The Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)

The Complexity of Extending Fair Allocations of Indivisible Goods

Argyrios Deligkas[1], **Eduard Eiben**[1], **Robert Ganian** [2], **Tiger-Lily Goldsmith**[1], **Stavros D. Ioannidis**[1]
[1]Department of Computer Science, Royal Holloway, University of London, UK
[2]Algorithms and Complexity Group, TU Wien, Austria
{argyrios.deligkas, eduard.eiben,tiger-lily.goldsmith, stavros.ioannidis}@rhul.ac.uk, rganian@ac.tuwien.ac.at

Abstract

We initiate the study of computing envy-free allocations of indivisible items in the extension setting, i.e., when some part of the allocation is fixed and the task is to allocate the remaining items. Given the known NP-hardness of the problem, we investigate whether—and under which conditions—one can obtain fixed-parameter algorithms for computing a solution in settings where most of the allocation is already fixed. Our results provide a broad complexity-theoretic classification of the problem which includes: (a) fixed-parameter algorithms tailored to settings with few distinct types of agents or items; (b) lower bounds which exclude the generalization of these positive results to more general settings. We conclude by showing that—unlike when computing allocations from scratch—the non-algorithmic question of whether more relaxed EFX allocations exist can be completely resolved in the extension setting.

Introduction

Finding a "fair" allocation of indivisible items or resources to a provided set of agents, each with their own preferences, is one of the central tasks arising in the area of computational social choice. The arguably most classical and established notion of fairness used in these settings is *envy-freeness*, where we ask for an allocation $\pi : M \to N$ from the set M of items to the set N of agents such that for each pair of agents $i, j \in N$, i prefers $\pi^{-1}(i)$ to $\pi^{-1}(j)$; in other words, no agent envies another agent. It is well known that an envy-free allocation need not exist, and in fact, determining whether one exists is NP-complete (Bouveret and Lang 2008).

The aforementioned intractability has led to a flurry of research to circumvent this issue. One approach that has been proposed is to consider "less restrictive" versions of envy-freeness instead, with the aim of not only ensuring that the computation of an assignment is tractable but—even more desirably—that one always exists. Notably, it is known that one can always compute, in polynomial time, an allocation which is *envy-free up to one item* (*EF1*) (Amanatidis, Markakis, and Ntokos 2020; Budish 2011; Igarashi et al. 2024). However, such allocations may sometimes be considered very far from "fair". A recently proposed intermediate notion between envy-free and EF1 allocations is *envy-free up to any item* (*EFX*) (Caragiannis et al. 2019; Feldman, Mauras, and Ponitka 2024), but it is not known whether EFX allocations always exist and their polynomial-time computability remains open as well.

Another notable approach to tackling the problem of computing envy-free allocations is to identify precise conditions under which it can be solved efficiently. This is typically done by investigating the problem through the lens of *parameterized complexity* (Downey and Fellows 2013; Cygan et al. 2015)—a refinement of the classical complexity paradigm where inputs are analyzed not only with respect to their size n, but also to a numerical parameter k which measures some well-defined quantity. While we cannot hope to obtain an $n^{\mathcal{O}(1)}$ algorithm for computing envy-free allocations in general, in this setting one aims to design so-called *fixed-parameter tractable* (or FPT) algorithms for the problem, that is, algorithms running in time $f(k) \cdot n^{\mathcal{O}(1)}$ for some computable function f. The simplest parameters considered in past works on fair division include, for example, the number of agents or items; however, such parameterizations place strong restrictions on the input instances, and so more recent works have focused on developing algorithms parameterized by the number of *agent types* (where two agents have the same type if they have identical preferences over the items), or the number of *item types* (where two items have the same type if they are valued the same by every agent) (Deligkas et al. 2021; Eiben et al. 2023a; Nguyen and Rothe 2023), see also (Brânzei, Lv, and Mehta 2016; Ganian, Ordyniak, and Rahul 2019).

While the aforementioned approaches have by now provided a fairly detailed understanding of computing an envy-free allocation from scratch, in this article we turn to the problem of extending a partial allocation—that is, computing an envy-free allocation when part (or even most) of the allocation is already fixed. This *envy-free allocation extension* problem arises naturally whenever one needs to deal with resources that have already been assigned or must be assigned to certain agents—consider, e.g., the case where a few new items are made available after an allocation has been fixed and we are not allowed to take items from agents, or the setting where most employees in a company already have fixed tasks, but we need to distribute a set of new tasks to recently hired employees. It is worth noting that while the complexity of extending partial solutions has been extensively researched in settings as diverse as data completion (Ganian et al. 2018; Koana, Froese, and Niedermeier 2021; Eiben et al. 2023b; Koana, Froese, and Niedermeier 2023) and graph drawing (Angelini et al. 2015; Arroyo, Derka, and Parada 2019; Eiben et al. 2020; Bhore et al. 2023), this has not yet been systematically studied in classical resource allocation in spite of the many situations in which one may need to deal with some items being preassigned[1].

Our Contribution

We study the envy-free allocation extension problem in the classical setting of additive utilities, as outlined below.

ENVY-FREE ALLOCATION EXTENSION	
Input:	A set M of indivisible items, a set N of n agents with additive valuations, and a partial allocation $\gamma : M' \to N$ of some subset $M' \subseteq M$ of items.
Question:	Does there exist an extension allocation $\pi : (M \setminus M') \to N$ of the remaining ("open") items such that $\gamma \cup \pi$ is envy-free?

Notice that ENVY-FREE ALLOCATION EXTENSION coincides with the usual problem of obtaining an envy-free allocation from scratch in the special case where $\gamma = \emptyset$, and hence is necessarily NP-hard. However, this intractability does not properly reflect typical usage scenarios of the problem: in many cases of interest, one may be dealing with allocations γ' which are "almost complete". Hence, we turn to the aforementioned parameterized paradigm and ask which parameterizations of the partial assignment γ allow us to achieve fixed-parameter tractability (i.e., obtain a fixed-parameter algorithm).

We begin our investigation by considering the complexity of the problem with the number $k = |M \setminus M'|$ of yet-to-be-allocated (or *open*) items as the parameter. It is worth noting that ENVY-FREE ALLOCATION EXTENSION can be solved via a trivial $n^{\mathcal{O}(k)}$ algorithm (a so-called XP *algorithm*), and that computing an envy-free allocation from scratch is trivially fixed-parameter tractable w.r.t. k—hence, one could have expected ENVY-FREE ALLOCATION EXTENSION to be fixed-parameter tractable w.r.t. k as well. However it turns out that this is not the case: we will later show (in Theorem 2) that the problem is W[1]-hard when parameterized by k, which excludes fixed-parameter tractability under well-established complexity assumptions.

Unsatisfied by this outcome, we ask whether one can at least obtain a fixed-parameter algorithm under some mild restrictions on the input instance. For example instances with a bounded number of item or agent types arise naturally in large-scale models: one often aggregates similar items into categories (depending on the setting, these may be, e.g., "cars" and "houses", or "PhD students" and "Postdocs"), and the inherent limitations of how preferences are collected may result in many agents being represented as having the same preferences. Inspired by related works on resource allocation (Deligkas et al. 2021; Eiben et al. 2023a; Nguyen and Rothe 2023), we consider additionally parameterizing by the number n_t of the agent types or the number m_t of item types. As our first positive result, we obtain an algorithm when we bound the number of agent types.

Theorem 1. ENVY-FREE ALLOCATION EXTENSION *is fixed-parameter tractable when parameterized by the number k of open items plus the number n_t of agent types.*

Turning to the analogous question for item types, we show that the situation is entirely different: allocating k open items is intractable even when there are only a few item types.

Theorem 2. ENVY-FREE ALLOCATION EXTENSION *is* W[1]-hard *when parameterized by the number k of open items plus the number m_t of item types.*

The proof of Theorem 1 relies on a combination of branching techniques with insights into the structure of hypothetical solutions; while it is not trivial, in this case it is the lower bound—Theorem 2—which presents a much more surprising result. Indeed, at first glance the problem seems far from intractable: we only need to assign k items to a set of agents, and for each open item i, we can easily test whether assigning it to an agent j would make any of the other agents in N envy j w.r.t. their current items. The only caveat is that we do not know which of the other agents will receive the remaining open items—and in spite of there being only at most k open items in total, we show that the uncertainty of how these will be allocated can be used to obtain a highly non-trivial reduction from the classical W[1]-hard MULTICOLORED CLIQUE problem (Downey and Fellows 2013).

In the second part of our article, we take aim at the setting where many items may be open, but we are only allowed to allocate these items to at most p-many agents. This is precisely the situation that arises when, e.g., allocating tasks to a few new employees, or needing to distribute new items to individuals while adhering to constraints on shipping costs or how many of the individuals should be contacted. Formally, this is captured by asking for an allocation π which assigns all open items to a set S of at most p-many agents. We will distinguish between whether the set S of potential recipients is provided in the input (in which case we speak of RESTRICTED ENVY-FREE ALLOCATION EXTENSION or simply REFAE) or can be selected along with π (giving rise to FREE ENVY-FREE ALLOCATION EXTENSION or simply FEFAE). Distinguishing these two variants of our problem is not only necessary to differentiate between distinct usage scenarios, but will also turn out to have an impact on the problem's complexity.

We note that while both variants are in NP, attempting to solve either of them when parameterized by p (the number of agents that receive items) alone is entirely hopeless—and this holds even in the strong setting, i.e., with unary-encoded valuations.

---

[1] In parallel to this article, a pre-print for a separate work investigating different aspects of fair allocation extension has very recently been published on arxiv (Prakash, Igarashi, and Vaish 2024).

Theorem 3. *Both* REFAE *and* FEFAE *are strongly* NP-hard *even when* $p \leq 2$.

This once again raises the question of whether one can use a bound on n_t or m_t to obtain efficient algorithms. Given Theorem 1 and Theorem 2, one might expect that agent types would once again allow us to obtain fixed-parameter algorithms in this case. And yet, we show that when dealing with p the situation is reversed: unexpectedly, here it is the number m_t of item types that provides the key to tractability. Indeed, one can establish strong W[1]-hardness as well as weak paraNP-hardness of REFAE and FEFAE when parameterized by the total number of agents in the instance (which upper-bounds both n_t and p) directly from existing results in the literature (Bliem, Bredereck, and Niedermeier 2016). Before proceeding, we at least show that the former lower bound is tight: the problems admit XP-algorithms in the case where the utilities are not encoded in binary.

Theorem 4. *If the valuations are encoded in unary, both* REFAE *and* FEFAE *are in* XP *when parameterized by the number* p *of recipients plus the number* n_t *of agent types.*

Turning to the setting with few item types, we provide a fixed-parameter algorithm for the restricted case (i.e., when the set S of recipients is already given) and an XP algorithm for the free case (when the task is to also find S on its own). Here, the latter result is the best one could hope for at this point, as fixed-parameter tractability is immediately excluded by Theorem 2. In contrast to Theorem 4, both of these algorithms can be applied regardless of whether the utilities are encoded in unary or binary.

Theorem 5. *When parameterized by* p *plus the number* m_t *of item types,* REFAE *and* FEFAE *are fixed-parameter tractable and* XP*-tractable, respectively.*

An overview of our complete complexity-theoretic classification of ENVY-FREE ALLOCATION EXTENSION under the considered restrictions is provided in Table 1.

While not part of our main technical contributions, in the final section of the article, we also make a few observations about the behavior of relaxed variants of envy-freeness in the extension setting that may be of general interest to the community. In particular, we prove that the known result guaranteeing the existence of an EF1 allocation can be strengthened to also hold when extending *any* provided partial allocation. Moreover, this result is tight: there exist partial EFX allocations which cannot be extended to a full EF1 allocation. Similarly, while the question of whether an EFX allocation "from scratch" is guaranteed to exist remains one of the main open questions in the field, we show that there exist partial envy-free allocations which cannot be extended to a full EFX allocation.

Preliminaries

For an integer ℓ, we use $[\ell]$ as shorthand for the set $\{1, \ldots, \ell\}$. We also use the standard $\mathcal{O}^*$ notation to suppress polynomial factors of the input size in the running time.

Our instances include a set of indivisible items $M = \{a_1, \ldots, a_m\}$ and a set of n agents $N = [n]$. Every agent i has an *additive valuation* function v_i that assigns a non-negative value $v_i(a)$ for every item $a \in M$ and for every subset, or *bundle*, of items $B \subseteq M$ we denote $v_i(B) := \sum_{j \in B} v_i(a_j)$. If $v_i(a) = v_j(a)$ for every $a \in M$, then we say that agents i and j are of the same type. We will use n_t to denote the number of different *agent types* in N. If for two items a and a' we have that $v_i(a) = v_i(a')$ for every agent i, then we say that a and a' are of the same type. We will use m_t to denote the number of different *item types* in M.

Partial and Extended Allocations. We will assume that M is partitioned into a set of *given* items M' and a set of *open* items A. Formally, $M = M' \cup A$ and $M' \cap A = \emptyset$. A *partial allocation* $\gamma = (\gamma_1, \gamma_2, \ldots, \gamma_n)$ is a partition of M' into n (potentially empty) sets, where γ_i is the bundle of given items to agent i. An allocation of open items $\pi = (\pi_1, \pi_2, \ldots, \pi_n)$ partitions A into n (again, potentially empty) sets, where agent i gets bundle π_i. Given a partial allocation γ and an allocation π, we get an *extended allocation* $\gamma \cup \pi$, where agent i gets allocated bundle $\gamma_i \cup \pi_i$.

Fairness Concepts. We will consider three different fairness notions. An extended allocation $\gamma \cup \pi$ is said to be:

- *envy-free*, denoted EF, if for any pair of agents i and j it holds that $v_i(\gamma_i \cup \pi_i) \geq v_i(\gamma_j \cup \pi_j)$;
- *envy-free up to any item*, denoted EFX, if for any pair of agents i and j and any item $a \in \gamma_j \cup \pi_j$ it holds that $v_i(\gamma_i \cup \pi_i) \geq v_i(\gamma_j \cup \pi_j \setminus a)$;
- *envy-free up to one item*, denoted EF1, if for any pair of agents i and j there exists an item $a \in \gamma_j \cup \pi_j$ such that $v_i(\gamma_i \cup \pi_i) \geq v_i(\gamma_j \cup \pi_j \setminus a)$.

Observe that EF1 is a more relaxed fairness notion compared to EFX, which in turn is more relaxed than EF.

Fair Allocation Extension Problems. We are interested in the computation of an allocation π of open items, such that the extended allocation is fair according to some of the above-mentioned criteria; we define the problems for envy-free solutions, but they could naturally be extended for EF1 and EFX. The input to each version of our problem is a set of agents with their valuations and a partial allocation γ. In the first version, termed ENVY-FREE ALLOCATION EXTENSION and already introduced in the Introduction, we do not constrain the extension allocation in any way. In the other two versions of the problem, the set of recipients is restricted in some way. One option is to restrict the set of agents that are allowed to receive an open item under π, which for brevity we denote REFAE.

RESTRICTED ENVY-FREE ALLOCATION EXTENSION	
Input:	A set M of indivisible items, a set N of n agents with additive valuations, a partial allocation $\gamma : M' \to N$ of some subset $M' \subseteq M$ of items, and a set $S \subseteq N$.
Question:	Does there exist an allocation π of the open items such that: (a) $\pi_i = \emptyset$ for every $i \notin S$; (b) $\gamma \cup \pi$ is envy-free?

	Few items k	Few Recipients p	
		REFAE	FEFAE
	W[1]-hard (Theorem 2)	NP-hard, $p \leq 2$ (Theorem 3)	
Agent types n_t	FPT (Theorem 1)	XP, unary valuations (Theorem 4)	
Item types m_t	W[1]-hard (Theorem 2)	FPT (Theorem 5)	XP (Theorem 5)

Table 1: An overview of our results. The W[1]-hardness results of Theorem 2 are complemented by the naive, brute force, XP algorithms. The XP algorithms for REFAE and FEFAE when parameterized by n_t are essentially tight, since matching strong W[1]-hardness and weak paraNP-hardness lower bounds follow from the literature (Bliem, Bredereck, and Niedermeier 2016).

A different option is to restrict just the number of recipients. We term this problem FEFAE.

FREE ENVY-FREE ALLOCATION EXTENSION

Input: A set M of indivisible items, a set N of n agents with additive valuations, a partial allocation $\gamma : M' \to N$ of some subset $M' \subseteq M$ of items, and $p \in \mathbb{N}$.

Question: Does there exist an allocation π of the open items such that: (a) $\pi_i \neq \emptyset$ for at most p agents; (b) $\gamma \cup \pi$ is envy-free?

While we formally study the decision variants of these problems for complexity-theoretic reasons, every algorithm obtained in this article is constructive and can also output a suitable allocation if one exists.

Parameterizing by the Number of Open Items

We start our investigation by considering the complexity of ENVY-FREE ALLOCATION EXTENSION when the number k of open items is included in the parameterization—in other words, we ask under which conditions one can efficiently solve the extension problem for only a few open items.

As our baseline, we observe that the problem admits a trivial XP algorithm: one can enumerate all possible assignments of the open items to agents in n^k time and check whether any of these is envy-free. However, such algorithms are considered highly inefficient in the community (Downey and Fellows 2013; Cygan et al. 2015), and the central question tackled by this section is whether (or under which conditions) one can achieve fixed-parameter tractability.

Before settling the problem when parameterized by k alone, we first provide a fixed-parameter algorithm for ENVY-FREE ALLOCATION EXTENSION when the parameterization also includes the number of agent types. Intuitively, this provides a positive result for the case where there are only a few open items and the agents can be partitioned into a few groups with identical preferences. We remark that this setting is not trivial, as agents which we consider to have identical preferences (e.g., due to polling limitations) can and will often have different pre-assigned items.

The following observation will be useful in the proof and follows directly from the definition of agent types and envy-free allocations.

Observation 1. *Assume we are given an envy-free partial allocation γ such that no pair of agents in agent type X envy each other. If a solution π assigns a set Q of positively valued items to an agent $j \in X$, then π must assign items of the same value as $v_j(Q)$ to* all *other agents in X.*

Theorem 1. ENVY-FREE ALLOCATION EXTENSION *is fixed-parameter tractable when parameterized by the number k of open items plus the number n_t of agent types.*

Proof Sketch. As the very first step, we observe that if there is an agent i envying agent j, then π must assign some open item to i. Hence, we begin by exhaustively branching, i.e., guessing, for each such agent i, which of the at most k items will be assigned to i, and restart our considerations with the instance updated accordingly. For the following, we hence consider that no agent envies another.

We begin by partitioning the set of agent types into *small* and *large* types. An agent type is classed as small when it contains at most k agents. We define Z as the set of agents in all of the small types, the size of Z is upper-bounded by $k \cdot n_t$.

Next, we exhaustively branch to determine the following information about the allocation of the k items to agents. First of all, we branch over all partitions of the open items into bundles, where we will assume that π assigns each of the bundles to distinct agents. Next, for each of the at most k bundles, we branch to either determine which of the agents in Z will receive it, or that it will not be assigned to any agent in Z. The overall branching factor up to this point is upper-bounded by $k^k \cdot (kn_t + 1)^k$.

At this point, we check whether any agent from Z envies another agent; if that is the case, then we can correctly reject the current branch, as we may assume to have precisely guessed the bundles assigned to all agents in Z. Similarly, no pair of agents outside of Z may envy each other due to the exhaustive procedure carried out in the first paragraph of the proof. At this point, we recall that by Observation 1 and the definition of large agent types, each agent outside of Z must receive a bundle that they value as 0. Hence, if an agent outside of Z were to envy an agent in Z, we may also correctly reject the current branch.

At this point, it remains to assign the remaining open items to the agents outside of Z without creating envy. For each of the bundles, we can determine whether assigning it to an agent $i \notin Z$ creates envy from any other agent in the instance (where for agents in Z we assume them to receive the bundles specified in our branching, while for agents outside of Z, we assume that nothing changed). To complete the proof, we construct an auxiliary bipartite graph G where:

- one side contains the set of all remaining bundles (i.e., those not assigned to Z),
- the other side contains the set of all agents outside of Z, and,
- there is an edge between bundle b and agent i if and only if b can be assigned to i without creating envy from any other agent (i.e., bundles of value 0).

To complete the proof, it suffices to check whether G admits a matching that saturates the set of all bundles. The overall running time can be upper-bounded by $(k \cdot n_t)^{\mathcal{O}(k)} \cdot n^2$. □

Turning back to ENVY-FREE ALLOCATION EXTENSION parameterized by k alone, our next result shows that the aforementioned trivial XP algorithm can be viewed as "optimal" in the sense of the problem not admitting any fixed-parameter algorithm under the complexity assumption of W[1] $\not\subseteq$ FPT. Naturally, one would then ask whether fixed-parameter tractability can be achieved at least when the parameterization is enriched by the number m_t of item types—a setting that can be seen as complementary to the one settled in Theorem 1. Surprisingly, we exclude fixed-parameter algorithms for the problem, even in this significantly more restrictive setting via a highly involved reduction.

Theorem 2. ENVY-FREE ALLOCATION EXTENSION *is* W[1]-hard *when parameterized by the number k of open items plus the number m_t of item types.*

Proof Sketch. We reduce from the classical W[1]-hard MULTICOLORED CLIQUE problem: given a q-partite (where each part is assigned a unique color) graph, decide whether the graph contains a clique of size q. Let $\mathcal{I} = ((\mathrm{V}, \mathrm{E}), q)$ be an instance of MULTICOLORED CLIQUE, where q is the number of colors and without loss of generality assume that E contains only edges adjacent to vertices of different colors. Let $\mathrm{V} = \bigcup_{1 \leq k \leq q} \mathrm{V}_k$, where V_k is the set that contains all vertices of color k and $\mathrm{E} = \bigcup_{1 \leq i < j \leq q} \mathrm{E}_{ij}$, where E_{ij} is the set that contains all edges of E that are adjacent to vertices that belong to sets V_i and V_j respectively. We will construct an instance $\mathcal{I}'$ of ENVY-FREE ALLOCATION EXTENSION and prove that an envy-free extension for $\mathcal{I}'$ exists if and only if $\mathcal{I}$ contains a clique of size q.

Construction: For each vertex set V_i we assume a *vertex-agent group* $V_i = \{\alpha^i_j | j \in \mathrm{V}_i\}$ and for each edge set E_{ij} an *edge-agent group* $E_{ij} = \{e^{ij}_\ell | \ell \in \mathrm{E}_{ij}\}$. For each vertex-agent group V_i, the construction will make use of three pre-allocated item types $\{\square^i, \triangle^i, \bigstar^i\}$. Each edge-agent group will also make use of three pre-allocated item types $\{\square^{ij}, \triangle^{ij}, \bigstar^{ij}\}$. The pre-allocation γ will only provide agents in vertex-agent group V_i with some combination of items from "their" types $\{\square^i, \triangle^i, \bigstar^i\}$, while agents in edge-agent group E_{ij} will receive items not only from "their" types $\{\square^{ij}, \triangle^{ij}, \bigstar^{ij}\}$ but also items of the types $\{\square^i, \triangle^i, \square^j, \triangle^j\}$.

Below, we provide the details of how the partial allocation γ is constructed. We assume that the agents in each group are provided in an arbitrary order, i.e., the agents in the vertex group V_i have the form $\alpha_1, \ldots, \alpha_{|V_i|}$ (and analogously for agents in edge groups).

Partial allocation γ. Each agent α_x within a vertex-agent group V_i holds a bundle consisting of (1) one copy of $\bigstar^i$, (2) x copies of $\square^i$, and (3) $2|V_i|^2 - x^2 - x$ copies of $\triangle^i$. Each agent η_z within an edge-agent group E_{ij} representing the edge between $\alpha_x \in V_i$ and $\alpha_y \in V_j$ holds a bundle consisting of (1) one copy of $\bigstar^{ij}$, (2) z copies of $\square^{ij}$, (3) $2|V_i|^2 - z^2 - z$ copies of $\triangle^{ij}$, and moreover (4) precisely the same number of $\square^i$ ($\square^j$) and $\triangle^i$ ($\triangle^j$) items as α_x (α_y).

The instance will furthermore contain precisely $q + \binom{q}{2}$ pairwise-distinct open items: $s_1, \ldots, s_q$ (one for each vertex agent group) and $t_{12}, t_{13}, \ldots, t_{(q-1)q}$ (one for each edge agent group). This leaves the total number of item types at $4q + 4\binom{q}{2}$. Intuitively, while the initial allocation γ is envy-free, the constructed instance of ENVY-FREE ALLOCATION EXTENSION will force us to allocate the open items in a way where (1) each vertex and each edge agent group receives precisely one open item, and (2) to prevent envy, the recipient in the edge agent group must represent an edge such that both incident vertex-agents are recipients as well. However, to achieve this, we require very careful calibration of γ and the specific valuations of the agents; in fact, each agent will have an entirely unique valuation function.

The valuation function. Each agent α_x within a vertex-agent group V_i values: (1) $\triangle^i$ as 1, (2) $\square^i$ as $2x + 1$, (3) $\bigstar^i$ as 0. This describes the valuation for items allocated to α_x by γ. Next, for the item types that are not pre-assigned to α_x, the agent uses the following valuation: (1) each $\triangle^\circ$ and each $\square^\circ$ for $\circ \neq i$ has a value of 0, (2) each $j \neq i$, $\bigstar^{ij}$ and/or $\bigstar^{ji}$ has a value of 0, and (3) each $\bigstar^\circ$ whose valuation is not defined up to now has a value that is precisely equal to the value of α_x's bundle as assigned by γ.

We observe that the construction at this point ensures that α_x values its own bundle identically to that of any agent from a different vertex agent group. It is less obvious—but provable—that α_x strictly prefers their own bundle to the one of every other agent in their own vertex agent group. Moreover, α_x values its own bundle identically to that of edge agents that are incident to α_x in the graph, while for the other agents in those edge agent groups it prefers its own bundle.

For the open items, α_x values s_i as 1, the items t_{ij} and/or t_{ji} for all $j \neq i$ as 1, and all remaining open items as 0. Each agent η_z within an edge-agent group E_{ij} values: (1) $\triangle^{ij}$ as 1, (2) $\square^{ij}$ as $2z + 1$, (3) $\bigstar^{ij}$ as well as all remaining items included in its bundle by γ as 0. Next, for the item types that are not pre-assigned to η_z, the agent uses the following valuation: (1) each $\triangle^\circ$ and each $\square^\circ$ for $\circ \neq ij$ has a value of 0, and (2) each $\bigstar^\circ$ where $\circ \neq ij$ has a value that is precisely equal to the value of η_z's bundle as assigned by γ.

We observe that by construction, η_z will value its bundle exactly the same as as that of every agent outside of its own group. Moreover, similarly as before, it is possible to prove that η_z strictly prefers their own bundle to the one of every other agent in their own edge agent group. Finally, for the open items, η_z values its "own" open item t_{ij} as 1 and all others as 0.

To complete the proof, two tasks remain. First, establish that the numbers are indeed set up in a way to ensure that each agent strictly prefers their own bundle to that of other agents in their own group. The second is then to show that every q-clique in the initial graph corresponds to an envy-free allocation in the constructed instance following the intuition provided at the end of the partial allocation γ. ☐

In summary, our results at this point prove that when extending allocations with k open items, having a small number of agent types helps achieve tractability but having a small number of item types does not.

Parameterizing by the Number of Recipients

In this section, we address the more complex situation of needing to allocate (a possibly large number of) open items to at most p recipients. As mentioned in the Introduction, here it will be important to distinguish whether the set of these recipients is fixed and provided on the input (REFAE), or whether the task also includes the identification of this set (FEFAE). Moreover—and unlike in the case studied in the previous section—the existence of efficient algorithms will sometimes depend on whether we may assume the valuations to be encoded in unary, or whether they are encoded in binary. It will be useful to recall that lower bounds achieved in the former (latter) setting are called *strong* (*weak*).

As regards the complexity of these problems when parameterized by p alone, it is easy to observe that both problems are weakly NP-hard already when $p \leq 2$ as this directly generalizes both the previously-studied setting of all items being open (Bliem, Bredereck, and Niedermeier 2016) and SUBSET SUM, both of which are weakly NP-hard. On the other hand, the problem is trivial for $p = 1$. Below, we show that in the extension setting studied here, NP-hardness holds for $p = 2$ even in the unary (i.e., strong) setting, contrasting the known existence of an XP algorithm for that case when all items are open (Bliem, Bredereck, and Niedermeier 2016).

Theorem 3. *Both* REFAE *and* FEFAE *are strongly* NP-hard *even when* $p \leq 2$.

Proof Sketch. We reduce from INDEPENDENT SET: Given an instance $\mathcal{I}$ consisting of a graph (V, E) and an integer ℓ, decide whether the graph contains an independent set of size precisely ℓ. We construct an instance $\mathcal{I}'$ of ENVY-FREE ALLOCATION EXTENSION and prove that an envy-free allocation extension for $\mathcal{I}'$ with only two recipients exists if and only if instance $\mathcal{I}$ contains an independent set of size ℓ.

Construction: We assume a set $N = \{1, \cdots, |E|+1, |E|+2\}$ of $|E|+2$ agents, consisting of a corresponding agent for each edge in E and two additional agents labeled for simplicity as $|E|+1$ and $|E|+2$. We assume a set of items partitioned into the sets $G = \{g_1, \cdots, g_{|E|+1}, g_{|E|+2}\}$, that contains one item for each agent in N and the set A $= \{a_1, \cdots, a_{|V|}\}$, that contains one item for each vertex of $\mathcal{I}$. Finally, we assume the partial allocation γ, where for each item $g_i \in G$ and for each agent $i \in N$, $\gamma(g_i) = i$, i.e, agent i's bundle in the partial allocation γ is item g_i, while the items in A are *open* and yet-to-be allocated.

The valuations for each agent $i \in N \setminus \{|E|+1, |E|+2\}$ are defined as: (*i.*) $v_i(g_i) = |V|$; (*ii.*) $v_i(g_{|E|+1}) = |V| - 1$ and $v_i(g_{|E|+2}) = 0$; (*iii.*) $v_i(a_j) = 1$, if the vertex j that corresponds to item a_j is adjacent to the edge represented by agent i, else 0.

The valuation for agent $|E|+1$, is defined as : (*i.*) $v_{|E|+1}(g_{|E|+1}) = |V| - 2\ell$; (*ii.*) $v_{|E|+1}(g) = 0, \forall g \in G \setminus \{g_{|E|+1}\}$; (*iii*) $v_{|E|+1}(a_j) = 1, \forall a_j \in A$;

Finally, the valuation for agent $|E|+2$ is defined as: (*i.*) $v_{|E|+2}(g_{|E|+2}) = |V|$ and $v_{|E|+2}(g_{|E|+1}) = 0$ (*ii.*) $v_{|E|+2}(g) = 2|V| - \ell$, for all $g \in G \setminus \{g_{|E|+1}, g_{|E|+2}\}$ (*iii.*) $v_{|E|+2}(a_j) = 1, \forall a_j \in A$;

To complete the proof, it suffices to show that (V, E) contains an independent set Q of size ℓ if and only if allocating the open items corresponding to Q to agent $|E|+1$ and all remaining open items to agent $|E|+2$ is a solution for $\mathcal{I}'$. To see that this holds, we note that if the items corresponding to both endpoints of an edge are allocated to the agent $|E|+1$, then $|E|+1$ would be envied by the corresponding edge agent. Moreover, if the number of items assigned to $|E|+2$ is less than $|V| - \ell$, then the agent $|E|+2$ would envy all of the edge-agents (possibly except for $|E|+1$). Finally, if the number of items assigned to $|E|+1$ is less than ℓ, then $|E|+1$ would envy $|E|+2$ in view of the previous sentence. ☐

We remark that Theorem 3 is tight in the sense that both problems are also in NP. Indeed, in both cases one can verify whether a provided complete allocation is envy-free in polynomial time. Having settled the intractability of REFAE and FEFAE w.r.t. p alone, we now ask whether one can achieve tractability at least in settings with a small number of agent or item types.

We begin by considering the former parameterization, which was the one that yielded fixed-parameter tractability for the setting investigated in the previous section. Here, we can immediately exclude fixed-parameter tractability for all of the problem variants considered in this section due to the aforementioned fact that finding an envy-free allocation of a set of open items to agents is known to be strongly W[1]-hard and weakly paraNP-hard when parameterized by the number of agents (Bliem, Bredereck, and Niedermeier 2016). Below, we at least show that both problems admit XP algorithms in the unary-valuation case, complementing the former lower bound.

Theorem 4. *If the valuations are encoded in unary, both* REFAE *and* FEFAE *are in* XP *when parameterized by the number* p *of recipients plus the number* n_t *of agent types.*

Proof Sketch. For FEFAE, we start by branching to determine which of the at most p agents will be the recipients, requiring a branching factor of at most n^p. For REFAE, we skip this step.

At this point, we now perform a dynamic programming subroutine for an input instance $\mathcal{I}$ and then argue that this subroutine, in fact, solves our problem. For the subroutine, we construct a table which stores the following information: for each recipient $i \in [p]$ and each agent type Z, we store the valuation of the bundle assigned to i from the perspective of agents of type Z. We note that each table entry consists of $p \cdot n_t$ unary-encoded values. We assume that the items are processed in an arbitrary but fixed order, and for each item we update the table by exhaustively assigning it to each of the possible recipients and updating the table accordingly.

At the end of this subroutine, we loop through each of the at most $|\mathcal{I}|^{\mathcal{O}(n_t \cdot p)}$ entries in the table, and notice that each such entry provides us with complete information about how each agent values each bundle in $\mathcal{I}$ in all possible assignments corresponding to that table entry; hence, it suffices to check, in polynomial time, whether at least one such table entry results in an envy-free assignment in $\mathcal{I}$. ☐

Surprisingly, when dealing with a small number of recipients, we prove that it is the number m_t of item types which yields better tractability results for the considered problems—a situation that is entirely opposite to that of the previous section. In particular, when parameterizing by $p + n_t$ we obtain a fixed-parameter algorithm for REFAE and an XP algorithm for FEFAE. We remark that the latter result is the best one could have at this point hoped for due to Theorem 2 ruling out fixed-parameter algorithms even in a strictly more restrictive setting.

Theorem 5. *When parameterized by p plus the number m_t of item types,* REFAE *and* FEFAE *are fixed-parameter tractable and* XP*-tractable, respectively.*

Proof Sketch. For FEFAE, we start again by branching to determine which of the at most p agents will be the recipients, requiring a branching factor of at most n^p. At this point, we can solve both problems via an encoding into an ILP using a number of variables that depend purely on the parameters. In particular, for each of the at most p recipients and each of the m_t item types, we construct a dedicated variable which will determine precisely how many items of this type will be assigned to the given recipient. The valuation of each agent can then be easily captured as a linear expression of the variables, and the requirement that the final allocation is envy-free can be captured via linear constraints. Since the number of variables is upper-bounded by $p \cdot m_t$, we can solve the resulting instance in time $(p \cdot m_t)^{\mathcal{O}(p \cdot m_t)} \cdot |\mathcal{I}|^{\mathcal{O}(1)}$ (Lenstra, Jr. 1983; Dadush 2012). ☐

Extending Allocations Beyond Envy-Freeness

While it seems that the algorithmic upper and lower bounds could be translated also to the analogous problems of extending EF1 and EFX allocations, in this section we turn towards a different, highly studied aspect of such allocations—specifically, the question of whether such allocations are guaranteed to exist (Budish 2011; Caragiannis et al. 2019). We show that in the extension setting, this question can be completely resolved based on the level of fairness we assume in the partial allocation.

First of all, it is obvious that if no fairness guarantees are provided for the partial allocation, then one cannot hope to guarantee an extension to an EF1 or EFX allocation (since, e.g., there might not be any open items left at all). On the other hand, if we assume the partial allocation to be envy-free, we obtain the following.

Proposition 1. *Every envy-free partial allocation can be extended to an* EF1 *allocation. At the same time, there exists an envy-free partial allocation of items to* 2 *agents which cannot be extended to an* EFX *allocation.*

Proof Sketch. To extend an envy-free allocation to EF1, we show that it is possible to compute an allocation π of open items such that the final allocation consisting of π and the partial allocation γ is EF1. This is done by executing a standard algorithm for computing EF1 allocations (Envy Cycle Elimination (Lipton et al. 2004)) on top of the partial allocation γ. For the second part of the claim, consider two agents with identical valuations such that γ assigns each an item that each of them values 1. There are three open items to be allocated that the agents value 3, 4, and 9 respectively. ☐

Interestingly, the latter non-existence result contrasts with the known fact that EFX allocations from scratch always exist for 2 agents (Plaut and Roughgarden 2020). The final remaining question is whether a partial EFX allocation extends to a complete EF1 allocation. We resolve this in the negative below.

Proposition 2. *There exists an* EFX *partial allocation to* 2 *agents which cannot be extended to an* EF1 *allocation.*

Proof Sketch. Consider the following instance with two agents and three items x, y, z. The valuations of the agents are as follows: $v_1(x) = 10, v_1(y) = 0, v_1(z) = 1$; $v_2(x) = 0, v_2(y) = 10, v_2(z) = 1$. Consider now the partial allocation that gives x to agent 2 and y to agent 1. Trivially, this is an EFX allocation since envy can be eliminated by removing the only item the other agent received. Now, no matter which agent receives item z, the constructed solution will not be EF1. ☐

Concluding Remarks

Our paper initiates the study of fairly extending partial allocations of indivisible items and the frontiers of tractability for this problem. The presented results showcase that the complexity of this task varies as it strongly depends on the chosen parameters—with the number of agent types and item types each leading to tractability in different settings. This naturally gives rise to the question of what other parameterizations can yield fixed-parameter algorithms for this natural but previously overlooked problem. For instance, can one obtain such algorithms when simultaneously parameterizing by the number m_t of item types *and* the number n_t of agent types? If this combination of parameters is not sufficient for tractability, can the inclusion of the number of recipients, p, in addition to n_t and m_t lead to tractability?

A different direction that has proven fruitful recently, is to consider the problem when agents form a social network (Bredereck, Kaczmarczyk, and Niedermeier 2022; Eiben et al. 2023a). In this setting, each agent compares their bundle against the bundles of their “friends”. What are the graph structures on the social network that make the problem tractable?

Acknowledgments

Argyrios Deligkas and Stavros D. Ioannidis acknowledge the support of the EPSRC grant EP/X039862/1. Robert Ganian acknowledges support from the Austrian Science Fund (FWF, projects 10.55776/Y1329 and 10.55776/COE12) and the Vienna Science Foundation (WWTF, project 10.47379/ICT22029).

References

Amanatidis, G.; Markakis, E.; and Ntokos, A. 2020. Multiple birds with one stone: Beating 1/2 for EFX and GMMS via envy cycle elimination. *Theoretical Computer Science*, 841: 94–109.

Angelini, P.; Battista, G. D.; Frati, F.; Jelínek, V.; Kratochvíl, J.; Patrignani, M.; and Rutter, I. 2015. Testing Planarity of Partially Embedded Graphs. *ACM Trans. Algorithms*, 11(4): 32:1–32:42.

Arroyo, A.; Derka, M.; and Parada, I. 2019. Extending Simple Drawings. In Archambault, D.; and Tóth, C. D., eds., *Graph Drawing and Network Visualization - 27th International Symposium, GD 2019, Prague, Czech Republic, September 17-20, 2019, Proceedings*, volume 11904 of *Lecture Notes in Computer Science*, 230–243. Springer.

Bhore, S.; Ganian, R.; Khazaliya, L.; Montecchiani, F.; and Nöllenburg, M. 2023. Extending Orthogonal Planar Graph Drawings Is Fixed-Parameter Tractable. In Chambers, E. W.; and Gudmundsson, J., eds., *39th International Symposium on Computational Geometry, SoCG 2023, June 12-15, 2023, Dallas, Texas, USA*, volume 258 of *LIPIcs*, 18:1–18:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Bliem, B.; Bredereck, R.; and Niedermeier, R. 2016. Complexity of Efficient and Envy-Free Resource Allocation: Few Agents, Resources, or Utility Levels. In Kambhampati, S., ed., *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, 102–108. IJCAI/AAAI Press.

Bouveret, S.; and Lang, J. 2008. Efficiency and Envy-freeness in Fair Division of Indivisible Goods: Logical Representation and Complexity. *J. Artif. Intell. Res.*, 32: 525–564.

Brânzei, S.; Lv, Y.; and Mehta, R. 2016. To Give or Not to Give: Fair Division for Single Minded Valuations. In Kambhampati, S., ed., *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, 123–129. IJCAI/AAAI Press.

Bredereck, R.; Kaczmarczyk, A.; and Niedermeier, R. 2022. Envy-free allocations respecting social networks. *Artif. Intell.*, 305: 103664.

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

Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Economics and Comput.*, 7(3): 12:1–12:32.

Cygan, M.; Fomin, F. V.; Kowalik, L.; Lokshtanov, D.; Marx, D.; Pilipczuk, M.; Pilipczuk, M.; and Saurabh, S. 2015. *Parameterized Algorithms*. Springer. ISBN 978-3-319-21274-6.

Dadush, D. 2012. *Integer Programming, Lattice Algorithms, and Deterministic Volume Estimation*. Ph.D. thesis, Georgia Institute of Technology.

Deligkas, A.; Eiben, E.; Ganian, R.; Hamm, T.; and Ordyniak, S. 2021. The Parameterized Complexity of Connected Fair Division. In Zhou, Z., ed., *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, 139–145. ijcai.org.

Downey, R. G.; and Fellows, M. R. 2013. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer. ISBN 978-1-4471-5558-4.

Eiben, E.; Ganian, R.; Hamm, T.; Klute, F.; and Nöllenburg, M. 2020. Extending Partial 1-Planar Drawings. In Czumaj, A.; Dawar, A.; and Merelli, E., eds., *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, 43:1–43:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Eiben, E.; Ganian, R.; Hamm, T.; and Ordyniak, S. 2023a. Parameterized complexity of envy-free resource allocation in social networks. *Artif. Intell.*, 315: 103826.

Eiben, E.; Ganian, R.; Kanj, I.; Ordyniak, S.; and Szeider, S. 2023b. On the parameterized complexity of clustering problems for incomplete data. *J. Comput. Syst. Sci.*, 134: 1–19.

Feldman, M.; Mauras, S.; and Ponitka, T. 2024. On Optimal Tradeoffs between EFX and Nash Welfare. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, 9688–9695. AAAI Press.

Ganian, R.; Kanj, I. A.; Ordyniak, S.; and Szeider, S. 2018. Parameterized Algorithms for the Matrix Completion Problem. In Dy, J. G.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, 1642–1651. PMLR.

Ganian, R.; Ordyniak, S.; and Rahul, C. S. 2019. Group Activity Selection with Few Agent Types. In Bender, M. A.; Svensson, O.; and Herman, G., eds., *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*, volume 144 of *LIPIcs*, 48:1–48:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Igarashi, A.; Kamiyama, N.; Suksompong, W.; and Yuen, S. M. 2024. Reachability of Fair Allocations via Sequential Exchanges. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence*, 9773–9780. AAAI Press.

Koana, T.; Froese, V.; and Niedermeier, R. 2021. Binary Matrix Completion Under Diameter Constraints. In Bläser, M.; and Monmege, B., eds., *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16-19, 2021, Saarbrücken, Germany (Virtual Conference)*, volume 187 of *LIPIcs*, 47:1–47:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Koana, T.; Froese, V.; and Niedermeier, R. 2023. The complexity of binary matrix completion under diameter constraints. *J. Comput. Syst. Sci.*, 132: 45–67.

Lenstra, Jr., H. W. 1983. Integer Programming with a Fixed Number of Variables. *Mathematics of Operations Research*, 8(4): 538–548.

Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, 125–131.

Nguyen, T. T.; and Rothe, J. 2023. Complexity Results and Exact Algorithms for Fair Division of Indivisible Items: A Survey. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*, 6732–6740. ijcai.org.

Plaut, B.; and Roughgarden, T. 2020. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2): 1039–1068.

Prakash, V.; Igarashi, A.; and Vaish, R. 2024. Fair and Efficient Completion of Indivisible Goods. *CoRR*, abs/2406.09468.