Fair Allocation with Initial Utilities

Niclas Boehmer, Luca Kreisel

Abstract

The problem of allocating indivisible resources to agents arises in a wide range of domains, including treatment distribution and social support programs. An important goal in algorithm design for these problems is fairness, where the focus in previous work has been on ensuring that the computed allocation provides equal treatment to everyone. However, this perspective disregards that agents may start from unequal initial positions, which is crucial to consider in settings where fairness is understood as *equality of outcome*. In such settings, the goal is to create an equal final outcome for everyone by leveling initial inequalities through the allocated resources. To close this gap, focusing on agents with additive utilities, we extend the classic model by assigning each agent an initial utility and study the existence and computational complexity of several new fairness notions following the principle of equality of outcome. Among others, we show that complete allocations satisfying a direct analog of envy-freeness up to one resource (EF1) may fail to exist and are computationally hard to find, forming a contrast to the classic setting without initial utilities. As our main contribution, we propose a new, always satisfiable fairness notion, called minimum-EF1-init and design a polynomial-time algorithm based on an extended round-robin procedure to compute complete allocations satisfying this notion.

1 Introduction

Allocating resources to agents is a central problem in algorithmic decision-making, with applications ranging from task and house allocation to treatment distribution and scheduling [45, 37, 26, 8, 36, 1, 22]. Thus, the problem has been widely studied across artificial intelligence, machine learning, multi-agent systems, operations research, and computational social choice. Much of this literature focuses on the design and analysis of algorithms that compute allocations satisfying formal fairness criteria, for instance, envy-freeness [2, 11, 35].

Both outside and inside computer science, the discourse on fairness distinguishes between two paradigms: On the one hand, *equality of opportunity* (also known as *formal equality* in law) demands equal treatment for everyone regardless of individual circumstances [21, 28]. On the other hand, *equality of outcome* (also known as *substantive equality* in law) acknowledges the existence of initial disparities and seeks to create an equal final outcome for everyone by leveling the existing inequalities [7, 36]. These two principles stand in an inherent conflict with each other, as countering initial disparities oftentimes necessitates the prioritization of those starting from an inferior position, contradicting the principle of equal treatment.[1]

In algorithmic fair allocation, an overwhelming majority of works adhere to equality of opportunity: They assess fairness by how the allocation distributes resources—irrespective of agents starting positions [2, 11, 35]. For example, under the classic notion of envy-freeness, an allocation is considered fair if no agent prefers the bundle allocated to another agent to their own. However, this ignores initial inequalities prevalent in many application areas of fair allocation. For example, in medical resource allocation programs, patients have different initial probabilities of recovery; educational interventions serve students with varying levels of prior knowledge or preparedness; and more generally, support programs often address individuals from diverse backgrounds and starting positions [17, 20, 22].

Motivated by the prevalence of such disparities, this paper initiates the study of equality of outcome in algorithmic fair allocation. In many real-world allocation problems, decision makers aim for equal

[1] This conflict recently surfaced in an executive order of US President Trump [40].

post-allocation outcomes instead of equal treatment [33, 39]. One example is the Fairer Scotland Duty [39], which places a statutory responsibility on certain public bodies to consider how their decisions can help reduce inequalities of outcome. Notably, existing, intensively studied formal models of fairness reflecting equality of opportunity fail to capture this goal of equality of outcome, a gap that we address in this work.

To give another concrete motivating example, consider the following real-world case: the Indian NGO ARMMAN runs large-scale mobile health programs for new mothers, enrolling millions across the country [3]. To boost program engagement, ARMMAN can make a limited number of weekly calls to participants, a sequential resource allocation problem [41]. Crucially, ARMMAN seeks to prioritize calls to mothers with lower initial engagement, often from historically marginalized communities: Rather than allocating calls proportionally across demographics, their goal is to ensure that engagement levels between different demographic groups are equalized post-intervention [42], thus following the principle of equality of outcome.

1.1 Our Contributions

We initiate the formal study of fair allocation for agents with initial disparities, focusing on fairness notions aligned with the principle of equality of outcome. We do so by equipping each agent i with an agent-specific *initial utility* value $b_i \in \mathbb{R}_{\geq 0}$ and measuring fairness in terms of the utility the agent derives from their bundle plus b_i. Throughout, we restrict attention to additive utilities, as is standard in the literature [2].

In Section 3, we present straightforward adaptations of classical fairness notions – envy-freeness (EF) and envy-freeness up to one resource (EF1) – which we term EF-init and EF1-init, respectively. Concretely, for EF-init (EF1-init), we require for every pair of agents i and j that agent i's initial utility plus i's utility for their own bundle is at least as large as agent j's initial utility plus agent i's utility for j's bundle (after removing a resource). Our results reveal that introducing initial utilities fundamentally changes the nature of the problem: even in simple cases with identical resources, complete EF1-init allocations may fail to exist, and deciding whether such an allocation exists is NP-complete. This stands in stark contrast to the classical setting, where complete EF1 allocations are always guaranteed and efficiently computable.

Nevertheless, we show that for a constant number of agents (under unary encoding), the existence of a complete EF-init and EF1-init allocations can be decided efficiently using dynamic programming. Furthermore, for the special case of identical resources, we present a polynomial-time algorithm for deciding the existence of EF-init allocations. From a practical standpoint, the allocation of identical resources is a natural and frequently encountered scenario, for example when allocating treatments or interventions (e.g., the allocation task faced by ARMMAN described above). From a theoretical perspective, our result highlights the added complexity introduced by initial utilities: while in the classical setting without initial utilities any envy-free allocation simply assigns the same number of resources to each agent, in our model, ensuring fairness requires a technically involved dynamic programming approach and a meticulous analysis of the interplay between initial utilities and allocated bundles.

The non-existence of complete EF1-init allocations highlights the need for fundamentally new approaches. In Section 4, we therefore search for an envy-based fairness criterion that is always satisfiable. As our main contribution, we present minimum-EF1-init, a new envy-based fairness notion for which a satisfying allocation is guaranteed to exist. The core idea is to relax EF1-init when evaluating whether an agent i with higher initial utility envies an agent j with lower initial utility. Instead of comparing the sums of the initial utilities and i's values for the bundles, we take a different approach: We first allow j to level the initial utility difference between i and j with a subset X^* of j's resources which, from j's perspective, equalizes this difference. Then, we compare i's bundle to j's bundle without X^* under

EF1. We show that minimum-EF1-init coincides with EF1-init in settings where resources' usefulness is diminishing in initial utility. We develop an adaptation of the classic round-robin algorithm for the setting with initial utilities that computes an allocation satisfying minimum-EF1-init. In our version, agents still select their most preferred available resource in rounds, but participation in each round is restricted: Initially, only agents with the lowest initial utility are present, and additional agents are added to the picking order as soon as all present agents have reached their initial utility. While our algorithm is simple to state, making it an appealing distribution algorithm in and of itself, the proof that it guarantees minimum-EF1-init is more involved, requiring a careful analysis of the position where newly added agents are inserted into the picking order.

The (complete) proofs of all statements can be found in the appendix.

1.2 Related Work

There is a rich body of work on the formal analysis of fair allocation of indivisible resources, considering various preference models and fairness concepts [2, 11, 35]. Our work aligns with two major strands in this literature. First, alongside share-based notions such as the maximin share [16, 32], envy-based concepts–as considered here–have arguably received the most attention [2, 34, 13]. Second, additive utility functions constitute the standard preference model in this domain. They are especially common in initial work on new allocation settings—such as ours—due to their relative simplicity and expressive power [14, 2, 11].

While fair allocation has been studied extensively, there are, to the best of our knowledge, only two works that extend fair allocation models to account for generally differing starting positions prior to the allocation process [38, 18]. Both focus on the following problem: given a partial initial allocation, can it be extended to a complete allocation satisfying a specified, traditional fairness criterion? They analyze the (parameterized) computational complexity of this extension problem under various preference models and fairness notions, highlighting its general intractability. Our work relates to theirs in that our model can be viewed as a special case of theirs, where all agents assign identical utilities to the initially allocated resources and each agent holds at most one resource in the initial allocation. As a result, some of their positive complexity results for restricted cases carry over to our setting (see Appendix B for a detailed discussion). However, our study fundamentally differs in two key aspects: First, instead of focusing on computational aspects of existing fairness axioms, we design new fairness notions (following the principle of equality of outcome) tailored to the setting with initial utilities. Second, by putting a focus on the case of initial utilities, we are able to obtain stronger axiomatic guarantees more relevant to our setting. Our negative results can be interpreted as strengthened versions of analogous results in the setting considered by Prakash HV et al. [38] and Deligkas et al. [18].

Moreover, our work connects to the study of fair allocation with subsidies [27], in which an additional divisible resource can be used to make an allocation of indivisible resources envy-free, and to lines of work on fair allocation with agents with different entitlements [23, 15] and budget constraints [43, 6]. More broadly, our work fits into the study of completion problems in computational social choice, such as possible and necessary winner problems in voting [44] and stable matching [19, 5].

2 Preliminaries

Fair Allocation of Indivisible Resources We consider the problem of allocating a set R of m *indivisible resources* among a set $A := [n]$ of $n \geq 2$ *agents*. We refer to subsets $X \subseteq R$ of resources as *bundles* and denote by 2^R the set of all bundles. Each agent $i \in A$ has a *utility function* $u_i : 2^R \to$

$\mathbb{R}_{\geq 0}$, where $u_i(\emptyset) = 0$.[2,3] We assume throughout that utility functions are *additive*, i.e., $u_i(X) = \sum_{r \in X} u_i(\{r\})$ for any $i \in A$ and $X \subseteq R$. In certain cases, we consider *identical resources*, where $u_i(\{r\}) = u_i(\{r'\})$ for all $i \in A$ and resources $r, r' \in R$. An *allocation* $\mathcal{X}$ is a tuple of n disjoint bundles $(X_1, \ldots, X_n)$ such that $X_i \cap X_j = \emptyset$ for all distinct $i, j \in A$, where X_i is *allocated/assigned* to agent $i \in A$. An allocation $\mathcal{X}$ is *complete* if $\bigcup_{i \in A} X_i = R$, which can be viewed as a very moderate efficiency criterion. We study the problem of finding fair and complete allocations. Two popular fairness notions for indivisible resources are *envy-freeness* (EF) and its relaxation, *envy-freeness up to one resource* (EF1). An allocation $\mathcal{X}$ is EF if $u_i(X_i) \geq u_i(X_j)$ for each $i, j \in A$. It is EF1 if, for every pair of agents $i, j \in A$, there is a resource $r \in X_j$ such that $u_i(X_i) \geq u_i(X_j \setminus \{r\})$ or $X_j = \emptyset$. For a variant of envy-freeness, we say that an agent $i \in A$ envies another agent $j \in A$ *under* this variant, if the pair $i, j \in A$ violates the conditions of the variant, e.g., i envies j under EF if $u_i(X_i) < u_i(X_j)$.

Initial Utilities We study the new problem of fairly allocating indivisible resources in the presence of *initial utilities*. That is, we assume that every agent $i \in A$ has an initial utility $b_i \in \mathbb{R}_{\geq 0}$. Thus, after the allocation $\mathcal{X}$, agent i attains utility $b_i + u_i(X_i)$. Conceptually, we assume that initial utilities are common knowledge and comparable across agents, i.e., unlike with bundles to which two agents might assign different utilities, the agents agree on their initial utility values. The classical setting is recovered by taking $b_i = 0$ for all $i \in A$. We group agents by their initial utilities as follows:

Definition 2.1. *The agents are partitioned into* $t \in [n]$ levels $L_1, \ldots, L_t$, *where each level* L_h *with* $h \in [t]$ *contains all agents with the same initial utility. The levels are indexed such that* $b_i < b_j$ *for any* $i \in L_h$, $j \in L_{h'}$, *and* $h < h'$.

3 A First Attempt at Adapting Envy Notions to Initial Utilities

In this section, we present first adaptations of the classical fairness notions of *envy-freeness* and its relaxation, *envy-freeness up to one resource*, to the setting with initial utilities. Following the equality of outcome principle, we are interested in measuring the fairness of an allocation $\mathcal{X}$ in terms of the total utility of agents after resources have been allocated, i.e., their initial utility plus the utility they derive from $\mathcal{X}$. Following this rationale, we adapt EF and EF1 as follows:

Definition 3.1 (EF-init)**.** *An allocation* $\mathcal{X}$ *is* EF-init*, if for every pair of agents* $i, j \in A$ *either* $X_j = \emptyset$[4] *or it holds that* $b_i + u_i(X_i) \geq b_j + u_i(X_j)$.

Definition 3.2 (EF1-init)**.** *An allocation* $\mathcal{X}$ *is* EF1-init*, if for every pair of agents* $i, j \in A$ *either* $X_j = \emptyset$ *or there exists a resource* $r \in X_j$ *such that* $b_i + u_i(X_i) \geq b_j + u_i(X_j \setminus \{r\})$.

In the classical setting without initial utilities, a complete EF1 allocation always exists, while a complete EF allocation may not exist, even with two agents and a single resource. In contrast, with initial utilities, even complete EF1-init allocations fail to exist:

Observation 3.3. *There exists an instance with two agents and identical resources in which no complete allocation satisfies EF1-init.*

Proof. Consider four identical resources R and two agents: agent j with $b_j = 1$ and $u_j(\{r\}) = 3$ for all $r \in R$, and agent i with $b_i = 10$ and $u_i(\{r\}) = 10$ for all $r \in R$. In any complete allocation where i

[2]To rule out trivial edge cases, we assume that for every agent $i \in A$, there exists some resource $r \in R$ such that $u_i(\{r\}) > 0$.

[3]Regarding the encoding of utility functions, all our hardness results hold even if utility functions are encoded in unary. If not stated otherwise, our positive results hold for both unary- and binary-encoded utility functions.

[4]This condition implies that an agent $i \in A$ with $b_i + u_i(X_i) < b_j$ does not envy an agent j under EF-init or EF1-init if $X_j = \emptyset$, as otherwise achieving an envy-free allocation would be impossible in settings with large initial utility disparities.

receives at least one resource, j has to get at least three resources, as otherwise j envies i under EF1-init. However, in allocation $\mathcal{X}$ where j receives three resources and i receives only one resource, i envies j under EF1-init, as $b_i + u_i(X_i) = 10 + 10 < 1 + 2 \cdot 10 = b_j + u_i(X_j \setminus \{r\})$ for any $r \in X_j$. $\square$

Given the non-existence of complete EF-init and EF1-init allocations, we analyze the complexity of determining whether a given instance with initial utilities admits such an allocation. We refer to these computational problems as EF-init Existence and EF1-init Existence, respectively.

EF1-init Existence We begin with EF1-init allocations, noting that, in the classical setting, EF1 allocations can be found via a simple round-robin algorithm in polynomial time. However, with initial utilities, the problem becomes computationally hard via a reduction from Graph Coloring [24]:

Theorem 3.4. *EF1-init Existence is NP-complete.*

EF1-init Existence/EF-init Existence for Few Agents On the positive side, following a dynamic programming approach, deciding the existence of a complete EF1-init or EF-init allocation becomes tractable when the number of agents is constant:

Proposition 3.5. *For a unary encoding of the utility functions and a constant number of agents, EF1-init Existence and EF-init Existence are polynomial-time solvable.*

EF-init Existence The complexity of deciding the existence of a (complete) envy-free allocation (without initial utilities) is well understood, and the problem is known to be hard, even for restrictive settings such as if there are only two agents with identical preferences or agents have 0/1 utilities for each resource [10]. By setting $b_i = 0$ for all $i \in A$, these results directly carry over to EF-init Existence. We therefore focus on the special case where all resources are identical, obtaining the following result:

Theorem 3.6. *EF-init Existence for identical resources can be decided in $\mathcal{O}(n^2 \cdot m^3)$ time and $\mathcal{O}(n \cdot m^2)$ space.*

To solve this problem, we examine the partitioning of agents by their initial utility into t levels, as introduced in Definition 2.1. Note that in any EF-init allocation, all agents in a level need to get the same number of resources, as all resources are identical. The following observation is crucial to prove the above result. If there are two agents $i, j \in A$ with $b_i < b_j$, where agent i has a strictly smaller utility for each resource than j, then agent j cannot get any resources in an EF-init allocation.[5] We call such a pair of agents $i, j \in A$ a *violating pair*.

Importantly, this observation provides the following useful way to structure the levels. Let $h^* \in [t]$ be the minimum level such that there is an agent $j \in L_{h^*}$ that is part of a violating pair (j is the agent with higher initial utility). If there is no violating pair in the whole instance, set $h^* := t + 1$. It follows that all agents in levels h^* and above cannot get any resources. Moreover, using the structure implied by the fact that there is no violating pair in the first $t^* := h^* - 1$ levels, we can show that within these levels, envy-freeness between pairs of agents behaves "transitively" with respect to the ordering of the levels.

Lemma 3.7. *Let $h, h', h'' \in [t^*]$ with $h' < h < h''$ and consider agents $i^* \in L_h$, $i \in L_{h'}$, and $j \in L_{h''}$. If in some allocation $\mathcal{X}$ there is no envy between agents i and i^*, and between i^* and j under EF-init, then there is no envy between i and j under EF-init.*

[5] Otherwise, agent i has to receive some number x of resources more than agent j to equalize the difference $b_j - b_i$ in initial utilities. However, as j has a strictly higher utility for each of these resources, this implies that j envies i under EF-init, as the utility that j assigns to the x additional resources is strictly greater than $b_j - b_i$.

This enables a dynamic programming approach to decide EF-INIT EXISTENCE for identical resources.

Proof Sketch of Theorem 3.6. We decide EF-INIT EXISTENCE using a dynamic program based on the following table D of size $\mathcal{O}(n \cdot m^2)$. For $a \in [m]$, $b \in [t^*]$ and $c \in [m]$, we store in entry $D[a][b][c]$ whether there is an allocation $\mathcal{X}$ without envy under EF-init between agents in the first b levels satisfying the following three properties: (i) exactly a resources are allocated, (ii) the set of agents that get at least one resource is the union of the first b levels $\bigcup_{\ell \in [b]} L_\ell$, (iii) and every agent $i \in L_b$ gets exactly $c = |X_i|$ resources. Initializing this table for $b = 1$ is trivial. Next, we sketch how to fill entry $D[a][b][c]$ for $b \geq 2$. By Lemma 3.7, it suffices to consider the entries for level $b-1$, and to check if there is an entry for this level (with $a' = a - c \cdot |L_b|$) where there is no envy under EF-init between the agents in L_{b-1} and the agents in L_b, when all agents in L_b get c resources each. After filling table D, the final step is to check if there is a "yes"-entry for $a = m$ where none of the agents that receive no resources are envious. □

A natural follow-up question to this result is whether it can be extended to decide EF1-INIT EXISTENCE for identical resources. However, directly adapting our dynamic programming approach seems challenging, since the crucial "transitivity" property from Lemma 3.7 does not hold any more for EF1-init: For three agents $i, i^*, j \in A$, under EF1-init, we are allowed to disregard a resource each when checking whether i envies i^* and whether i^* envies j. As a result, even when there is no envy between i and i^* and between i^* and j under EF1-init, a single resource may not be sufficient to eliminate envy between i and j. Moreover, it does not hold that no agent in level h^* (the first level containing an agent j from a violating pair) and above can get a resource, if the differences in the utility agents from a violating pair derive for a resource are only small.

4 A Satisfiable Envy Notion

In Observation 3.3, we showed that a complete EF1-init allocation may not exist, even in simple instances, and we proved that finding such an allocation is computationally intractable. This limits the practical applicability of EF1-init. Motivated by this, we propose a relaxation of EF1-init that is always satisfiable (see Section 4.1) and show that an allocation satisfying this relaxed notion can be computed efficiently (see Section 4.2). Furthermore, we identify a natural special case where this relaxation coincides with EF1-init.

4.1 Minimum EF1 with Initial Utilities

To derive a meaningful fairness notion that is always satisfiable, we revisit Observation 3.3. The key challenge illustrated by this counterexample is that agent j with lower initial utility also derives a lower utility from each resource than agent i, so the agents fundamentally disagree on how many resources are needed to bridge their initial utility gap. Thus, there will always be an agent who is envious under EF1-init. To circumvent this, we need to deviate from the comparison approach taken by EF1-init where we compare initial utilities plus bundle values. Instead, when checking whether an agent i with higher initial utility envies another agent j with lower initial utility, we first allow agent j with lower utility to receive a subset of resources, $X^* \subseteq X_j$, that, from their perspective, are worth less than the initial utility gap $b_i - b_j$, and disregard X^* when evaluating if agent i is envious of j's bundle. Formally, one could adjust the EF1-init definition for the case $b_i > b_j$ as follows:

$$\text{If } b_i > b_j, \text{ there is } X^* \subseteq X_j \text{ with } u_j(X^*) < b_i - b_j \text{ and } r \in X_j \text{ so that } u_i(X_i) \geq u_i(X_j \setminus (X^* \cup \{r\})). \tag{1}$$

However, this relaxation falls short if there are multiple agents with low initial utility values that have very different utility functions:

Example 4.1. *Consider an instance with agents $A = \{1, 2, 3\}$ and $m = 10$ identical resources. The agents have initial utilities $b_1 = b_2 = 0$ and $b_3 = 10$. Agents 1 and 3 have $u_1(\{r\}) = u_3(\{r\}) = 100$ for every $r \in R$, while agent 2 has $u_2(\{r\}) = 1$. Since agents 1 and 2 have the same initial utility, we need to have $|X_2| - |X_1| \leq 1$. Moreover, Condition (1) requires that $|X_1| - |X_3| \leq 1$, since only the set $X^* = \emptyset$ satisfies $u_1(X^*) < 10 - 0$. However, this implies that 2 envies 3.*

Intuitively, the example shows that it is insufficient to only consider each pair of agents individually when restricting the set X^*, since this may lead to the introduction of envy between agents with similar initial utility, if the difference in initial utility can be offset more easily for some agents than others. To overcome this, when checking whether an agent i with higher initial utility envies an agent j with lower initial utility, we further relax the restriction on X^* by considering, for each resource in X^*, the minimum utility it gives to any agent with a lower initial utility than i. In Example 4.1, by considering the minimum utility, both of the agents 1 and 2 with low initial utility can choose a set X^* of the same size, allowing them to get the same number of resources more than agent 3 and enabling the existence of a fair allocation. This yields the following notion:

Definition 4.2. *An allocation $\mathcal{X}$ is* minimum EF1-init *(min-EF1-init) if for every pair $i, j \in A$, $X_j = \emptyset$ or it holds that:*

(C1) If $b_i \leq b_j$, then $b_i + u_i(X_i) \geq b_j + u_i(X_j \setminus \{r\})$ for some $r \in X_j$.

(C2) If $b_i > b_j$, then there exists a resource $r \in X_j$ and a subset $X^* \subseteq X_j$ with

$$\sum_{r' \in X^*} \min_{j' \in A \colon b_{j'} < b_i} u_{j'}(\{r'\}) < b_i - b_j,$$

such that $u_i(X_i) \geq u_i(X_j \setminus (X^* \cup \{r\}))$ holds.

While we believe that min-EF1-init is, in a sense, the strongest always-satisfiable fairness notion in this setting, it is also clearly a weaker requirement than (1), and may permit allocations that agents with high initial utilities could perceive as unfair:

Example 4.3. *Consider three agents $A = \{1, 2, 3\}$ with $b_1 = b_2 = 0$, $b_3 = 10$, and $m = 100$ resources. Agent 1 has $u_1(\{r^*\}) = 500$ for a particular $r^* \in R$ and $u_1(\{r\}) = 0$ for all other $r \in R$. Agents 2 and 3 have $u_2(\{r\}) = u_3(\{r\}) = 50$ for all $r \in R$. The allocation $X_1 = \{r^*\}$, $X_2 = R \setminus \{r^*\}$, $X_3 = \emptyset$ is min-EF1-init. Agents 1 and 2 do not envy agent 3 since $X_3 = \emptyset$ and agent 3 does not envy the others (since $|X_1| = 1$ and $u_1(X_2) = 0$). Further, it is easy to see that agent 1 does not envy agent 2 and vice versa. However, agent 3 may regard such an allocation as unfair.*

On the positive side, we show that min-EF1-init implies EF1-init if agents' utility for a resource diminishes in their initial utility.

Definition 4.4. *We call utilities* diminishing, *if $b_i < b_j$ implies $u_i(\{r\}) \geq u_j(\{r\})$ for all $i, j \in A$ and $r \in R$.*

Diminishing utilities arise, for instance, in support programs (like the ARMMAN program featured in the introduction), where aid typically has a stronger effect on those who are initially worse off.

Proposition 4.5. *Let $\mathcal{X}$ be an allocation that is min-EF1-init. If utilities are diminishing, then allocation $\mathcal{X}$ is EF1-init.*

Proof. In the following, to highlight that an inequality follows from Definition 4.4, we will write $\geq_{\text{dim}}$ and $\leq_{\text{dim}}$, and analogously $\leq_{\text{(C2)}}$ to highlight that an inequality follows from (C2) in Definition 4.2. We need to verify that for every pair of agents $i, j \in A$, either $X_j = \emptyset$ or there exists $r \in X_j$ such that

$$b_i + u_i(X_i) \geq b_j + u_i(X_j \setminus \{r\}). \quad (*)$$

Algorithm 1 Round-Robin with Initial Utilities

Partition set of agents A into levels $L_1, \dots, L_t$ and initialize $\mathcal{X}$ as the empty allocation.
Initialize set of *active* agents $\mathcal{L} \subseteq A$ to L_1.
Initialize linear order $\preceq$ over $\mathcal{L}$ to some arbitrary linear order over L_1.
while there are unallocated resources **do**
    Start a new *round*:
    while an agent in $\mathcal{L}$ has not picked in the current round **do**
        Pick first agent $i \in \mathcal{L}$ according to $\preceq$ that has not picked in this round.
        Agent i picks unallocated resource $r \in R$ with maximum utility $u_i(\{r\})$: $X_i \leftarrow X_i \cup \{r\}$
        if all resources are allocated **then return** $\mathcal{X}$
        else if all $i' \in \mathcal{L}$ reached the agents in some level L_ℓ with $\ell \in [t]$ and $L_\ell \cap \mathcal{L} = \emptyset$ **then**
            Let L_ℓ be the level with minimum $\ell \in [t]$ that satisfies the stated condition.
            $\mathcal{L} \leftarrow \mathcal{L} \cup L_\ell$
            Extend $\preceq$ (maintaining the order of agents in $\mathcal{L} \setminus L_\ell$) to the updated $\mathcal{L}$ such that:
                ▷ all agents that have already picked in this round are before the agents in L_ℓ,
                ▷ all agents in $\mathcal{L} \setminus L_\ell$ that have not yet picked are after the agents in L_ℓ.

Since $\mathcal{X}$ is min-EF1-init, conditions (C1) and (C2) from Definition 4.2 hold. If $b_i \leq b_j$, (C1) implies ($*$). If $b_i > b_j$, we have that (C2) holds for some set $X^* \subseteq X_j$ and resource $r \in X_j$. Note that $u_i(X^*) = \sum_{r' \in X^*} u_i(\{r'\}) \leq_{\text{dim}} \sum_{r' \in X^*} \min_{j' \in A:\, b_{j'} < b_i} u_{j'}(\{r'\}) <_{\text{(C2)}} b_i - b_j$. Since $u_i(X_i) \geq u_i(X_j \setminus (X^* \cup \{r\}))$ by (C2), it follows that $u_i(X_j \setminus \{r\}) \leq u_i(X_j \setminus (X^* \cup \{r\})) + u_i(X^*) \leq_{\text{(C2), dim}} u_i(X_i) + b_i - b_j$, which implies ($*$). $\square$

4.2 A Round-Robin Algorithm Guaranteeing Minimum EF1 with Initial Utilities

In this section, we propose an algorithm that computes a complete minimum-EF1-init allocation. For this, we extend the well-known "round-robin" algorithm for the setting with initial utilities.[6]

Idea In our algorithm, we use the concept that the agents are partitioned into t levels by their initial utility, as introduced in Definition 2.1. Our extension is based around the idea that at any point of our round-robin algorithm, only a subset of the agents is *active* and can be assigned a pick. We maintain a picking order over all currently active agents in the algorithm. In the beginning, only the agents with the lowest initial utility in level L_1 are active. After each pick, we check whether all currently active agents have *reached* the agents in some (not yet active) level L_ℓ with initial utility value b. By this, we mean that for each currently active agent i, the sum of i's initial utility and i's utility for their bundle is at least the initial utility b (all active agents believe that they have "reached the initial situation" of all agents from L_ℓ).

Definition 4.6. *We say that an agent $i \in A$ with bundle $X_i \subseteq R$ has* reached *an agent $j \in A$ if $b_i + u_i(X_i) \geq b_j$.*

If all active agents have reached the agents in L_ℓ, the agents in L_ℓ become active and are inserted into the picking order after the agents that have already picked in this round, meaning that they get to pick directly after being activated. This ensures that each agent from L_ℓ prefers the resources picked by them over the resources picked by any other agent that was active before them after L_ℓ's activation. Another observation crucial to the inner workings of the algorithm is that if agent j is activated after another agent i picked some resource r, then i had not yet reached j in the round before, implying

[6]In the standard round-robin algorithm, we fix an ordering of the agents. In each round of the algorithm, following this ordering, agents pick their favorite so far unallocated resource. Once there are no more resources left, the algorithm terminates. In the classical setting without initial utilities, the resulting allocation is guaranteed to satisfy EF1.

that the value of i for their current bundle without r is less than the initial utility difference $b_j - b_i$ (see Observation 4.9). A full description is given in Algorithm 1. We show that Algorithm 1 returns a complete allocation satisfying min-EF1-init, implying that a min-EF1-init allocation always exists.

Theorem 4.7. *The allocation computed by Algorithm 1 is complete and satisfies min-EF1-init. Further, Algorithm 1 runs in polynomial time.*

4.3 Proof of Theorem 4.7

In this subsection, we give an overview of the proof of Theorem 4.7. We start with defining notation and formalizing two observations (Section 4.3.1) before we present the crucial activation gap lemma (Section 4.3.2) and finally give a proof sketch (Section 4.3.3).

4.3.1 Notation and Inital Observations

We say that an agent $i \in A$ is *active* in a round if $i \in \mathcal{L}$ at the end of the round. An agent $i \in A$ is *activated* in a round (after/by a pick) if i is added to $\mathcal{L}$ at some point during the round (after/by the pick). A *picking sequence* starting with/after some pick is the contiguous subsequence of all picks until the end of the algorithm after this pick.[7] We start by observing the following properties of Algorithm 1.

Observation 4.8. *Consider Algorithm 1 and let $\preceq$ be the final linear order when the algorithm terminates. Let $i, j \in A$ be two distinct agents. For $k \in \{i, j\}$, let X_k^* be the bundle picked by k in a given subset $\mathcal{R}$ of rounds in which both i and j are active and where j picks in at least one round in $\mathcal{R}$. Then it holds that:*

1. *Assume that i picks in the last round[8] in $\mathcal{R}$. For the resource $r \in X_i^*$ picked by agent i in the last round in $\mathcal{R}$ and some $r' \in X_j^*$, we have that $u_i(X_i^* \setminus \{r\}) \geq u_i(X_j^* \setminus \{r'\})$.*
2. *If i does not pick in the last round in $\mathcal{R}$, we have that $u_i(X_i^*) \geq u_i(X_j^* \setminus \{r'\})$ for some $r' \in X_j^*$.*
3. *If $i \preceq j$, then $u_i(X_i^*) \geq u_i(X_j^*)$.*

Secondly, we formalize an observation from above:

Observation 4.9. *Consider Algorithm 1 and let $i, j \in A$ be two agents with $b_j > b_i$ such that j is activated after agent i picked some resource $r \in R$. Let X_i^* be the bundle of i after this pick. Then, it holds that $u_i(X_i^* \setminus \{r\}) < b_j - b_i$.*

4.3.2 The Activation Gap Lemma

The following lemma gives a bound on the value of the bundle assigned to an active agent j at the point when some agent i is activated. More concretely, we show that the sum of the minimum utility of an agent with lower initial utility than i for each resource in j's bundle is at most the initial utility gap between i and j after removing some resource r. This is crucial in our proof of Theorem 4.7, since it implies that j's current bundle without r satisfies the restriction on X^* in min-EF1-init (C2), allowing to disregard the resources picked by j in rounds when i was not active when checking whether i envies j.

Lemma 4.10 (Activation Gap Lemma)**.** *Consider Algorithm 1 and let $i, j \in A$ be two agents that were active in some round with $b_i \geq b_j$ and $b_i > \min_{i' \in A} b_{i'}$. For $k \in A$, let X_k^* be the bundle of k after the pick that activated i. Then, either $X_j^* = \emptyset$ or there exists $r \in X_j^*$ such that*

$$\sum_{r' \in X_j^* \setminus \{r\}} \min_{j' \in A\colon b_{j'} < b_i} u_{j'}(\{r'\}) < b_i - b_j.$$

[7] We will also consider picking sequences starting with a given pick up to another pick, by which we mean the contiguous subsequence of all picks in the algorithm between the two given picks.

[8] The ordering of the rounds in $\mathcal{R}$ that we refer to here is given by the order in which they occur in Algorithm 1.

Proof. For agents $a, b \in A$ with $a \in L_h$ and $b \in L_{h'}$ for $1 \leq h' \leq h \leq t$, we define $\Delta_{a,b} := h - h'$. We prove the claim by induction over $\Delta_{i,j}$. For $\Delta_{i,j} = 0$, the claim holds trivially, as both agents are activated in the same round and thus $|X_j^*| \leq 1$.

Induction Step. Assume the claim holds for all $i, j \in A$ with $b_i \geq b_j$ and $b_i > \min_{i' \in A} b_{i'}$, and with $\Delta_{i,j} \leq \Delta^*$ for some $\Delta^* \in [t-1]$. We show that the claim then also holds for $i, j \in A$ with $b_i \geq b_j$ and $b_i > \min_{i' \in A} b_{i'}$, and with $\Delta_{i,j} = \Delta^* + 1$. If $X_j^* = \emptyset$, then the claim holds trivially, so we assume in the following that $X_j^* \neq \emptyset$. Note that we have that $b_i > b_j$ since $\Delta_{i,j} \geq 1$. Let $j^* \in A$ be the agent after whose pick agent i was activated and let $r \in X_{j^*}^*$ be the resource picked by j^* in this pick. Let $\preceq$ be the final linear order when the algorithm terminates. If $j = j^*$, the claim holds, since $b_i - b_j > u_j(X_j^* \setminus \{r\})$ for the resource $r \in X_j^*$ picked by j in the round when i was activated by Observation 4.9 and $u_j(X_j^* \setminus \{r\}) \geq \sum_{r' \in X_j^* \setminus \{r\}} \min_{j' \in A:\, b_{j'} < b_i} u_{j'}(\{r'\})$ since $b_j < b_i$. Thus, assume that $j^* \neq j$. Note that we have that $j^* \preceq i$. We make a case distinction based on b_j and b_{j^*}.

Case 1 ($b_j = b_{j^*}$). First, consider that $b_j = b_{j^*}$. Note that this implies that j and j^* are active in the same rounds. Let $\mathcal{R}$ be the set of rounds in which j and j^* are active up to (and including) the round in which i is activated. In these rounds, agent j^* picks the resources in $X_{j^*}^*$ and agent j picks a superset of the resources X_j^* (since agent j may pick after i is added in the last round in $\mathcal{R}$). Observation 4.8 (1) implies that $u_{j^*}(X_{j^*}^* \setminus \{r\}) \geq u_{j^*}(X_j^* \setminus \{r'\})$ for $r \in X_{j^*}^*$ picked last by j^* and some $r' \in X_j^*$. Moreover, it holds that $u_{j^*}(X_{j^*}^* \setminus \{r\}) < b_i - b_j$ by Observation 4.9. Thus, the claim holds in this case since $b_i - b_j > u_{j^*}(X_j^* \setminus \{r'\}) \geq \sum_{r'' \in X_j^* \setminus \{r'\}} \min_{j' \in A:\, b_{j'} < b_i} u_{j'}(\{r''\})$.

Case 2 ($b_j > b_{j^*}$). Next, consider that $b_j > b_{j^*}$ and let $\tilde{X}_{j^*}^*$ be the bundle of j^* after the pick that activated j. Since j was activated, it needs to hold that

$$u_{j^*}(\tilde{X}_{j^*}^*) \geq b_j - b_{j^*}. \tag{2}$$

Recall that agent j^* picked the resource $r \in X_{j^*}^*$ in the round in which i is activated. Thus, the bundle of agent j^* in the round prior is $X_{j^*}^* \setminus \{r\}$. By Observation 4.9, we get that $u_{j^*}(X_{j^*}^* \setminus \{r\}) < b_i - b_{j^*}$. Note that $r \notin \tilde{X}_{j^*}^*$, as $\tilde{X}_{j^*}^*$ is the bundle of j^* when j was activated and r is picked by j^* when i is activated, and it holds that $b_i > b_j$. Using Equation (2), this implies that

$$u_{j^*}(X_{j^*}^* \setminus (\tilde{X}_{j^*}^* \cup \{r\})) < b_i - b_j, \tag{3}$$

Note that it suffices to show that

$$b_i - b_j > u_{j^*}(X_j^* \setminus \{r'\}) \tag{4}$$

for some $r' \in X_j^*$, since $u_{j^*}(X_j^* \setminus \{r'\}) \geq \sum_{r'' \in X_j^* \setminus \{r'\}} \min_{j' \in A:\, b_{j'} < b_i} u_{j'}(\{r''\})$, as $b_{j*} < b_i$. To show that there exists a resource such that Equation (4) holds, we now make a case distinction between $j \preceq j^*$ and $j^* \preceq j$. Note that $j^* \preceq j$ implies that $j^* \preceq i \preceq j$, since $b_j < b_i$ and there can only be agents with initial utility at least b_i between j^* and i in $\preceq$: Since i was activated directly after j^*'s pick, only agents that became active after or by the same pick as i can be between j^* and i.

If $j^* \preceq i \preceq j$, consider the set $\mathcal{R}$ of all rounds after the round in which j is activated up to (excluding) the round in which i is activated. The resources picked by agent j^* in the rounds in $\mathcal{R}$ is the set $X_{j^*}^* \setminus (\tilde{X}_{j^*}^* \cup \{r\})$. Note that agent j has not yet picked in the round when i is activated, since $j^* \preceq i \preceq j$. Thus, the resources picked by j in the rounds in $\mathcal{R}$ is exactly the set $X_j^* \setminus \{r'\}$ for the resource $r' \in X_j^*$ picked first by j. Since $j^* \preceq j$, we can apply Observation 4.8 (3) to show that Equation (4) holds for resource r', as[9]

$$b_i - b_j >_{(3)} u_{j^*}(X_{j^*}^* \setminus (\tilde{X}_{j^*}^* \cup \{r\})) \geq_{(4.8)} u_{j^*}(X_j^* \setminus \{r'\}).$$

[9]In the following, to highlight that an inequality follows from a previous result, definition, or equation, we add a reference in the subscript, e.g., by writing $\leq_{(*)}$ for an inequality that follows from a previously introduced equation $(*)$.

If $j \preceq j^*$, we consider the set $\mathcal{R}$ of all rounds in which j is active up to (now including) the round in which i is activated. Note that as $j \preceq j^*$, agent j^* has not yet picked in the round when j is activated, so $\tilde{X}^*_{j^*}$ only contains resources picked by j^* in rounds in which j is not active. Thus, the resources picked by agent j^* in the rounds in $\mathcal{R}$ is the set $X^*_{j^*} \setminus \tilde{X}^*_{j^*}$ and agent j picks the resources in X^*_j. Using Observation 4.8 (1), it follows there exists a resource $r' \in X^*_j$ such that Equation (4) holds, since $b_i - b_j >_{(3)} u_{j^*}(X^*_{j^*} \setminus (\tilde{X}^*_{j^*} \cup \{r\})) \geq_{(4.8)} u_{j^*}(X^*_j \setminus \{r'\})$, for resource $r' \in X^*_j$ from Observation 4.8 (1).

Case 3 ($b_j < b_{j^*}$). It remains to consider that $b_j < b_{j^*}$. Note that this implies that b_{j^*} is strictly greater than the minimum initial utility of all agents. Since $b_{j^*} < b_i$, we have that $\Delta_{j^*,j} \leq \Delta^*$, so we can apply the induction hypothesis for j^* and j. Let $\tilde{X}^*_j$ be the bundle of j after the pick that activated j^*. Note that it cannot be the case that $\tilde{X}^*_j = \emptyset$, since $b_j < b_{j^*}$. Thus, there exists $r' \in \tilde{X}^*_j$ such that

$$b_{j^*} - b_j > \sum_{r'' \in \tilde{X}^*_j \setminus \{r'\}} \min_{j' \in A:\, b_{j'} < b_{j^*}} u_{j'}(\{r''\}). \tag{5}$$

Moreover, by Observation 4.9, we have that $u_{j^*}(X^*_{j^*} \setminus \{r\}) < b_i - b_{j^*}$ for the resource $r \in X^*_{j^*}$ picked last by j^*. Consider the picking sequence starting with agent j^*'s first pick up to (excluding) the pick by j^* that activates i. Agent j picks the resources in $X^*_j \setminus \tilde{X}^*_j$, and agent j^* the resources in $X^*_{j^*} \setminus \{r\}$. Since j^* picks first in the picking sequence, it follows that $u_{j^*}(X^*_{j^*} \setminus \{r\}) \geq u_{j^*}(X^*_j \setminus \tilde{X}^*_j)$, so we have

$$b_i - b_{j^*} >_{(4.9)} u_{j^*}(X^*_{j^*} \setminus \{r\}) \geq u_{j^*}(X^*_j \setminus \tilde{X}^*_j). \tag{6}$$

This implies the claim, since

$$\begin{aligned} b_i - b_j = (b_i - b_{j^*}) + (b_{j^*} - b_j) &>_{(6),(5)} u_{j^*}(X^*_j \setminus \tilde{X}^*_j) + \sum_{r'' \in \tilde{X}^*_j \setminus \{r'\}} \min_{j' \in A:\, b_{j'} < b_{j^*}} u_{j'}(\{r''\}) \\ &\geq \sum_{r'' \in X^*_j \setminus \{r'\}} \min_{j' \in A:\, b_{j'} < b_i} u_{j'}(\{r''\}). \end{aligned}$$

□

4.3.3 Putting the Pieces Together

We now sketch the proof that Algorithm 1 computes a complete min-EF1-init allocation. Intuitively, the main idea is that an agent i does not envy another agent j with higher initial utility, since i is allowed to pick resources until i has reached j before j gets activated and makes their first pick. Moreover, i does not envy an agent with lower initial utility, since the bound from Lemma 4.10 allows to disregard the resource picked by this agent before i was activated in min-EF1-init, and agent i picks before the other agent in the picking sequence after this pick.

Theorem 4.7. *The allocation computed by Algorithm 1 is complete and satisfies min-EF1-init. Further, Algorithm 1 runs in polynomial time.*

Proof (Sketch). Let $\mathcal{X}$ be the computed allocation and let $\preceq$ be the final linear order. We show that for any two agents $i, j \in A$, agent i does not envy agent j under min-EF1-init in $\mathcal{X}$. Thus, we need to show that $X_j = \emptyset$ or condition (C1) holds if $b_i \leq b_j$ or (C2) holds if $b_i > b_j$. In the following, we sketch the proof for the case $X_j \neq \emptyset$ and $b_i < b_j$ and the case $X_j \neq \emptyset$ and $b_i > b_j$.

(Case $X_j \neq \emptyset$ and $b_i \leq b_j$) For $k \in A$, let X_k^* be the bundle of k after j's first pick and let $r \in R$ be the first resource picked by j, i.e., $X_j^* = \{r\}$. Since i had reached j when agent j was activated:

$$b_i + u_i(X_i^*) \geq b_j + \underbrace{u_i(X_j^* \setminus \{r\})}_{=0}.$$

Moreover, since i picks before j in the picking sequence starting with the pick after j's first pick, agent i prefers the resources they pick in this picking sequence (i.e., $X_i \setminus X_i^*$) over those picked by j (i.e., $X_j \setminus X_j^*$). Thus, $u_i(X_i \setminus X_i^*) \geq u_i(X_j \setminus X_j^*)$ and (C1) follows for resource r.

(Case $X_j \neq \emptyset$ and $b_i > b_j$) We need to argue that condition (C2) holds. We make a case distinction: First, we consider the case that agent i was activated at some point during the algorithm. Let X_k^* be the bundle of agent $k \in A$ after the pick that activated i. As $b_i > b_j$ and j has reached i before their activation, we have that $X_j^* \neq \emptyset$. Thus, Lemma 4.10 implies that there is a resource $r \in X_j^*$ such that

$$\sum_{r' \in X_j^* \setminus \{r\}} \min_{j' \in A\colon b_{j'} < b_i} u_{j'}(\{r'\}) < b_i - b_j.$$

Importantly, this implies that the set $X^* = X_j^* \setminus \{r\}$ satisfies the condition from (C2). Consider the picking sequence of all picks starting with (including) agent i's first pick. In this sequence, agent i picks the resources in X_i and agent j picks the resources in $X_j \setminus X_j^*$. As agent i picks first, we have that $u_i(X_i) \geq u_i(X_j \setminus X_j^*) = u_i(X_j \setminus (X^* \cup \{r\}))$, so (C2) holds for resource r and $X_j^* \setminus \{r\}$.

If agent i was not activated, then there needs to be an agent $j^* \in A$ with $b_{j^*} < b_i$ that never reaches i, i.e., $u_{j^*}(X_{j^*}) < b_i - b_{j^*}$. In this case, using a more involved case distinction, one can show that it is always possible to choose a resource $r \in X_j$ and a set $X^* \subseteq X_j$ satisfying the constraint on X^* from (C2) such that $X^* \cup \{r\} = X_j$, which directly implies $u_i(X_i) \geq 0 = u_i(X_j \setminus (X^* \cup \{r\}))$. □

5 Conclusion

We initiated the study of fair allocation with initial utilities, focusing on fairness notions that implement the principle of equality of outcome. Our results show that incorporating initial utilities fundamentally alters and complicates the nature of fair allocation problems. Nevertheless, positive algorithmic and axiomatic results can be recovered by tailoring fairness notions to this setting, most notably through our always-satisfiable notion of min-EF1-init, and by focusing on special cases, such as our polynomial-time algorithm for EF-init Existence for identical resources.

Our work opens several directions for future research. First, it would be interesting to obtain a more fine-grained understanding of the complexity of EF1-init Existence, for instance, in the case when resources are identical (see discussion at the end of Section 3) or agents value resources as 0 or 1. Second, it is worth exploring further relaxations of envy-freeness for initial utilities. Conceptually, min-EF1-init relaxes EF1 in a way that allows more flexibility to assign resources to agents with lower initial utility without creating envy, potentially at the expense of agents with higher initial utility (see Example 4.3). It would be interesting to study whether always-satisfiable envy notions can also incorporate favorable treatment of agents with higher initial utility or interpolate between both directions. Third, adapting fairness notions beyond envy-freeness, such as the maximin share (MMS), to equality of outcome deserves attention.[10] Lastly, the idea of initial utilities and equality of outcome also deserve attention in other allocation domains, including the allocation of indivisible resources with non-additive utilities [25], chore division [4], online settings [31, 9], and the allocation of divisible resources [12].

[10] Note that in Appendix B we use the work of Prakash HV et al. [38] to show that we cannot guarantee the existence of an α-approximation of an adapted version of MMS for any $\alpha > 0$ in the presence of initial utilities

References

[1] Martin Aleksandrov and Toby Walsh. Online fair division: A survey. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI '20)*, pages 13557–13562. AAAI Press, 2020.

[2] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.*, 322:103965, 2023.

[3] ARMMAN. Helping mothers and children. https://armman.org, 2025. Accessed: 2025-05-01.

[4] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI '19)*, pages 53–59. ijcai.org, 2019.

[5] Haris Aziz, Péter Biró, Tamás Fleiner, Serge Gaspers, Ronald de Haan, Nicholas Mattei, and Baharak Rastegari. Stable matching with uncertain pairwise preferences. *Theor. Comput. Sci.*, 909:1–11, 2022.

[6] Siddharth Barman, Arindam Khan, Sudarshan Shyam, and K. V. N. Sreenivas. Finding fair allocations under budget constraints. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI 23)*, pages 5481–5489. AAAI Press, 2023.

[7] Catherine Barnard and Bob Hepple. Substantive equality. *The Cambridge Law Journal*, 59(3): 562–585, 2000.

[8] Hamsa Bastani, Kimon Drakopoulos, Vishal Gupta, Ioannis Vlachogiannis, Christos Hadjichristodoulou, Pagona Lagiou, Gkikas Magiorkinis, Dimitrios Paraskevis, and Sotirios Tsiodras. Efficient and targeted covid-19 border testing via reinforcement learning. *Nature*, 599(7883): 108–113, 2021.

[9] Gerdus Benadè, Daniel Halpern, and Alexandros Psomas. Dynamic fair division with partial information. In *Proceedings of the Thirty-Sixth Annual Conference on Neural Information Processing Systems (NeurIPS '22)*, 2022.

[10] Sylvain Bouveret and Jérôme Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *J. Artif. Intell. Res.*, 32:525–564, 2008.

[11] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, pages 284–310. Cambridge University Press, 2016.

[12] Steven J Brams and Alan D Taylor. *Fair Division: From cake-cutting to dispute resolution.* Cambridge University Press, 1996.

[13] 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 2019 ACM Conference on Economics and Computation (EC '19)*, pages 527–545. ACM, 2019.

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

[15] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. Weighted envy-freeness in indivisible item allocation. *ACM Trans. Economics and Comput.*, 9(3):18:1–18:39, 2021.

[16] Ilan Reuven Cohen, Alon Eden, Talya Eden, and Arsen Vasilyan. Plant-and-steal: Truthful fair allocations via predictions. In *Proceedings of the Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS '24)*, 2024.

[17] Commission on Social Determinants of Health. *Closing the Gap in a Generation: Health Equity Through Action on the Social Determinants of Health.* World Health Organization, Geneva, 2008.

[18] Argyrios Deligkas, Eduard Eiben, Robert Ganian, Tiger-Lily Goldsmith, and Stavros D. Ioannidis. The complexity of extending fair allocations of indivisible goods. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI 25)*, pages 13745–13753. AAAI Press, 2025.

[19] Vânia M. Félix Dias, Guilherme Dias da Fonseca, Celina M. H. de Figueiredo, and Jayme Luiz

Szwarcfiter. The stable marriage problem with restricted pairs. *Theor. Comput. Sci.*, 306(1-3): 391–405, 2003.
[20] Anmei Dong, Morris Siu-Yung Jong, and Ronnel B King. How does prior knowledge influence learning engagement? the mediating roles of cognitive load and help-seeking. *Frontiers in psychology*, 11:591203, 2020.
[21] Gideon Elford. Equality of Opportunity. In Edward N. Zalta and Uri Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2023 edition, 2023.
[22] Ezekiel J Emanuel and Govind Persad. The shared ethical framework to allocate scarce medical resources: a lesson from covid-19. *The Lancet*, 401(10391):1892–1902, 2023.
[23] Alireza Farhadi, Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Sébastien Lahaie, David M. Pennock, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods to asymmetric agents. *J. Artif. Intell. Res.*, 64:1–20, 2019.
[24] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
[25] Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods: Beyond additive valuations. *Artif. Intell.*, 303:103633, 2022.
[26] Jonathan R. Goldman and Ariel D. Procaccia. Spliddit: unleashing fair division algorithms. *SIGecom Exch.*, 13(2):41–46, 2014.
[27] Daniel Halpern and Nisarg Shah. Fair division with subsidy. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT '19)*, pages 374–389. Springer, 2019.
[28] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In *Proceedings of the Thirtieth Annual Conference on Neural Information Processing Systems (NIPS '16)*, pages 3315–3323, 2016.
[29] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, Jun Wang, and Lirong Xia. Fair division through information withholding. *CoRR*, abs/1907.02583, 2019. URL http://arxiv.org/abs/1907.02583.
[30] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, Hejun Wang, and Lirong Xia. Fair division through information withholding. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI 20)*, pages 2014–2021. AAAI Press, 2020.
[31] Ian A. Kash, Ariel D. Procaccia, and Nisarg Shah. No agent left behind: Dynamic fair division of multiple resources. *J. Artif. Intell. Res.*, 51:579–603, 2014.
[32] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. *J. ACM*, 65(2):8:1–8:27, 2018.
[33] Haylee Lane, Mitchell Sarkies, Jennifer Martin, and Terry Haines. Equity in healthcare resource allocation decision making: A systematic review. *Social Science & Medicine*, 175:11–27, 2017.
[34] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC '04)*, pages 125–131. ACM, 2004.
[35] Shengxin Liu, Xinhang Lu, Mashbat Suzuki, and Toby Walsh. Mixed fair division: A survey. *J. Artif. Intell. Res.*, 80:1373–1406, 2024.
[36] Hervé Moulin. *Fair division and collective welfare*. MIT press, 2004.
[37] Hervé Moulin. Fair division in the internet age. *Annual Review of Economics*, 11(1):407–441, 2019.
[38] Vishwa Prakash HV, Ayumi Igarashi, and Rohit Vaish. Fair and efficient completion of indivisible goods. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI 25)*, pages 14045–14053. AAAI Press, 2025.
[39] Scottish Government. Fairer scotland duty: Guidance for public bodies, 2021.
[40] Donald J. Trump. Restoring equality of opportunity and meritocracy, 2025. Executive Order 14281.
[41] Shresth Verma, Gargi Singh, Aditya Mate, Paritosh Verma, Sruthi Gorantla, Neha Madhiwalla, Aparna Hegde, Divy Thakkar, Manish Jain, Milind Tambe, and Aparna Taneja. Expanding impact

of mobile health programs: SAHELI for maternal and child care. *AI Mag.*, 44(4):363–376, 2023.
[42] Shresth Verma, Yunfan Zhao, Sanket Shah, Niclas Boehmer, Aparna Taneja, and Milind Tambe. Group fairness in predict-then-optimize settings for restless bandits. In *Proceedings of the 40th Conference on Uncertainty in Artificial Intelligence (UAI '24)*, volume 244 of *Proceedings of Machine Learning Research*, pages 3448–3469. PMLR, 2024.
[43] Xiaowei Wu, Bo Li, and Jiarui Gan. Budget-feasible maximum nash social welfare allocation is almost envy-free. *CoRR*, abs/2012.03766, 2020. URL https://arxiv.org/abs/2012.03766.
[44] Lirong Xia and Vincent Conitzer. Determining possible and necessary winners given partial orders. *J. Artif. Intell. Res.*, 41:25–67, 2011.
[45] Shengwei Zhou, Rufan Bai, and Xiaowei Wu. Multi-agent online scheduling: MMS allocations for indivisible items. In *Proceedings of the Fortieth International Conference on Machine Learning (ICML '23)*, volume 202 of *Proceedings of Machine Learning Research*, pages 42506–42516. PMLR, 2023.

A Appendix

In the following, to highlight that an inequality follows from a previous result, definition, or equation, we add a reference in the subscript, e.g., by writing $\leq_{(*)}$ for an inequality that follows from a previously introduced equation $(*)$. Given an instance I, we denote the set of all allocations by $\Pi(I) \coloneqq \big\{(X_1, \ldots, X_n) \in (2^R)^n \,\big|\, \forall i, j \in A,\, i \neq j \implies X_i \cap X_j = \emptyset\big\}$.

B Additional Related Work

As mentioned in Section 1.2, the setting with initial utilities is formally closely related to the completion (or extension) setting considered by Prakash HV et al. [38], Deligkas et al. [18]. In their setting, in addition to the resources R and agents A with utility functions, an instance of the completion problem they consider contains a *frozen* (or partial) allocation $\bar{\mathcal{X}}$ of a subset $F \subseteq R$ of the resources (to which they refer as frozen resources). All other resources are called *open*. A straightforward way to translate an instance in our setting to an instance in their setting is the following construction: Given an instance with agents A with utility function u_i and initial utility b_i for $i \in A$ and resources R, we construct the following instance in the completion setting, which we call the *derived* instance. The derived instance has the same set of agents A. The set of resources is $R' = R \cup \{r_i^* \mid i \in A\}$. Agent $i \in A$ has utility function $u_i' : 2^{R'} \to \mathbb{R}_{\geq 0}$ with $u_i'(\{r\}) = u_i(\{r\})$ for all $r \in R$ and $u_i'(\{r_j^*\}) = b_j$ for all $j \in A$. In the frozen allocation $\bar{\mathcal{X}}$, each agent $i \in A$ is assigned the resource r_i^* and the resources in R are unassigned.

First, we discuss the work by Deligkas et al. [18], who focus on the problem of deciding whether the frozen allocation can be completed to a complete allocation without envy under (standard) EF. As this problem is known to be NP-hard even without any frozen resources, they study the parameterized complexity of this problem with the number of open resources (i.e., $|R \setminus F|$) as parameter. They find that the problem is W[1]-hard for this parameter. To obtain a positive result, they consider *agent types* (two agents are of the same type if they have the same utility for each resource), showing that the completion problem is fixed-parameter tractable when parameterized by the number of open resources plus the number of agent types. Moreover, they consider variants of the completion problem where the open resources can be allocated to at most p different agents, and study the parameterized complexity of this problem for parameter p. While their positive parameterized complexity results could technically be translated to our setting, the parameter regarding the number of open resources becomes vacant, since there is only one frozen resource per agent in the derived instance and all resources from the original instance are open.

We continue by discussing how our work relates to the work by Prakash HV et al. [38]. They also consider the completion problem for standard fairness notions, but, in addition to EF, consider EF1 and the maximin share (MMS) [16, 32]. Their work too focuses on the computational complexity of determining whether an allocation completing the given frozen allocation exists that satisfies these fairness notions, and they prove that, in contrast to the standard setting, many such problems become computationally intractable. We discuss how their results relate to our setting in the following. First, we remark that while the setting with initial utilities can be seen as a special case of the completion setting as described above, an EF1 allocation in the completion setting does not correspond to an EF1-init allocation in the initial utility setting: Since EF1 does not distinguish between open and frozen resources, it is possible to simply "disregard" the initial disparities by removing the corresponding frozen resources when comparing two agents bundles to check for envy. This is not possible in our adapted EF1-init notion, which requires that the difference in initial utility is equalized and that the removed resource needs to come from the resources to be distributed (or, that the agent with higher initial utility does not get any resources). Therefore, our new envy-based fairness notions for initial utilities are better suited to our specialized setting. Moreover, since our setting is a special case of their completion setting, their negative results do not immediately translate to our setting.

In addition to EF1, Prakash HV et al. [38] consider MMS. As a first step towards studying MMS in the initial utility setting, we show that their result that even an arbitrarily bad approximation of MMS cannot be guaranteed extends to our initial utility setting. An intuitive adaptation of MMS to our setting with initial utilities is the following.

Definition B.1. *The* max-min-share *of an agent $i \in A$ in instance I with initial utilities is*

$$\mu_i := \max_{\mathcal{X} \in \Pi(I)} \min_{j \in A} b_j + u_i(X_j).$$

We say that an allocation $\mathcal{X}$ is max-min-share fair *(MMS-init) if it holds that $b_i + u_i(X_i) \geq \mu_i$ for all $i \in A$. An allocation $\mathcal{X}$ is α-MMS-init for $0 < \alpha \leq 1$, if it holds that $b_i + u_i(X_i) \geq \alpha \cdot \mu_i$ for all $i \in A$.*

This definition matches the idea behind the extension of MMS to the completion setting by Prakash HV et al. [38], who define the max-min-share as

$$\mu_i^* := \max_{\mathcal{X} \in \Pi^*(I)} \min_{j \in A} u_i(X_j),$$

where $\Pi^*(I)$ denotes the set of all possible allocations that complete the frozen allocation $\bar{\mathcal{X}}$, that is, $\bar{X}_i \subseteq X_i$ for all $\mathcal{X} \in \Pi^*(I)$ and $i \in A$. If we reduce an allocation instance with initial utilities to the completion setting as described above, then the max-min-share μ_i in the original instance is equal to the max-min-share μ_i^* in the derived instance for any $i \in A$: It holds that $u_i(\bar{X}_j) = b_j$ and $\bar{X}_i \subseteq X_i$ for all $\mathcal{X} \in \Pi^*(I)$ and $i, j \in A$. Moreover, let $\mathcal{X}$ be an allocation in the original instance and define allocation $\mathcal{X}'$ with $X_i' = X_i \cup \{r_i^*\}$ for all $i \in A$ in the derived instance. Then, $\mathcal{X}$ is MMS-init if and only if $\mathcal{X}'$ is MMS, since $b_i + u_i(X_i) = u_i'(X_i')$ for all $i \in A$. Using this fact, it follows from Proposition 1 by Prakash HV et al. [38] that a Pareto-optimal (PO) and MMS-init allocation always exists for binary additive valuations and initial utilities $b_i \in \{0, 1\}$ for all $i \in A$: they show that in their setting such an allocation always exists in this case when the frozen allocation is PO, which is the case in any derived instance.

Proposition B.2. *For binary additive valuations and initial utilities $b_i \in \{0, 1\}$ for all $i \in A$, an MMS-init and PO allocation always exists and can be computed in polynomial time.*

On the negative side, Prakash HV et al. [38] show that for every $0 < \alpha \leq 1$, there exists an instance in their completion setting that does not admit an α-MMS allocation. Since their construction introduces an asymmetry in the utility of the agents for the frozen resources, we cannot immediately translate their result to our setting, since all agents agree on the initial utilities. However, using an adapted construction, we can prove the result in our setting, using their key idea that each agent $i \in [n]$ believes that the first i agents should share the resources equally, since the remaining agents are already "rich" enough. We note that in both constructions, the number of agents grows exponentially in α^{-1}.

Proposition B.3. *For every $0 < \alpha \leq 1$, there exists an instance that does not admit an α-MMS-init allocation, even when all resources are identical.*

Proof. Let $0 < \alpha \leq 1$. For every $\ell \in \mathbb{N}$, we denote by $H_\ell := \sum_{i \in [\ell]} \frac{1}{i}$ the ℓ-th harmonic number. We choose the number of agents n such that $1 < \frac{\alpha}{3} \cdot H_n$. The set R contains $m = 2n$ identical resources. Agent $i \in [n]$ has an initial utility and a utility for a resource of $b_i = u_i(\{r\}) = m^{(i-1)}$ for all $r \in R$. Observe that for every agent $i \in [n]$, it holds that $\mu_i \geq \lfloor \frac{m}{i} \rfloor \cdot m^{(i-1)}$, which can be achieved by distributing the resources equally among the first i agents. Note that all remaining agents $j \in A \setminus [i]$ already have at least initial utility $b_j \geq m^i \geq \mu_i$.

Now, suppose for the sake of contradiction that there exists an α-MMS-init allocation $\mathcal{X}$. Then, it needs to hold that $b_i + u_i(X_i) \geq \alpha \cdot \mu_i \geq \alpha \cdot \lfloor \frac{m}{i} \rfloor \cdot m^{(i-1)} \geq \frac{\alpha \cdot m}{2i} \cdot m^{(i-1)}$ for all $i \in A$. It follows that for all $i \in A$,

$$b_i + u_i(X_i) \geq \frac{\alpha \cdot m}{2i} \cdot m^{(i-1)} \implies u_i(X_i) \geq \frac{\alpha \cdot m}{2i} \cdot m^{(i-1)} - m^{(i-1)} \implies |X_i| \geq \frac{\alpha \cdot m}{2i} - 1.$$

Since $\mathcal{X}$ can allocate at most all m resources, it follows that

$$\begin{aligned} m &\geq \sum_{i \in A} (\frac{\alpha \cdot m}{2i} - 1) = m(\sum_{i \in A} \frac{\alpha}{2i}) - n \\ \iff 1 &\geq (\sum_{i \in A} \frac{\alpha}{2i}) - \frac{n}{m} = (\sum_{i \in A} \frac{\alpha}{2i}) - \frac{1}{2} \\ \iff \frac{3}{2} &\geq \frac{\alpha}{2} \sum_{i \in A} \frac{1}{i} \iff 1 \geq \frac{\alpha}{3} \cdot H_n. \end{aligned}$$

However, this contradicts our choice of n, which completes the proof. □

C Additional Material for Section 3

We first prove the NP-hardness of EF1-init Existence for additive utility functions.

Theorem 3.4. *EF1-init Existence is NP-complete.*

Our reduction is from Equitable Coloring. To define Equitable Coloring, we first introduce some necessary definitions.

Definition C.1. *For some $\ell \in \mathbb{N}$, an ℓ-*coloring *of a graph $G = (V, E)$ assigns a color $c(v) \in [\ell]$ to every vertex $v \in V$. We call an ℓ-coloring* proper *if it holds for every edge $e = \{u, v\}$ that $c(u) \neq c(v)$. The* color class *$V(h) \subseteq V$ of color $h \in [\ell]$ is the subset of vertices colored with color h, formally, $V(h) := \{v \in V \mid c(v) = h\}$. We say that an ℓ-coloring is* equitable *if it is proper and for any two colors $h, h' \in [\ell]$, it holds that $|V(h)| \leq |V(h')| + 1$.*

For some given $\ell \in \mathbb{N}$, the Equitable Coloring problem asks whether a given graph admits an equitable ℓ-coloring.

Equitable Coloring
Input: A graph $G = (V, E)$ and a number $\ell \in \mathbb{N}$.
Question: Is there an equitable ℓ-coloring of G?

The NP-hardness of Equitable Coloring can be shown by a straightforward reduction (outlined below) from Graph Coloring, which is defined as follows and is known to be NP-hard (see, e.g., Garey and Johnson [24]).

Graph Coloring
Input: A graph $G = (V, E)$ and a number $\ell \in \mathbb{N}$.
Question: Is there a proper ℓ-coloring of G?

We shortly sketch the reduction from Graph Coloring to Equitable Coloring. Given a graph $G = (V, E)$ and $\ell \in \mathbb{N}$, we construct a graph G' from G by adding $(\ell - 1) \cdot |V|$ isolated vertices. Clearly, G admits a proper ℓ-coloring if and only if G' admits an equitable ℓ-coloring: Any proper ℓ-coloring of G can be extended to an equitable ℓ-coloring of G' by coloring the added isolated vertices such that every color class has size $|V|$. Conversely, any equitable ℓ-coloring of G' needs to induce a proper ℓ-coloring when restricted to G.

Now, we are ready to prove that EF1-init Existence is NP-complete. Our reduction is inspired by a reduction due to Hosseini et al. [30][11], who reduce from a variant of Equitable Coloring where all color classes are required to be of the same size to proof NP-hardness of deciding the existence of an envy-free allocation (without initial utilities) for binary utilities[12].

Proof of Theorem 3.4. First, we observe that EF1-init Existence is in NP: Given an allocation $\mathcal{X}$, we can check for any pair of agents $i, j \in A$ whether agent i envies j under EF1-init. To show NP-hardness, we reduce from Equitable Coloring. Given a graph $G = (V, E)$ and a number $\ell \in \mathbb{N}$, we construct an instance for EF1-init Existence with agents A and resources R as follows:

- The set of agents A is the union of two sets A_E and A_C. The set A_E contains an *edge agent* e for each edge $e \in E$. A_C contains a *color agent* h for every color $h \in [\ell]$.
- For every vertex $v \in V$, the set R contains a *vertex resource* v.
- The utility functions and the initial utilities of the agents are defined as follows. All color agents from A_C have an initial utility of 0 and get a utility of 1 for every resource, that is, $u_i(X) = |X|$ for any $i \in A_C$ and $X \subseteq R$. All edge agents $e \in A_E$ have an initial utility of $|V| + 1$. An edge agent e for edge $e = \{u, v\} \in E$ gets a utility of $|V| + 2$ for the two vertex resources corresponding to edge e's endpoints u and v. All other resources do not increase the utility of e. Formally, $u_e(X) = |e \cap X| \cdot (|V| + 2)$ for any $X \subseteq R$.

This construction can clearly be computed in polynomial-time. It remains to prove that G admits an equitable ℓ-coloring if and only if the constructed instance has a complete EF1-init allocation.

($\Rightarrow$): Given an equitable ℓ-coloring of G, we construct an allocation $\mathcal{X}$ by assigning the vertex resource v to the color agent $c(v)$. As every vertex is colored, the constructed allocation $\mathcal{X}$ is complete. Furthermore, note that only color agents receive resources.

Next, we show that $\mathcal{X}$ is EF1-init. Consider any two agents $i, j \in A$. If $j \in A_E$, then, as $X_j = \emptyset$, agent i does not envy agent j under EF1-init. Next, suppose that $i \in A_E$ and $j \in A_C$. If $X_j = \emptyset$, then agent i does not envy agent j under EF1-init. Otherwise, we choose a resource $r \in X_j$ as follows. First, assume that $X_j \cap i \neq \emptyset$, that is, one of the endpoints of the edge corresponding to i is assigned color j. Then, we choose a resource $r \in X_j \cap i$, otherwise, we choose an arbitrary resource $r \in X_j$. Note that it needs to hold that $|X_j \cap i| \leq 1$, as the coloring would otherwise not be proper. Thus, we have that $(X_j \setminus \{r\}) \cap i = \emptyset$ and thus $b_j + u_i(X_j \setminus \{r\}) = 0$. Therefore, agent i does not envy agent j under EF1-init.

[11]The reduction can be found in the full version of their paper [29].

[12]Agents either approve or disapprove a resource. The utility for a bundle $X \subseteq R$ is the number of approved resources in X.

Finally, assume that $i, j \in A_C$. Since the coloring is equitable, we have that $|X_j| \leq |X_i| + 1$. Therefore, either $X_j = \emptyset$ or it holds for any $r \in X_j$ that

$$b_i + u_i(X_i) = u_i(X_i) = |X_i| \geq |X_j| - 1 \geq b_j + u_i(X_j \setminus \{r\}),$$

so agent i does not envy agent j under EF1-init. As we have exhausted all possible cases, it follows that $\mathcal{X}$ is EF1-init.

($\Leftarrow$): Let $\mathcal{X}$ be any complete EF1-init allocation for the constructed instance. We first prove that no edge agent can have a resource in $\mathcal{X}$. Suppose for the sake of contradiction that there is an edge agent $j \in A_E$ with $X_j \neq \emptyset$. Pick any color agent $i \in A_C$. Note that for all $r \in X_j$, it holds that

$$b_i + u_i(X_i) < |V| < |V| + 1 = b_j \leq b_j + u_i(X_j \setminus \{r\}).$$

Thus, agent j envies agent i under EF1-init. Therefore, only color agents can be allocated resources in $\mathcal{X}$.

Next, we construct an ℓ-coloring in G from the allocation $\mathcal{X}$ by coloring every vertex $v \in V$ with the color corresponding to the color agent $h \in A_C$ such that $v \in X_h$. As no edge agent can have a resource and $\mathcal{X}$ is complete, every vertex is assigned a color. First, we argue that the coloring is proper. Consider any edge $e = \{u, v\} \in E$. Suppose that u and v have the same color h. Then, $e \subseteq X_h$. However, this implies that the edge agent for e envies the color agent h under EF1-init, as $X_h \neq \emptyset$, and for all $r \in X_h$,

$$b_e + u_e(X_e) = |V| + 1 < |V| + 2 \leq b_h + u_e(X_h \setminus \{r\}).$$

This contradicts that $\mathcal{X}$ is EF1-init, therefore the constructed coloring is proper.

It remains to show that the coloring is also equitable. Assume for the sake of contradiction that there are two colors $h, k \in [\ell]$ such that $|V(h)| > |V(k)| + 1$. However, this implies that the color agent k envies color agent h under EF1-init. We have that $X_h \neq \emptyset$, and for all $r \in X_h$, it holds that

$$b_k + u_k(X_k) = |X_k| < |X_h| - 1 = b_h + u_k(X_h \setminus \{r\}).$$

Thus, the constructed coloring is an equitable ℓ-coloring of G, which completes the proof. $\square$

Next, we show that the hardness result cannot be extended for a constant number of agents, assuming a unary encoding of the utility functions.

Proposition 3.5. *For a unary encoding of the utility functions and a constant number of agents, EF1-INIT EXISTENCE and EF-INIT EXISTENCE are polynomial-time solvable.*

Proof. Both problems can be solved by a similar dynamic programming approach for a constant number of agents $n = |A|$. We first prove the statement for EF-init, and then show how the proof can be adapted for EF1-init. We fix an arbitrary ordering of the resources, i.e., $R = \{r_1, r_2, \dots, r_m\}$. Since we assume a unary encoding of the utility functions, without loss of generality, we assume that $u_i(X) \in \mathbb{N}$ for all $i \in A$ and $X \subseteq R$, and define $s := \max_{i \in A} u_i(R)$.

The dynamic program for EF-init is based on a boolean table $T[\ell, (v_{i,j})_{i,j=1}^n]$, where $\ell \in [m]$ and $v_{i,j} \in [s]$ for all $i, j \in A$. An entry should be set to "true" if there is an allocation $\mathcal{X}$ of the resources $\{r_1, \dots, r_\ell\}$, where it holds that $u_i(X_j) = v_{i,j}$ for all $i, j \in A$. Note that as n is constant and s is polynomial in the input size since we assume a unary encoding, it follows that the size of the table in $\mathcal{O}(m \cdot s^{n^2})$ is polynomial in the size of the input. We initialize the entry for $\ell = 0$ and $v_{i,j} = 0$ for all $i, j \in A$ to "true" and all other entries to "false". The table entries are filled in ascending order for $\ell \in \{1, \dots, m\}$ as follows. To fill an entry for given indices $\ell \in \{1, \dots, m\}$ and $v_{i,j} \in [s]$ with $i, j \in A$, check if there exists an agent $i^* \in A$, such that the entry for $\ell' = \ell - 1$, $v'_{j,i^*} = v_{j,i^*} - u_j(\{r_\ell\})$

for all $j \in A$, and $v'_{i,j} = v_{i,j}$ for all $i \in A$ and $j \neq i^*$ is set to "true". Intuitively, this means that the ℓ-th resource is given to agent i^*, and it remains to check if there is an (already computed) entry that represents that the first $\ell - 1$ resources can be distributed accordingly. Clearly, each entry can be computed in polynomial time. After computing the table, to decide EF-init Existence, check if there is an entry for $\ell = m$, where for all $i, j \in A$, it holds that $b_i + v_{i,i} \geq b_j + v_{i,j}$ or $v_{i',j} = 0$ for all $i' \in A$ (note that this implies that $X_j = \emptyset$ in the corresponding allocation $\mathcal{X}$).

To adapt the approach for EF1-init, we introduce additional dimensions $p_{i,j}$ for each pair of agents $i, j \in A$ used to store the value of the resource in j's bundle that i likes the most, yielding a table $T[\ell, (v_{i,j})_{i,j=1}^n, (p_{i,j})_{i,j=1}^n]$, where $\ell \in [m]$ and $v_{i,j}, p_{i,j} \in [s]$ for all $i, j \in A$. We initialize the entry for $\ell = 0$ and $v_{i,j} = p_{i,j} = 0$ for all $i, j \in A$ to "true" and all other entries to "false". Again, table entries are filled in ascending order for $\ell \in \{1, \ldots, m\}$. Analogously as before, to fill an entry for given indices $\ell \in \{1, \ldots, m\}$ and $v_{i,j}, p_{i,j} \in [s]$ with $i, j \in A$, check that there is an agent $i^* \in A$ such that it holds that $p_{j,i^*} \geq u_j(\{r_\ell\})$ for all $j \in A$, and such that there is a "true"-entry for i^* with the following indices:

- $\ell' = \ell - 1$,
- $v'_{j,i^*} = v_{j,i^*} - u_j(\{r_\ell\})$ for all $j \in A$, and $v'_{i,j} = v_{i,j}$ for all $i \in A$ and $j \neq i^*$,
- $p'_{j,i^*} = p_{j,i^*}$ if $p_{j,i^*} > u_j(\{r_\ell\})$ and otherwise $p'_{j,i^*} \in [p_{j,i^*}]$ for all $j \in A$,
- $p'_{i,j} = p_{i,j}$ for all $i \in A$ and $j \neq i^*$.

Again, it is clear that the table size and the time to compute each entry is polynomial in the input size. To decide EF1-init Existence, after filling the table, check if there is an entry for $\ell = m$, where for all $i, j \in A$, it holds that $b_i + v_{i,i} \geq b_j + v_{i,j} - r_{i,j}$ or $v_{i',j} = 0$ for all $i' \in A$. □

Finally, we give the complete proof for Theorem 3.6, which we sketched in the main section of the paper.

Theorem 3.6. *EF-init Existence for identical resources can be decided in $\mathcal{O}(n^2 \cdot m^3)$ time and $\mathcal{O}(n \cdot m^2)$ space.*

For agent $i \in A$, let $v_i := u_i(\{r\})$ for any $r \in R$ be the *value* that agent i derives from a single resource. We begin by formalizing the following observations, which we informally introduced already in the main part of the paper.

Observation C.2. *Let $h \in [t]$. Let $\mathcal{X}$ be an allocation where there is no envy between any pair of agents $i, j \in \bigcup_{\ell \in [h]} L_\ell$ in the first h levels. Then, the following holds.*

1. *If for some $h' \in [h]$, there is an agent $i \in L_{h'}$ with $|X_i| \geq 1$, then $|X_j| \geq 1$ for all $j \in L_{h''}$ with $h'' \leq h'$.*
2. *For any level $L_{h'}$ with $h' \in [h]$, it holds that any two agents $i, j \in L_{h'}$ are allocated the same number of resources $|X_i| = |X_j|$.*
3. *Let $i, j \in \bigcup_{\ell \in [h]} L_\ell$ be two agents in the first h levels. If $b_i < b_j$ and $v_i < v_j$, we have that $|X_j| = 0$.*

Proof. It is easy to see that the first two properties need to hold. To see why the third property holds, consider two agents $i, j \in \bigcup_{\ell \in [h]} L_\ell$ with $b_i < b_j$ and $v_i < v_j$. Now, assume for the sake of contradiction that $|X_j| > 0$. We have that $|X_i| > 0$, since otherwise agent i envies agent j under EF-init, as i has a lower initial utility. Since agent i does not envy j under EF-init in $\mathcal{X}$, we have that

$$\begin{aligned} & b_i + u_i(X_i) \geq b_j + u_i(X_j) \\ \iff & b_i + v_i \cdot |X_i| \geq b_j + v_i \cdot |X_j| \\ \iff & |X_i| - |X_j| \geq \frac{b_j - b_i}{v_i}. \end{aligned}$$

However, since agent j does not envy i under EF-init, we need to have that

$$\begin{aligned} & b_j + u_j(X_j) \geq b_i + u_j(X_i) \\ \iff & b_j + v_j \cdot |X_j| \geq b_i + v_j \cdot |X_i| \\ \iff & \frac{b_j - b_i}{v_j} \geq |X_i| - |X_j|. \end{aligned}$$

Note that this is a contradiction: Since $b_j - b_i > 0$ and $v_i < v_j$, we have that

$$|X_i| - |X_j| \leq \frac{b_j - b_i}{v_j} < \frac{b_j - b_i}{v_i} \leq |X_i| - |X_j|.$$

Therefore, it needs to hold that $|X_j| = 0$. $\square$

Note that for $h = t$, the allocation $\mathcal{X}$ from Observation C.2 is EF-init. Then, the properties hold for any level and all agents.

As defined in the main part of the paper, let $h^* \in [t]$ be the minimal level such that there is a violating pair of agents $i, j \in A$ with $b_i < b_j$ and $v_i < v_j$ where $j \in L_{h^*}$ (if such a pair exists). Then, by Observation C.2, it needs to hold for any EF-init allocation $\mathcal{X}$ and agent $i' \in L_{h'}$ with $h' \geq h^*$ that $|X_{i'}| = 0$. That is, no agent outside the first $h^* - 1$ levels gets a resource in any EF-init allocation. If $h^* = 1$, we know that the given instance for EF-init Existence is a no-instance, since there cannot be a complete EF-init allocation.

Otherwise, we define $t^* := h^* - 1$. Let $i \in \bigcup_{\ell \in \{h^*, \ldots, t\}} L_\ell$ be an agent outside the first t^* levels. For every level $L_{h'}$ with $h' \in [t^*]$, we know that an agent $j \in L_{h'}$ can get at most $\lfloor \frac{b_i - b_j}{v_i} \rfloor$ resources, as otherwise agent i would envy j under EF-init. By defining $k_{h'}$ as the minimum of these values for all $i \in \bigcup_{\ell \in \{h^*, \ldots, t\}} L_\ell$, we get an upper-bound on the number of allocated resources for each agent in the first t^* levels in any EF-init allocation. If there is no violating pair in the whole instance, we set $t^* := t$ and $k_{h'} := m$ for all $h' \in [t]$.

Now, we can show the following crucial lemma.

Lemma 3.7. *Let $h, h', h'' \in [t^*]$ with $h' < h < h''$ and consider agents $i^* \in L_h$, $i \in L_{h'}$, and $j \in L_{h''}$. If in some allocation $\mathcal{X}$ there is no envy between agents i and i^*, and between i^* and j under EF-init, then there is no envy between i and j under EF-init.*

Proof. We prove the statement by proving the following two implications.

1. If i does not envy i^* and i^* does not envy j, then i does not envy j (under EF-init).
2. If j does not envy i^* and i^* does not envy i, then j does not envy i (under EF-init).

Firstly, assume that i does not envy i^* and i^* does not envy j under EF-init. If $X_{i^*} = \emptyset$, then also $X_j = \emptyset$, as i^* does not envy j under EF-init and $b_j > b_{i^*}$, which implies that i does not envy j under EF-init. Thus, we assume that $X_{i^*} \neq \emptyset$ and $X_j \neq \emptyset$. As i does not envy i^* under EF-init, we have that

$$\begin{aligned} b_i + u_i(X_i) \geq b_{i^*} + u_i(X_{i^*}) &\iff b_i + v_i \cdot |X_i| \geq b_{i^*} + v_i \cdot |X_{i^*}| \\ &\iff |X_i| - |X_{i^*}| \geq \frac{b_{i^*} - b_i}{v_i}. \end{aligned}$$

Furthermore, agent i^* does not envy j under EF-init, so

$$\begin{aligned} b_{i^*} + u_{i^*}(X_{i^*}) \geq b_j + u_{i^*}(X_j) &\iff b_{i^*} + v_{i^*} \cdot |X_{i^*}| \geq b_j + v_{i^*} \cdot |X_j| \\ &\Rightarrow |X_{i^*}| - |X_j| \geq \frac{b_j - b_{i^*}}{v_{i^*}} \geq \frac{b_j - b_{i^*}}{v_i}, \end{aligned}$$

where the last inequality holds since $b_j - b_{i^*} > 0$ and $v_{i^*} \leq v_i$ (as there is no violating pair in the first t^* levels and $h' < h$). This implies that i does not envy j under EF-init, since

$$|X_i| - |X_j| = |X_i| - |X_{i^*}| + |X_{i^*}| - |X_j| \geq \frac{b_j - b_{i^*}}{v_i} + \frac{b_{i^*} - b_i}{v_i} = \frac{b_j - b_i}{v_i}$$
$$\iff b_i + v_i \cdot |X_i| \geq b_j + v_i \cdot |X_j|.$$

Secondly, assume symmetrically that j does not envy i^* and i^* does not envy i under EF-init. Since j does not envy i^* under EF-init, if $X_{i^*} \neq \emptyset$, then it holds that

$$b_j + u_j(X_j) \geq b_{i^*} + u_j(X_{i^*}) \iff b_j + v_j \cdot |X_j| \geq b_{i^*} + v_j \cdot |X_{i^*}|$$
$$\iff \frac{b_j - b_{i^*}}{v_j} \geq |X_{i^*}| - |X_j|.$$

If $X_{i^*} = \emptyset$, then $|X_{i^*}| - |X_j| \leq 0 < \frac{b_j - b_{i^*}}{v_j}$, so in any case, it holds that $\frac{b_j - b_{i^*}}{v_j} \geq |X_{i^*}| - |X_j|$.

Furthermore, we have that i^* does not envy i under EF-init. If $X_i = \emptyset$, then this directly implies that j does not envy i under EF-init. Otherwise, it follows that

$$b_{i^*} + u_{i^*}(X_{i^*}) \geq b_i + u_{i^*}(X_i) \iff b_{i^*} + v_{i^*} \cdot |X_{i^*}| \geq b_i + v_{i^*} \cdot |X_i|$$
$$\Rightarrow |X_i| - |X_{i^*}| \leq \frac{b_{i^*} - b_i}{v_{i^*}} \leq \frac{b_{i^*} - b_i}{v_j},$$

where the last inequality holds since $b_{i^*} - b_i > 0$ and $v_{i^*} \geq v_j$ (since there is no violating pair in the first t^* levels and $h'' > h$). This implies that j does not envy i under EF-init, since

$$|X_i| - |X_j| = |X_i| - |X_{i^*}| + |X_{i^*}| - |X_j| \leq \frac{b_{i^*} - b_i}{v_j} + \frac{b_j - b_{i^*}}{v_j} = \frac{b_j - b_i}{v_j}$$
$$\iff b_i + v_j \cdot |X_i| \leq b_j + v_j \cdot |X_j|.$$

□

Now, we are ready to prove that EF-init Existence for identical resources and additive utility functions can be decided in time polynomial in n and m.

Proof of Theorem 3.6. We prove the statement by giving a dynamic-programming algorithm based on the following table D of size $\mathcal{O}(n \cdot m^2)$. For $a \in [m]$, $b \in [t^*]$, and $c \in [m]$, we store in entry $D[a][b][c]$ whether there is an allocation $\mathcal{X}$ that satisfies the following requirements, which we call *fitting* for this entry. The table entry may be omitted when it is clear from context.

- Allocation $\mathcal{X}$ allocates exactly a resources,
- the set of agents that get at least one resource is the union of the first b levels $\bigcup_{\ell \in [b]} L_\ell$,
- every agent $i \in L_b$ gets exactly $c = |X_i|$ resources,
- there is no pair of agents $i, j \in \bigcup_{\ell \in [b]} L_\ell$ in the first b levels such that i envies j in $\mathcal{X}$ under EF-init,
- and it holds that $|X_i| \leq k_h$ for any $h \in [t^*]$ and agent $i \in L_h$ (recall that we defined these upper-bounds above when defining t^*).

Algorithm 2 Decide EF-init Existence

1: **Global Table:** $D[a][b][c]$, for $a \in [m]$, $b \in [t^*]$, and $c \in [m]$.
2: **for** $a \in [m]$, $b \in [t^*]$, and $c \in [m]$ **do**
3: $\quad D[a][b][c] \leftarrow$ "no" $\quad \triangleright$ Initialize Table D entries to "no".
4: **for** $a \in [m]$ and $c \in [k_1]$ **do**
5: $\quad$ **if** $|L_1| \cdot c = a$ **then**
6: $\quad\quad D[a][1][c] \leftarrow$ "yes" $\quad \triangleright$ Initialize entries for the first level.
7: **for** $b \in \{2, \dots, t^*\}$ **do**
8: $\quad$ **for** $a \in [m]$ and $c \in [k_b]$ **do**
9: $\quad\quad$ **if** $a - |L_b| \cdot c \geq 0$ **then**
10: $\quad\quad\quad a' \leftarrow a - |L_b| \cdot c$
11: $\quad\quad\quad b' \leftarrow b - 1$
12: $\quad\quad\quad$ Choose $i \in L_{b'}$ with minimal v_i and $j \in L_b$ with maximal v_j
13: $\quad\quad\quad c'_{min} \leftarrow c + \lceil \frac{b_j - b_i}{v_i} \rceil$ $\quad \triangleright$ Lower bound ensures that i does not envy j.
14: $\quad\quad\quad c'_{max} \leftarrow c + \lfloor \frac{b_j - b_i}{v_j} \rfloor$ $\quad \triangleright$ Upper bound ensures that j does not envy i.
15: $\quad\quad\quad$ **if** $c'_{min} \leq c'_{max}$ **then**
16: $\quad\quad\quad\quad$ **for** $c' \in \{c'_{min}, \dots, c'_{max}\}$ **do**
17: $\quad\quad\quad\quad\quad$ **if** $D[a'][b'][c'] =$ "yes" **then**
18: $\quad\quad\quad\quad\quad\quad D[a][b][c] \leftarrow$ "yes"
19: **for** $b \in [t^*]$ and $c \in [k_b]$ **do**
20: $\quad$ **if** $D[m][b][c] =$ "yes" and CheckEFInit(m, b, c) = "yes" **then**
21: $\quad\quad$ **Accept** $\quad \triangleright$ Accept if fitting allocation is EF-init.
22: **Reject**
23: **function** CheckEFInit(a, b, c)
24: $\quad$ Choose arbitrary $j \in L_b$.
25: $\quad$ **if** $b < t^*$ **then**
$\quad\quad \triangleright$ Check if an agent in levels $L_{b+1}, \dots, L_{t^*}$ (that receives no resources) envies j.
26: $\quad\quad$ **for** $i \in \bigcup_{h \in \{b+1, \dots, t^*\}} L_h$ **do**
27: $\quad\quad\quad$ **if** $c > \frac{b_j - b_i}{v_i}$ **then return** "no"
28: $\quad$ **return** "yes"

Next, we give Algorithm 2 to decide EF-init Existence. The algorithm first fills table D and then checks if there is a "yes"-entry for which a fitting allocation is EF-init. Note that for a fitting allocation for an entry for level $b \in [t^*]$, we only require that there is no envy under EF-init between agents $i, j \in \bigcup_{\ell \in [b]} L_\ell$ in the first b levels, so we have to check if agents in higher levels are envious. To prove the correctness of Algorithm 2, we first prove that table D has been filled correctly when we check if there is a fitting EF-init allocation for a "yes"-entry after Line 19.

Claim C.3. *After Line 19, for any $a \in [m]$, $b \in [t^*]$, and $c \in [m]$, it holds that entry $D[a][b][c] =$ "yes" if and only if there is a fitting allocation for entry $D[a][b][c]$.*

Proof. Note that for any $a \in [m]$ and $b \in [t^*]$, a fitting allocation can only exist if $c \leq k_b$. Moreover, in Algorithm 2, for any $b \in [t^*]$, we only introduce "yes"-entries for $c \leq k_b$. Thus, for all $b \in [t^*]$, entries with $c > k_b$ are correctly set to "no", and it suffices to consider entries with $c \in [k_b]$. We prove this claim by induction over $b \in [t^*]$.

First, we consider the base case of $b = 1$. Note that for any $a \in [m]$ and $c \in [k_1]$, a fitting allocation $\mathcal{X}$ for $D[a][1][c]$ allocates exactly a resources only to agents in L_1, $|X_i| = c$ for every agent $i \in L_1$ and $\mathcal{X}$ satisfies that there is no envy between any pair $i, j \in L_1$ under EF-init. Clearly, there is such an allocation if and only if $|L_1| \cdot c = a$. As we set $D[a][1][c] =$ "yes" if and only if $|L_1| \cdot c = a$ in Line 6, the base case holds.

For the induction step, assume that for some $1 \leq h < t^*$, the claim holds for all table entries $D[a][b][c]$ with $b \leq h$, $a \in [m]$ and $c \leq k_b$. We need to show that for any $a \in [m]$ and $c \in [k_{h+1}]$, entry $D[a][h+1][c] =$ "yes" if and only if there is a fitting allocation $\mathcal{X}$ for this entry.

($\Leftarrow$): Recall that we set $D[a][h+1][c] =$ "yes" in Line 18 if we find a suitable "yes"-entry for the preceding level L_h. Suppose that there is a fitting allocation $\mathcal{X}$ for $D[a][h+1][c]$. In $\mathcal{X}$, every agent in level $h+1$ gets exactly c resources and agents in $A \setminus \bigcup_{\ell \in [h+1]} L_\ell$ do not receive any resources. Next, we will construct a fitting allocation $\mathcal{X}'$ for an entry in the previous level h. Clearly, by setting $X'_i = \emptyset$ for all $i \in L_{h+1}$ and $X'_i = X_i$ for all $i \notin L_{h+1}$, allocation $\mathcal{X}$ induces an allocation $\mathcal{X}'$ that allocates $a' = a - |L_{h+1}| \cdot c \geq 0$ resources and where the set of agents that get a resource is $\bigcup_{\ell \in [h]} L_\ell$. Furthermore, in $\mathcal{X}'$, there cannot be any envy under EF-init between agents $i, j \in \bigcup_{\ell \in [h]} L_\ell$, as these agents would also be envious in $\mathcal{X}$. Finally, by Observation C.2, in $\mathcal{X}'$, all agents in L_h need to get the same number c' of resources. It needs to hold that $c' \leq k_h$, since $\mathcal{X}$ is fitting. Thus, $\mathcal{X}'$ is fitting for $D[a'][h][c']$. By the induction assumption, this implies that $D[a'][h][c'] =$ "yes". As in the algorithm, choose an agent $i^* \in L_h$ with minimal v_{i^*} and $j^* \in L_{h+1}$ with maximal v_{j^*}, and let $c'_{min} = c + \lceil \frac{b_{j^*} - b_{i^*}}{v_{i^*}} \rceil$ and $c'_{max} = c + \lfloor \frac{b_{j^*} - b_{i^*}}{v_{j^*}} \rfloor$. Note that since i^* does not envy j^* under EF-init in $\mathcal{X}$, it needs to hold that

$$\begin{aligned} & b_{i^*} + u_{i^*}(X_{i^*}) \geq b_{j^*} + u_{i^*}(X_{j^*}) \\ \iff & b_{i^*} + v_{i^*} \cdot |X'_{i^*}| \geq b_{j^*} + v_{i^*} \cdot c \\ \iff & |X'_{i^*}| - c \geq \frac{b_{j^*} - b_{i^*}}{v_{i^*}} \iff |X'_{i^*}| \geq c'_{min}. \end{aligned}$$

Furthermore, j^* does not envy i^* under EF-init in $\mathcal{X}$, so we have that

$$\begin{aligned} & b_{j^*} + u_{j^*}(X_{j^*}) \geq b_{i^*} + u_{j^*}(X_{i^*}) \\ \iff & b_{j^*} + v_{j^*} \cdot c \geq b_{i^*} + v_{j^*} \cdot |X'_{i^*}| \\ \iff & \frac{b_{j^*} - b_{i^*}}{v_{j^*}} \geq |X'_{i^*}| - c \iff |X'_{i^*}| \leq c'_{max}. \end{aligned}$$

Note that as shown above, $c'_{min} \leq c' = |X'_{i^*}| \leq c'_{max}$. Thus, we consider the "yes"-entry for $\mathcal{X}'$ in the algorithm and set entry $D[a][h+1][c] =$ "yes" in Line 18.

($\Rightarrow$): Conversely, assume that $D[a][h+1][c] =$ "yes". Let $a' = a - |L_{h+1}| \cdot c$. Again, choose $i^* \in L_h$ with minimal v_{i^*} and $j^* \in L_{h+1}$ with maximal v_{j^*}, and let $c'_{min} = c + \lceil \frac{b_{j^*} - b_{i^*}}{v_{i^*}} \rceil$ and $c'_{max} = c + \lfloor \frac{b_{j^*} - b_{i^*}}{v_{j^*}} \rfloor$. As we only set the entry to "yes" in Line 18, there needs to exist some $c' \in \{c'_{min}, \dots, c'_{max}\}$ such that $D[a'][h][c'] =$ "yes". By the induction assumption, this implies that there is a fitting allocation $\mathcal{X}$ for $D[a'][h][c']$. Extending $\mathcal{X}$ by allocating c resources to every agent in level $h+1$ yields an allocation $\mathcal{X}'$ that allocates exactly $a' + c \cdot |L_{h+1}| = a$ resources and where the set of agents that receive a resource is $\bigcup_{\ell \in [h+1]} L_\ell$. It remains to prove that there is no envy under EF-init between any agents $i, j \in \bigcup_{\ell \in [h+1]} L_\ell$ in $\mathcal{X}'$. Clearly, there cannot be a pair of agents $i, j \in \bigcup_{\ell \in [h]} L_\ell$ where i envies j, as i would also envy j (under EF-init) in $\mathcal{X}$. Next, consider two agents $i \in L_h$ and

$j \in L_{h+1}$. Note that

$$
\begin{aligned}
&|X_i'| = c' \geq c'_{min} = c + \lceil \frac{b_{j^*} - b_{i^*}}{v_{i^*}} \rceil \geq |X_j'| + \lceil \frac{b_j - b_i}{v_i} \rceil \\
\Rightarrow &|X_i'| - |X_j'| \geq \frac{b_j - b_i}{v_i} \\
\iff &b_i + |X_i'| \cdot v_i \geq b_j + |X_j'| \cdot v_i,
\end{aligned}
$$

so agent i does not envy j under EF-init. Furthermore, we have that

$$
\begin{aligned}
&|X_i'| = c' \leq c'_{max} = c + \lfloor \frac{b_{j^*} - b_{i^*}}{v_{j^*}} \rfloor \leq |X_j'| + \lfloor \frac{b_j - b_i}{v_j} \rfloor \\
\Rightarrow &\frac{b_j - b_i}{v_j} \geq |X_i'| - |X_j'| \\
\iff &b_j + |X_j'| \cdot v_j \geq b_i + |X_i'| \cdot v_j,
\end{aligned}
$$

so agent j does not envy i under EF-init. It remains to prove that there is no envy under EF-init between an agent $i \in L_{h'}$ with $h' < h$ and an agent $j \in L_{h+1}$. Consider an agent $j' \in L_h$. We have that there is no envy between agents i and j', as there is no envy between any agents in $\bigcup_{\ell \in [h]} L_\ell$ (under EF-init), since $\mathcal{X}$ is fitting. Moreover, as argued above, there is no envy between agent $j' \in L_h$ and $j \in L_{h+1}$. With Lemma 3.7, it follows that there is no envy between agent $i \in L_{h'}$ with $h' < h$ and $j \in L_{h+1}$. Thus, $\mathcal{X}'$ is fitting for $D[a][h+1][c]$. ♦

Now, we are ready to prove the correctness of the algorithm in the following claim.

Claim C.4. *Algorithm 2 accepts if and only if there is a complete EF-init allocation.*

Proof. We prove the claim by proving both directions of the equivalence.

(⇒): First, suppose that Algorithm 2 accepts. Note that we only accept in Line 21 if there is an entry $D[m][b][c] =$ "yes" for some $b \in [t^*]$ and $c \in [m]$. By Claim C.3, there exists a fitting allocation $\mathcal{X}$ for $D[m][b][c]$. Clearly, this allocation is complete.

Note that no agent in $A \setminus \bigcup_{\ell \in [b]} L_\ell$ gets a resource in $\mathcal{X}$. Thus, only agents in $\bigcup_{\ell \in [b]} L_\ell$ can be envied under EF-init by any agent. Since $\mathcal{X}$ is fitting for $D[m][b][c]$, no agent $i \in \bigcup_{\ell \in [b]} L_\ell$ can be envious of another agent under EF-init. Therefore, we only need to check if there is an agent $i \in A \setminus \bigcup_{\ell \in [b]} L_\ell$ that envies an agent $j \in \bigcup_{\ell \in [b]} L_\ell$ under EF-init. First, observe that no agent $i \in \bigcup_{\ell \in \{t^*, \dots, t\}} L_\ell$ envies an agent $j \in \bigcup_{\ell \in [b]} L_\ell$ under EF-init: Consider some agent $i \in \bigcup_{\ell \in \{t^*, \dots, t\}} L_\ell$ and an agent $j \in L_\ell$ for some $\ell \in [b]$. Since $\mathcal{X}$ is fitting, the allocation needs to satisfy the upper-bounds on the number of resources for agents in the first t^* levels that we defined in the beginning. Formally, it needs to hold that $|X_j| \leq k_\ell$. Thus, we have that

$$
|X_j| \leq k_\ell \leq \frac{b_i - b_j}{v_i} \iff b_i \geq b_j + v_i \cdot |X_j| \iff b_i + u_i(X_i) \geq b_j + u_i(X_j),
$$

so agent i does not envy j under EF-init. Note that if $b = t^*$, we already have shown that there is no envy between any pair of agents under EF-init. Otherwise, if $b < t^*$, it remains to consider an agent $i \in \bigcup_{\ell \in \{b+1, \dots, t^*\}} L_\ell$. Let $j^* \in L_b$ be an agent in L_b. Recall that j^* does not envy any agent $j \in \bigcup_{\ell \in [b]} L_\ell$ under EF-init, since $\mathcal{X}$ is fitting. Thus, if i does not envy j^*, by Observation C.2, i does not envy j (under EF-init). Note that i does not envy j^* under EF-init if and only if $c = |X_{j^*}| \leq \frac{b_{j^*} - b_i}{v_i}$, as

$$
|X_{j^*}| \leq \frac{b_{j^*} - b_i}{v_i} \iff b_{j^*} + v_i \cdot |X_{j^*}| \leq b_i \iff b_{j^*} + u_i(X_{j^*}) \leq b_i + u_i(X_i).
$$

Since we accept in Line 21 only if this inequality holds for all agents $i \in \bigcup_{\ell \in \{b+1, \dots, t^*\}} L_\ell$, allocation $\mathcal{X}$ is EF-init.

(⇐): Suppose there is a complete EF-init allocation $\mathcal{X}$. If the algorithm accepts, we are done. In the following, assume for the sake of contradiction that the algorithm does not accept. Note that it needs to hold for all $h \in [t^*]$ and $i \in L_h$ that $|X_i| \leq k_h$ as argued in the beginning: If $t^* < t$, all agents in $\bigcup_{\ell\in\{t^*+1,\ldots,t\}} L_\ell$ cannot get a resource in $\mathcal{X}$. Therefore, if $|X_i| > k_h$ for some $h \in [t^*]$ and $i \in L_h$, then there is some agent $j \in \bigcup_{\ell\in\{t^*+1,\ldots,t\}} L_\ell$ that envies i under EF-init. If $t^* = t$, we set $k_h = m$ (recall that this is the case if there is no violating pair in the instance), so $|X_i| \leq k_h = m$ holds trivially.

Let h be the maximal level $h \in [t]$ such that an agent $i \in L_h$ receives a resource in $\mathcal{X}$. Note that $h \leq t^*$, since only agents in $\bigcup_{\ell\in[t^*]} L_\ell$ can get a resource in an EF-init allocation. Moreover, by Observation C.2, all agents in $\bigcup_{\ell\in[t]} L_\ell$ need to get at least one resource and every agent in L_h needs to get exactly $|X_i|$ resources in $\mathcal{X}$. Clearly, since $\mathcal{X}$ is EF-init, there is no envy under EF-init between any pair of agents in $\mathcal{X}$. It follows that $\mathcal{X}$ is fitting for $D[m][h][|X_i|]$ and thus $D[m][h][|X_i|] =$ "yes" by Claim C.3.

If $h = t^*$, CheckEFInit($m, h, |X_i|$) returns "yes" and we accept in Line 21, which is a contradiction. Thus, assume that $h < t^*$ and consider an agent $i' \in \bigcup_{\ell\in\{h+1,\ldots,t^*\}} L_\ell$ and an agent $j^* \in L_h$. As argued above, i' does not envy j^* under EF-init if and only if $|X_{j^*}| \leq \frac{b_{j^*}-b_{i'}}{v_{i'}}$. Thus, this inequality needs to hold for all agents $i' \in \bigcup_{\ell\in\{h+1,\ldots,t^*\}} L_\ell$ in allocation $\mathcal{X}$. Therefore, CheckEFInit($m, h, |X_i|$) = "yes". Again, this implies that we accept in Line 21, which is a contradiction. Therefore, the algorithm needs to accept either for this entry or for an entry that was checked earlier. ◆

Finally, it remains to show that Algorithm 2 runs in $\mathcal{O}(n^2 \cdot m^3)$ time and $\mathcal{O}(n \cdot m^2)$ space. Recall that we proceed in two steps: Firstly, we partition the agents into the levels, compute $t^* \in [t]$ and the upper-bounds k_h for all $h \in [t^*]$. Secondly, we decide EF-INIT EXISTENCE using Algorithm 2.

Partitioning the agents into the levels and computing t^* and the upper-bounds can be done in time $\mathcal{O}(n^2 \cdot m^3)$. Moreover, we need to store at most $t^* \leq n$ upper-bounds. Next, we consider Algorithm 2. First, since $t \leq n$, the table D has size $\mathcal{O}(n \cdot m^2)$. Clearly, the initialization of the table within Line 2 to Line 4 can be done in running time $\mathcal{O}(n \cdot m^3)$. To compute a table entry for level L_b with $b \in \{2, \ldots, t\}$, we need to look at at most m entries for level L_{b-1}. Note that we also need to find the agents $i \in L_{b-1}$ with minimal value v_i and $j \in L_b$ with maximal value v_j. However, for every one of the at most n levels, we can store the agents with minimal and maximal value when partitioning the agents into the levels. Thus, each one of the $\mathcal{O}(n \cdot m^2)$ entries can be computed in $\mathcal{O}(m)$, so the algorithm fills table D in time $\mathcal{O}(n \cdot m^3)$. Lastly, we analyze the running time of checking if there is a "yes"-entry with a fitting EF-init allocation within the loop in Line 19. First, observe that we check at most $\mathcal{O}(n \cdot m)$ entries. For every "yes"-entry, we check whether a fitting allocation is EF-init using function CheckEFInit, which takes time at most $\mathcal{O}(n \cdot m)$. □

Note that the running time and space consumption from above may be improved by a more careful analysis. For example, in Algorithm 2, the space consumption of table D can be improved by storing only the $\mathcal{O}(m^2)$ entries of the previous level. Whenever we introduce a "yes"-entry, we can directly check if a fitting allocation is also EF-init, in which case we can directly accept. Moreover, the algorithm can also be modified to compute an EF-init allocation by storing a fitting allocation for "yes"-entries. Note that it suffices to store an arbitrary fitting allocation for an entry: Extending a fitting allocation for an entry for level L_b with $b \in [t]$ to level L_{b+1} and checking if the allocation is EF-init only depends on the number of resources allocated to an agent in L_b, which is the same for all fitting allocations for this entry.

D Additional Material for Section 4

In this section, we provide the complete proof for Theorem 4.7.

Theorem 4.7. *The allocation computed by Algorithm 1 is complete and satisfies min-EF1-init. Further, Algorithm 1 runs in polynomial time.*

We start by formalizing properties of Algorithm 1 that we informally remarked in the main part, before restating the crucial Lemma 4.10 (Activation Gap Lemma). Using this, we then complete the proof of Theorem 4.7.

We observe the following.

Observation D.1. *Consider Algorithm 1 and let $\mathcal{X}$ be the allocation returned by the algorithm. Let $i, j \in A$ be two distinct agents, and for each $k \in \{i, j\}$, let X_k^* denote the bundle of k after a given pick in the algorithm. If either $X_j^* = \emptyset$ or there exists a resource $r \in X_j^*$ such that $b_i + u_i(X_i^*) \geq b_j + u_i(X_j^* \setminus \{r\})$, and after this pick agent i picks before j's next pick, then it holds that $b_i + u_i(X_i) \geq b_j + u_i(X_j \setminus \{r\})$ in the final allocation $\mathcal{X}$.*

Lemma 4.10 (Activation Gap Lemma). *Consider Algorithm 1 and let $i, j \in A$ be two agents that were active in some round with $b_i \geq b_j$ and $b_i > \min_{i' \in A} b_{i'}$. For $k \in A$, let X_k^* be the bundle of k after the pick that activated i. Then, either $X_j^* = \emptyset$ or there exists $r \in X_j^*$ such that*

$$\sum_{r' \in X_j^* \setminus \{r\}} \min_{j' \in A \colon b_{j'} < b_i} u_{j'}(\{r'\}) < b_i - b_j.$$

Using the above lemma, we can now prove Theorem 4.7.

Proof of Theorem 4.7. Let $\mathcal{X}$ be the allocation computed by the algorithm and let $\preceq$ be the final linear order. We show that for any two agents $i, j \in A$, agent i does not envy agent j under min-EF1-init in $\mathcal{X}$. Thus, we need to show that $X_j = \emptyset$ or condition (C1) holds if $b_i \leq b_j$ or (C2) holds if $b_i > b_j$. Therefore, in the following, assume that $X_j \neq \emptyset$.

Case 1 ($b_i = b_j$). In this case, we need to argue that there is a resource $r \in X_j$ such that condition (C1) holds. Note that i and j are active in the same rounds. Thus, condition (C1) holds by applying Observation 4.8 on the set of all rounds: At least one of the cases is applicable, and every case implies that $u_i(X_i) \geq u_i(X_j \setminus \{r\})$ for some $r \in X_j$.

Case 2 ($b_i < b_j$). For $k \in A$, let X_k^* be the bundle of k after j's first pick and let $r \in R$ be the first resource picked by j, i.e., $X_j^* = \{r\}$. Since i had reached j when agent j was activated, it holds that

$$b_i + u_i(X_i^*) \geq b_j + \underbrace{u_i(X_j^* \setminus \{r\})}_{=0}.$$

Moreover, since i picks before j in the picking sequence starting with the pick after j's first pick, agent i prefers the resources they pick in this picking sequence (i.e., $X_i \setminus X_i^*$) over those picked by j (i.e., $X_j \setminus X_j^*$). Thus, $u_i(X_i \setminus X_i^*) \geq u_i(X_j \setminus X_j^*)$. With Observation D.1, condition (C1) follows for resource r.

Case 3 ($b_i > b_j$). In this case, we need to argue that condition (C2) holds for some resource $r^* \in X_j$, that is, that there exists a subset $X^* \subseteq X_j$ with

$$\sum_{r \in X^*} \min_{j' \in A \colon b_{j'} < b_i} u_{j'}(\{r\}) < b_i - b_j, \tag{7}$$

such that it holds that $u_i(X_i) \geq u_i(X_j \setminus (X^* \cup \{r^*\}))$.

We make a case distinction whether agent i was activated at some point during the algorithm. First, suppose that agent i was activated. Let X^*_k be the bundle of agent $k \in A$ after the pick that activated i. Note that $b_i > b_j$ implies that $b_i > \min_{i' \in A} b_{i'}$, so by Lemma 4.10, we have that either $X^*_j = \emptyset$ or it holds that

$$\sum_{r' \in X^*_j \setminus \{r\}} \min_{j' \in A \colon b_{j'} < b_i} u_{j'}(\{r'\}) < b_i - b_j$$

for some resource $r \in X^*_j$. Note that it needs to hold that $X^*_j \neq \emptyset$ since $b_i > b_j$. Thus, let $X^* = X^*_j \setminus \{r\}$ and $r^* = r$. Note that X^* satisfies Equation (7). Next, consider the picking sequence of all picks starting with (including) agent i's first pick. In this sequence, agent i picks the resources in X_i and agent j picks the resource in $X_j \setminus X^*_j$. Since agent i picks first, we have that $u_i(X_i) \geq u_i(X_j \setminus X^*_j) = u_i(X_j \setminus (X^* \cup \{r^*\}))$, so X^* and r^* satisfy (C2).

Finally, suppose that agent i was not activated. In this case, we will argue that it is always possible to choose X^* and resource r^* in (C2) such that $X^* \cup \{r^*\} = X_j$, which implies that $u_i(X_i) \geq 0 = u_i(X_j \setminus (X^* \cup \{r^*\}))$ holds trivially. If agent i was not activated, there needs to be an agent $j^* \in A$ with $b_{j^*} < b_i$ such that

$$u_{j^*}(X_{j*}) < b_i - b_{j^*}. \tag{8}$$

If $j = j^*$, this implies that (C2) holds for $X^* = X_{j*}$ and any arbitrary resource $r^* \in X_{j*}$, since

$$b_i - b_j >_{(8)} u_{j^*}(X_{j*}) \geq \sum_{r' \in X_{j^*}} \min_{j' \in A \colon b_{j'} < b_i} u_{j'}(\{r'\}).$$

Thus, it remains to consider the case that $j \neq j^*$. First, we assume that $b_j = b_{j^*}$. Note that in this case, agents j and j^* are active in the same set of rounds $\mathcal{R}$. Thus, using Observation 4.8, we get that

$$u_{j^*}(X_j \setminus \{r\}) \leq_{(4.8)} u_{j^*}(X_{j^*}) <_{(8)} b_{j^*} - b_i = b_j - b_i$$

for some resource $r \in X_j$. Thus, condition (C2) holds for $X^* = X_j \setminus \{r\}$ and resource $r^* = r$, since $X^* \cup \{r\} = X_j$ and X^* satisfies Equation (7). We consider the remaining case that $b_j \neq b_{j^*}$.

First, assume that $b_j > b_{j^*}$. Let $X^*_{j^*}$ be the bundle of agent j^* when agent j was activated. It needs to hold that

$$u_{j^*}(X^*_{j^*}) \geq b_j - b_{j^*}. \tag{9}$$

From Equations (8) and (9), it follows that

$$u_{j^*}(X_{j^*} \setminus X^*_{j^*}) < b_i - b_j. \tag{10}$$

Let $r \in X_j$ be the first resource picked by agent j. Consider the picking sequence of all picks starting with the pick after agent j's first pick. In this picking sequence, agent j picks the resources in $X_j \setminus \{r\}$ and agent j^* picks the resources in $X_{j^*} \setminus X^*_{j^*}$. Since agent j^* picks before j in this picking sequence, we have that

$$u_{j^*}(X_{j^*} \setminus X^*_{j^*}) \geq_{(4.8)} u_{j^*}(X_j \setminus \{r\}). \tag{11}$$

With this, it follows that

$$b_i - b_j >_{(10)} u_{j^*}(X_{j^*} \setminus X^*_{j^*}) \geq_{(11)} u_{j^*}(X_j \setminus \{r\}) \geq \sum_{r' \in X_j \setminus \{r\}} \min_{j' \in A \colon b_{j'} < b_i} u_{j'}(\{r'\}).$$

This implies that condition (C2) holds for $X^* = X_j \setminus \{r\}$ and resource $r^* = r$.

Finally, assume that $b_j < b_{j^*}$. Let X_j^* be the bundle of agent j when agent j^* was activated. Note that it cannot be the case that $X_j^* = \emptyset$, since $b_j < b_{j^*}$. Moreover, $b_j < b_{j^*}$ implies that $b_{j^*} > \min_{i' \in A} b_{i'}$. By Lemma 4.10, we thus have that there exists a resource $r \in X_j^*$ such that

$$b_{j^*} - b_j > \sum_{r' \in X_j^* \setminus \{r\}} \min_{j' \in A \colon b_{j'} < b_{j^*}} u_{j'}(\{r'\}). \tag{12}$$

Consider the picking sequence starting with agent j^*'s first pick. In this sequence, agent j^* picks the resources in X_{j^*} and agent j picks the resources in $X_j \setminus X_j^*$. Since agent j^* picks first, we get that

$$u_{j^*}(X_j \setminus X_j^*) \leq u_{j^*}(X_{j^*}) <_{(8)} b_i - b_{j^*}. \tag{13}$$

It follows from Equations (12) and (13) that

$$\begin{aligned} b_i - b_j = (b_i - b_{j^*}) + (b_{j^*} - b_j) &>_{(13),(12)} u_{j^*}(X_j \setminus X_j^*) + \sum_{r' \in X_j^* \setminus \{r\}} \min_{j' \in A \colon b_{j'} < b_{j^*}} u_{j'}(\{r'\}) \\ &\geq \sum_{r' \in X_j \setminus \{r\}} \min_{j' \in A \colon b_{j'} < b_i} u_{j'}(\{r'\}). \end{aligned}$$

Thus, condition (C2) holds for $X^* = X_j \setminus \{r\}$ and resource $r^* = r$. This completes the proof. $\square$

Niclas Boehmer
Hasso-Plattner-Institute
Potsdam, Germany
Email: Niclas.Boehmer@hpi.de

Luca Kreisel
Hasso-Plattner-Institute
Potsdam, Germany
Email: Luca.Kreisel@hpi.de