

According to Lemma 3, if we show that esw-EF-PA is strongly NP-hard for ternary utility values $0 < v < u$, then we automatically also get the strong NP-hardness of $\mathcal{E}$-EF-PA for each $\mathcal{E} \in \{\text{mcar}, \text{usw}, \text{size}\}$. 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. *Let $\mathcal{E} \in \{\text{esw}, \text{mcar}, \text{usw}, \text{size}\}$ and let $v, u \in \mathbb{N}$ be fixed with $0 < v < u$. Then, $\mathcal{E}$-EF-PA is strongly NP-hard, even if each agent assigns only values from $\{0, v, u\}$ to the resources and $t = 1$.*

By Lemma 3, it suffices to show the strong NP-hardness for esw-EF-PA. To this end, the proof serves as a case distinction over the values of u and v. Each lemma shows a different reduction from the NP-hard Exact Cover by 3-Sets (X3C) problem [21]. Given a multiset $X = \{x_1, x_2, \ldots, x_{3n}\}$ 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'. The detailed proof of Theorem 7 is given in the appendix.

6 Conclusions

We studied how to allocate indivisible resources to agents in an envy-free manner by relaxing the common requirement that all resources must be allocated. We considered envy-free partial allocations that provide at least some utility or allocate some resources from both systematic or individual perspectives, and we obtained comprehensive results under various classes of utilities. While most of the problems we considered are generally NP-hard, we identified several tractable results for binary utilities by establishing interesting connections to matching problems on bipartite graphs. Notably, our tractable results imply that, at least for binary utilities, if the goal is to allocate some resources or provide some utility to agents, then the problem of finding envy-free partial allocations (or confirming their non-existence) can be efficiently solved. Complementing the well-known NP-hardness of finding envy-free complete allocations, our results provide a more fine-grained understanding of the computational complexity of finding efficient envy-free allocations.

Our work can be extended in several directions. First, we show a stark contrast: some cases are tractable under binary utilities but all scenarios become NP-hard under ternary utilities. It is worth further exploring this frontier, in particular, bivalued utilities other than the combination of 0 and 1, that lie between binary and ternary utilities. In the appendix, we provide some initial results for $1/2$ utilities. 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, we assumed all resources are goods. A natural extension is to study chores or mixed resources. For chores, the case of a planner who wants to distribute as many tasks to agents as possible well justifies our measures size and mcar. 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 [26]. Finally, applying our setting for alternative fairness notions, such as equitability, instead of envy-freeness offers another research direction. We note that for identical utilities, these two fairness notions are equivalent.

References

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

[2] Haris Aziz, Serge Gaspers, Simon Mackenzie, and Toby Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.

[3] Haris Aziz, Ildikó Schlotter, and Toby Walsh. Control of fair division. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI '16)*, pages 67–73, 2016.

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

[5] Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI '21)*, pages 5119–5126, 2021.

[6] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Greedy algorithms for maximizing nash social welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, (AAMAS '18)*, pages 7–13, 2018.

[7] Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. Almost full efx exists for four agents. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI '22)*, pages 4826–4833, 2022.

[8] Aurélie Beynier, Yann Chevaleyre, Laurent Gourvès, Ararat Harutyunyan, Julien Lesca, Nicolas Maudet, and Anaëlle Wilczynski. Local envy-freeness in house allocation problems. *Autonomous Agents and Multi-Agent Systems*, 33(5):591–627, 2019.

[9] Bernhard Bliem, Robert Bredereck, and Rolf Niedermeier. Complexity of efficient and envy-free resource allocation: Few agents, resources, or utility levels. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence, (IJCAI '16)*, pages 102–108, 2016.

[10] Niclas Boehmer, Robert Bredereck, Klaus Heeger, Dusan Knop, and Junjie Luo. Multivariate algorithmics for eliminating envy by donating goods. *Autonomous Agents and Multi-Agent Systems*, 38(2):43, 2024.

[11] Sylvain Bouveret and Jérôme 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.

[12] Sylvain Bouveret, Yann Chevaleyre, and Nicolas 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.

[13] Xiaolin Bu, Zihao Li, Shengxin Liu, Jiaxin Song, and Biaoshuai Tao. On the complexity of maximizing social welfare within fair allocations of indivisible goods. *CoRR*, abs/2205.14296, 2022.

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

[15] Ioannis Caragiannis, Nick Gravin, and Xin Huang. 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 '19)*, pages 527–545. ACM, 2019.

[16] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. *ACM Transactions on Economics and Computation*, 7(3):12:1–12:32, 2019.

[17] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. A little charity guarantees almost envy-freeness. *SIAM Journal on Computing*, 50(4):1336–1358, 2021.

[18] Britta Dorn, Ronald de Haan, and Ildikó Schlotter. Obtaining a proportional allocation by deleting items. *Algorithmica*, 83(5):1559–1603, 2021.

[19] Zack Fitzsimmons, Vignesh Viswanathan, and Yair Zick. On the hardness of fair allocation under ternary valuations. *CoRR*, abs/2403.00943, 2024.

[20] Jiarui Gan, Warut Suksompong, and Alexandros A. Voudouris. Envy-freeness in house allocation problems. *Mathematical Social Sciences*, 101:104–106, 2019.

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

[22] Daniel Halpern, Ariel D. Procaccia, Alexandros Psomas, and Nisarg Shah. Fair division with binary valuations: One rule to rule them all. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE '20)*, pages 370–383, 2020.

[23] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, Hejun Wang, and Lirong Xia. Fair division through information withholding. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI '20)*, pages 2014–2021, 2020.

[24] Naoyuki Kamiyama, Pasin Manurangsi, and Warut Suksompong. On the complexity of fair house allocation. *Operations Research Letters*, 49(4):572–577, 2021.

[25] Lyle Ramshaw and Robert E Tarjan. On minimum-cost assignments in unbalanced bipartite graphs. *Technical report*, 20, 2012.

[26] Biaoshuai Tao, Xiaowei Wu, Ziqi Yu, and Shengwei Zhou. On the existence of EFX (and pareto-optimal) allocations for binary chores. In *Procedings of the 18th International Joint Conference on Frontiers of Algorithmics (IJTCS-FAW '24)*, pages 33–52, 2024.

[27] Toby Walsh. Challenges in resource and cost allocation. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI '15)*, pages 4073–4077, 2015.

Contact Information

Robert Bredereck
Institut für Informatik
TU Clausthal
Clausthal-Zellerfeld, Germany
Email: robert.bredereck@tu-clausthal.de

Andrzej Kaczmarczyk
Department of Computer Science
The University of Chicago
Chicago, USA
Email: akaczmarczyk@uchicago.edu

Junjie Luo
School of Mathematics and Statistics,
Beijing Jiaotong University, Beijing, China
Email: jjluo1@bjtu.edu.cn

Bin Sun
Institut für Informatik
TU Clausthal
Clausthal-Zellerfeld, Germany
Email: bin.sun@tu-clausthal.de

Appendix

7 Additional Material for Section 5

7.1 Proof of Thm. 7

By Lemma 3, it suffices to show the strong NP-hardness for esw-EF-PA. To this end, the following Lemma 3 to 6 serve as a case distinction over the values of u and v. Each lemma shows a different reduction from the NP-hard Exact Cover by 3-Sets (X3C) problem [21]. Given a multiset $X = \{x_1, x_2, \dots, x_{3n}\}$ and a collection $C = \{S_1, S_2, \dots, S_m\}$ 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 $m > 3n$, as we can always add dummy 3-sets to guarantee this.

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

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

- There are m *cover agents* $\mathcal{A}_C = \{a_1, a_2, \dots, a_m\}$ and a *special agent* a^*, i.e., $\mathcal{A} = \mathcal{A}_C \cup \{a^*\}$.
- There are $3n$ *normal resources* $\mathcal{R}_N = \{r_1, r_2, \dots, r_{3n}\}$, $(k-3)n$ *small resources* $\mathcal{R}_S = \{s_1, s_2, \dots, s_{(k-3)n}\}$, $(m-n)$ *dummy resources* $\mathcal{R}_D = \{d_1, \dots, d_{m-n}\}$, and a *special resource* s^*, i.e., $\mathcal{R} = \mathcal{R}_N \cup \mathcal{R}_S \cup \mathcal{R}_D \cup \{s^*\}$.
- 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. Each cover agent values each dummy resource and the special resource s^* as kv. Finally, the special agent a^* values the special resource s^* as kv and values all other resources as 0.

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

$(\Longrightarrow)$ Assume that (X, C) is a YES-instance, then there is a subset $C' \subseteq C$ with $|C'| = n$ such that each $e \in X$ occurs in exactly one member of C'. For each $S_j \in C$, if $S_j \in C'$, then we allocate the 3 corresponding normal resources to a_j resulting in value that a_j gets being exactly 3 for now. Then, a_j will also get $k-3$ small resources, finally getting the value $u\,(= kv)$. If $S_j \notin C'$, then we allocate 1 dummy resource to a_j, which also results in value u. In addition, the special agent will get the special resource which is valued at exactly u. It is easy to check that every agent gets utility $u = kv$ and values other agents' bundle by at most $u = kv$. Thus, $\mathcal{I}$ is also a YES-instance.

$(\Longleftarrow)$ Assume that there is a solution for the constructed instance $\mathcal{I}$ of esw-EF-PA. Since in this solution each agent has to get a bundle with a positive value, the special agent will get the special resource s^*. Then, each of the cover agents will require a bundle of value at least $u = kv$. Since the total value that all the m cover agents can receive is at most $3nv + (k-3)nv + (m-n)kv = mkv$, the value that each cover agent receives should be exactly kv. Notice that $m-n$ dummy resources can be allocated to $m-n$ cover agents, so the remaining n agents get all the normal and small resources. Since each remaining agent can receive at most value $3v$ from the normal resources, we conclude that each of them gets 3 normal resources they like and $k-3$ small resources. Let $I_j = \{i_{ja}, i_{jb}, i_{jc}\}$ be the normal resources received by each remaining agent a_j. Then we can find n corresponding sets $S_j = \{x_{ja}, x_{jb}, x_{jc}\}$ from C, which are pairwise disjoint. This induces a feasible solution C' for (X, C). □

Next we consider the case with $u = 2v$. The distinctive feature of the following proof, lies in our creation of standard agents and special resources as benchmarks, ensuring that the value of the bundle

Table 3: Agent's utility functions in the proof of Lemma 5.

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

desired by each agent exceeds a certain constant value. Additionally, we introduce a large number of special "observer" agents 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* $\{0, v, u\}$, $u = 2v > 0$*, 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 in Table 3 and the standard agents are not interested in any of the other resources:
- Each observer assigns value 2 to each blank resource and each special resource. In particular, there are three different kinds of observers. Listing only resources for which the observers have a non-zero value, we define them as follows: (1) Each observer $w_{i,j;k}$ of type 1 values the two normal resources r_i, r_j, and one dummy resource d_k at $2v$, respectively. (2) Each observer $w'_{i;j;k}$ of type 2 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 $\binom{3n}{2} \cdot m + 3m \cdot n \cdot 2(m-n) + 1$ observers. Thus, there are $O(m^2 n)$ numbers of observers and blank resources. Assuming that there is a solution for the constructed instance $\mathcal{I}$ of esw-EF-PA, we have the following observations.

Ob. 1. We first consider the standard agents. Since each agent has to get a positive value, standard agent d will get r_3^*. Then, standard agent b will get r_1^* and 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 the cover agents are not interested in the blank resources and the sum of the value that these three kinds of resources can provide is at most $4mv$.

Ob. 3. Each cover agent gets utility exactly $4v$. Thus, if a cover agent gets 2 dummy resources, they 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. Otherwise, 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. Otherwise, the agent needs another resource. Yet, neither can it be a dummy nor a normal resource because of, respectively, the type 2 and 1 observers.

Ob. 9. It follows from Observations 3 and 6–8 above, that if some cover agent gets a dummy resource then they 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 some normal resources and small resources and each of them gets exactly 1 small resource. This is because the value that each of them can get from the 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 is 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 each a_j such that the value that a_j can get is exactly $3v$ for now. In addition, each a_j 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 $4v$. Further, 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 agent is envious. Thus, $\mathcal{I}$ is also a YES-instance.

($\Longleftarrow$) Since $\mathcal{I}$ 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 sets $S_j = \{x_{ja}, x_{jb}, x_{jc}\}$. We can find exactly n such disjoint sets, which 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. *With ternary utility values* $\{0, v, u = kv + c\}$ *for* $v > c > 0, k > 0$ esw-*EF-PA is strongly NP-hard and* $t = 1$.

Proof. This final case builds up on ideas from Case 2, but has some more technicalities and uses more involved auxiliary agents and resources. The hardness proof is again realized via a reduction from X3C. Given an instance (X, C) of the X3C problem, we construct an instance $\mathcal{I} = (\mathcal{R}, \mathcal{A}, (u_a)_{a \in \mathcal{A}}, t = 1)$ of the esw-EF-PA problem as follows.

We have the following resources.

- A set of *element resources* $\mathcal{R}_X = \{r_1, r_2, \ldots, r_{3n}\}$ as well as a set of *dummy resources* $\mathcal{R}_D = \{d_1, \ldots, d_{3(m-n)}\}$.
- Three sets of *special resources*: namely *booster resources* $\mathcal{R}_B$ as well as *guard resources* $\mathcal{R}_G^{\text{rrd}}$ and $\mathcal{R}_G^{\text{rdd}}$, whose cardinalities are specified later.
- Altogether, $\mathcal{R} = \mathcal{R}_X \cup \mathcal{R}_D \cup \mathcal{R}^B \cup \mathcal{R}_G^{\text{rrd}} \cup \mathcal{R}_G^{\text{rdd}}$.

We have the following agents.

- A set of m *cover agents* $\mathcal{A}_C = \{a_1, a_2, \ldots, a_m\}$.
- A set of three *(utility) booster agents* $\mathcal{A}_B = \{b_1, b_2, b_3\}$.
- A set of $\binom{3n}{2} \cdot 3(m-n)$ *rrd-guards* $\mathcal{A}_G^{\text{rrd}} = \{g^{\text{rrd}}(r, r', d) \mid r, r' \in \mathcal{R}_X, d \in \mathcal{R}_D, r \neq r'\}$.
- A set of $\binom{3(m-n)}{2} \cdot 3n$ *rdd-guards* $\mathcal{A}_G^{\text{rdd}} = \{g^{\text{rrd}}(r, d, d') \mid r \in \mathcal{R}_X, d, d' \in \mathcal{R}_D, d \neq d'\}$.

Before we go into the formal proof, we provide some intuition:

- The cover agents together with the element and dummy resources are meant to encode the X3C solution: A cover agent corresponding to a set selected in the X3C solution gets a bundle of three element resources and a cover agent corresponding to a set that is not selected gets a bundle of three dummy resources.
- Booster agents are used to boost the minimum utility value of the other agents' bundle: They need to obtain predetermined bundles of booster resources and the other agents also have some value for some of these resources. To avoid envy towards the booster agents, the bundles of the other agents need to have a specific value for them.
- Guard agents ensure that cover agents can only get bundles that either contain three element or three dummy resources. Each guard "forbids" a specific combination mixed bundles.

Next, we fully specify the special resources. To do so, we define that $k_1 = k + 1$ and $k_2 = \min\{k' \in \mathbb{N} \mid k' * v > 2u\}$. Now, $\mathcal{R}_B = \mathcal{R}_B^{b_1} \cup \mathcal{R}_B^{b_2} \cup \mathcal{R}_B^{b_3}$ with $\mathcal{R}_B^{b_1} = \{s^{b_1}\}$, $\mathcal{R}_B^{b_2} = \{s_1^{b_2}, s_2^{b_2}, \ldots, s_{k_1}^{b_2}\}$, and $\mathcal{R}_B^{b_3} = \{s_1^{b_3}, s_2^{b_3}, \ldots, s_{k_2}^{b_3}\}$. Moreover, the guards resources are $\mathcal{R}_G^{\text{rrd}} = \{s_1^{\text{rrd}}, s_2^{\text{rrd}}, \ldots, s_{k_2 \cdot |\mathcal{A}_G^{\text{rrd}}|}^{\text{rrd}}\}$ and $\mathcal{R}_G^{\text{rdd}} = \{s_1^{\text{rdd}}, s_2^{\text{rdd}}, \ldots, s_{k_2 \cdot |\mathcal{A}_G^{\text{rdd}}|}^{\text{rdd}}\}$. Note that $v \leq kv < u < k_1 v < 2u < k_2 v < 3u$.

We now define the *utility functions* of the agents by specifying non-zero utility values (that is, in all combinations not specified here, the agent assigns utility zero to the resource).

- For each cover agent a_j and each element resource r_i, we have $u_{a_j}(r_i) = u$ if $x_i \in S_j$ (and $u_{a_j}(r_i) = 0$ otherwise). Moreover, cover agents assign utility u to each dummy resource and to $s_1^{b_3}$, $s_2^{b_3}$, and $s_3^{b_3}$.
- The booster agent b_1 assigns utility u to s^{b_1} only.
- The booster agent b_2 assigns utility u to s^{b_1} and utility v to each resource from $\mathcal{R}_B^{b_2}$.
- The booster agent b_3 assigns utility u to $s_1^{b_2}$ and $s_2^{b_2}$, and utility v to each resource from $\mathcal{R}_B^{b_3}$.
- The rrd-guard $g^{\text{rrd}}(r, r', d)$ assigns utility u to the element resources r and r' and to the dummy resource d and utility v to each of the resources from $\mathcal{R}_G^{\text{rrd}}$. They also assigns utility u to $s_1^{b_3}$ and $s_2^{b_3}$.

- The rdd-guard $g^{\mathrm{rrd}}(r, d, d')$ assigns utility u to the element resource r and to the dummy resources d and d' and utility v to each of the resources from $\mathcal{R}_G^{\mathrm{rdd}}$. They also assigns utility u to $s_1^{b_3}$ and $s_2^{b_3}$.

Finally, there are overall $m + 3 + \binom{3n}{2} \cdot 3(m-n) + \binom{3(m-n)}{2} \cdot 3n$ agents and $3n + 3(m-n) + 1 + k_1 + k_2 + k_2 \cdot \binom{3n}{2} \cdot 3(m-n) + k_2 \cdot \binom{3(m-n)}{2} \cdot 3n$ resources. The the construction can definitely be performed in polynomial time.

Assuming that there is a solution for the constructed instance $\mathcal{I}$, we have the following observations.

Ob. 1. The agent b_1 will get s^{b_1}, otherwise the value of their bundle will be 0.

Ob. 2. The agent b_2 will get all resources from $\mathcal{R}_B^{b_2}$. Otherwise, they will envy b_1 who gets value of u in their eyes.

Ob. 3. The agent b_3 will get all resources of $\mathcal{R}_B^{b_3}$. Otherwise, they will envy b_2 who gets value of $2u$ in their eyes.

Ob. 4. Each of the cover agents will get exactly three resources from $\mathcal{R}_X \cup \mathcal{R}_D$ and get value of $3u$. Otherwise, some of them will envy booster agent b_3 who gets value of $3u$ in their eyes. By the pigeonhole principle, this implies that other agents (who are not cover agents) cannot get any element or dummy resource.

Ob. 5. Each rrd-guard gets k_2 arbitrary resources from $\mathcal{R}_G^{\mathrm{rrd}}$ and each rdd-guard gets k_2 arbitrary resources from $\mathcal{R}_G^{\mathrm{rdd}}$. If any of the guard agent gets fewer than k_2 of the corresponding resources, then the value of their bundle would be smaller than $2u$ and they would envy booster agent b_3. By pigeonhole principle, no guard agent can get more than k_2 of the corresponding resources.

Ob. 6. It is not possible for any cover agent to get a bundle with resources from both R_X and R_D. Assume that such agent a' exists and gets $\{r, r', d\}$ with $r \neq r' \in \mathcal{R}_X$ and $d \in \mathcal{R}_D$. Then, the rrd-guard $g(r, r', d)$ would require at least $k_2 + 1$ resources from $\mathcal{R}_G^{\mathrm{rrd}}$ (to not envy a'); a contradiction to the previous observation. Analogously (replacing rrd-guards by rdd-guards), no cover agent can get $\{r, d, d'\}$ with $r \in \mathcal{R}_X$ and $d \neq d' \in \mathcal{R}_D$.

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

($\Longrightarrow$) Assume that (X, C) is a YES-instance, then there exists a subset $C' \subseteq C$ with $|C'| = n$ such that every element of X occurs in exactly one member of C'. We construct a solution for $\mathcal{I}$ as follows. Booster agents will get resources as discussed in Observations 1 to 3 and guard-agents will get get resources as discussed in Observation 5. For each $S_j \in C$, if $S_j \in C'$, we allocate the three corresponding element resources to a_j. If $S_j \notin C'$, we allocate three arbitrary dummy resources to a_j. Since there are no combinations of element resources and dummy resources (observation 6), no guard-agent will envy (observation 4). Therefore, no agent will envy in this case. Thus, $\mathcal{I}$ is also a YES-instance.

($\Longleftarrow$) Assume that $\mathcal{I}$ is a YES-instance. According to the above observations, there is a set A^* of n cover agents who each get a bundle of three element resources. Let $A^* = \{a_{j_1}, a_{j_2}, \ldots, a_{j_n}\}$. Now, observe that $C^* = \{S_{j_1}, S_{j_2}, \ldots, S_{j_n}\}$ clearly form a feasible solution for (X, C): Each element is covered exactly once since each element resource is assigned to exactly one of these cover agents. Moreover, the bundles of A^* correspond exactly to the item subsets from C^*. Thus, (X, C) is a YES-instance. $\square$

The claim of Thm. 7 follows from Lemmas 3 to 6.

8 Bivalued valuations

When $u_i(a_j) \in \{1, 2\}$ for any agent $a_i \in \mathcal{A}$ and any resource $r_j \in \mathcal{R}$, the four efficiency measures are equivalent.

Lemma 7. *When $t = 1$, for any $\mathcal{E}_1, \mathcal{E}_2 \in \{\mathrm{usw}, \mathrm{esw}, \mathrm{size}, \mathrm{mcar}\}$, $\mathcal{E}_1$-EF-PA and $\mathcal{E}_2$-EF-PA are equivalent.*

Proof. Since $u_i(a_j) \in \{1, 2\}$, if an envy-free allocation $\boldsymbol{\pi}$ satisfies that $\mathrm{usw}(\boldsymbol{\pi}) \geq 1$, then it follows that $\mathrm{size}(\boldsymbol{\pi}) \geq 1$. By envy-freeness, $\mathrm{size}(\boldsymbol{\pi}) \geq 1$ implies that $\mathrm{mcar}(\boldsymbol{\pi}) \geq 1$. By positive value of every resource, $\mathrm{mcar}(\boldsymbol{\pi}) \geq 1$ implies that $\mathrm{esw}(\boldsymbol{\pi}) \geq 1$, which implies that $\mathrm{usw}(\boldsymbol{\pi}) \geq 1$. □

We show that we can reduce the problem to the case where each agent gets at most two resources.

Lemma 8. *When $t = 1$, for $\mathcal{E}$-EF-PA with any $\mathcal{E} \in \{\mathrm{usw}, \mathrm{esw}, \mathrm{size}, \mathrm{mcar}\}$, there exists a desired envy-free allocation if and only if*

1. *there exists a desired envy-free allocation where each agent gets exactly one resource, or*
2. *there exists a desired envy-free allocation where some agents get one resource they like and the remaining agents get two resources.*

Proof. According to Lemma 7, it suffices to consider one efficiency measure, say esw. The “if” direction is trivial and we show the “only if” direction. We define a new utility function u' for each agent such that $u'_i(r_j) = u_i(r_j) - 1$ for any resource $r_j \in \mathcal{R}$. Since the value of each resource is decreased by 1, it is easy to see that there exists an envy-free allocation where each agent gets exactly one resource under the original utility function u if and only if there exists such an allocation under u'. Notice that u' is a binary valuation, and according to Lemma 2, there exists such an allocation under u' if and only if $|X| \leq |Y_L|$. Suppose $|X| > |Y_L|$, which means case (1) does not happen, we show that case (2) must happen. Since case (1) does not happen, for any envy-free allocation, at least one agent gets more than one resource. Moreover, since $|X_S| > |Y_S|$, we have $|Y| < 2|X|$. Since $u_i(r_j) \in \{1, 2\}$, if one agent gets more than two resources, than all the other agents should get at least 2 resources to be envy-free, which overall needs more than $2|X| > |Y| = |R|$ resources and is impossible. Therefore, in any envy-free allocation, every agent gets either one resource or two resources. Moreover, every agent who gets one resource must get one resource they like, as otherwise they will envy agents who get two resources. □

Open Question: For any $\mathcal{E} \in \{\mathrm{usw}, \mathrm{esw}, \mathrm{size}, \mathrm{mcar}\}$, is $\mathcal{E}$-EF-PA with bivalued valuations (1 and 2) and $t = 1$ NP-hard?