Repeated Fair Allocation of Indivisible Items

Ayumi Igarashi, Martin Lackner, Oliviero Nardi and Arianna Novaro

Abstract

The problem of fairly allocating a set of indivisible items is a well-known challenge in the field of (computational) social choice. In this scenario, there is a fundamental incompatibility between notions of fairness (such as envy-freeness and proportionality) and economic efficiency (such as Pareto-optimality). However, in the real world, items are not always allocated once and for all, but often repeatedly. For example, the items may be recurring chores to distribute in a household. Motivated by this, we initiate the study of the repeated fair division of indivisible goods and chores and propose a formal model for this scenario. In this paper, we show that, if the number of repetitions is a multiple of the number of agents, we can always find (i) a sequence of allocations that is envy-free and complete (in polynomial time), and (ii) a sequence of allocations that is proportional and Pareto-optimal (in exponential time). On the other hand, we show that irrespective of the number of repetitions, an envy-free and Pareto-optimal sequence of allocations may not exist. For the case of two agents, we show that if the number of repetitions is even, it is always possible to find a sequence of allocations that is overall envy-free and Pareto-optimal. We then prove even stronger fairness guarantees, showing that every allocation in such a sequence satisfies some relaxation of envy-freeness.

1 Introduction

In a variety of real-life scenarios, a group of agents has to divide among themselves a set of items, that can be desirable (goods) or undesirable (chores), over which they have (heterogeneous) preferences. In the case of *goods*, we can think, for instance, of employees having to share access to some common infrastructure, such as meeting rooms or computing facilities. In the case of *chores*, we can think of roommates having to split household duties or teams having to split admin tasks. In some cases, we may have a set of *mixed items*, where the agents may consider some items good and others bad: for instance, when assigning teaching responsibilities, some courses may be desirable to teach for some while being undesirable (negative) for others.

The examples above are all instances of problems of fair allocation of indivisible items (see the handbook of Brandt et al. [7] for a survey). One of the challenges of fair division problems is that in many cases, given the preferences of the agents, it may be impossible to find an allocation of the items which is both fair (e.g., no agent envies the bundle received by another agent) and efficient (e.g., no other allocation would make some agents better off, without making anyone worse off). However, a crucial feature which has so far received little attention, is that many fair allocation problems have a *repeated* nature, in the sense that the same items will need to be assigned multiple times to the agents. For instance, university courses are usually offered every year, meeting rooms may be needed daily or weekly by the same teams of employees, and household chores need to be done on a regular basis.

In this paper, we thus focus precisely on those settings where a set of items has to be repeatedly allocated to a set of agents. This opens the field to new and exciting research directions, revolving around the following central question: "Can some fairness and efficiency notions, which may be impossible to achieve in the standard setting of fair division, be guaranteed when taking a global perspective on the overall sequence of repeated allocations?"

condition		result	
#agents (n)	#rounds (k)	fairness guarantee	reference
$n \geqslant 2$	$k \in n\mathbb{N}$	EF overall	Prop. 4
$n \geqslant 2$	$k \notin n\mathbb{N}$	~~PROP overall~~	Prop. 5
$n > 2$	$k \in \mathbb{N}$	~~EF+PO overall~~	Thm. 6
$n \geqslant 2$	$k \in n\mathbb{N}$	PROP+PO overall	Thm. 7
$n = 2$	even	EF+PO overall	Thm. 10
$n = 2$	$k > 2$	~~EF+PO overall and per-round EF1~~	Prop. 12
$n = 2$	$k = 2$	EF+PO overall and per-round EF1	Cor. 14
$n = 2$	even	EF+PO overall and per-round weak EF1	Cor. 16
$n = 2$	even	EF overall and per-round EF1	Thm. 17

Table 1: Overview of our results regarding envy-freeness (EF), envy-freeness up to one item (EF1), weak EF1, proportionality (PROP), and Pareto-optimality (PO). Crossed-out results cannot be guaranteed under the stated conditions. We write $k \in n\mathbb{N}$ to denote that the number of rounds is a multiple of the number of voters n.

Contribution and outline. Our main contribution is the definition of a new model for the repeated fair allocation of goods and chores, which we present in Section 2. In particular, we specify how sequences of allocations will be evaluated with respect to classical axioms of fairness and efficiency: we distinguish between sequences satisfying some axioms *per-round* (i.e., for every allocation composing them), or *overall* (i.e., when considering the collection of all the bundles every agent has received in the sequence), and prove some preliminary results relating these two variants.

In Section 3 we study our model for n agents and we show that: if the number of rounds is a multiple of n, we can always find both a sequence of allocations that is envy-free overall in polynomial time (Proposition 4), and a sequence of allocations that is proportional and Pareto-optimal in exponential time (Theorem 7). This is essentially optimal for the allocation of chores: for any number $n > 2$ of agents and any number k of rounds, there is an instance where an envy-free and Pareto-optimal sequence of allocations does not exist (Theorem 6).

In Section 4, we restrict our model to the case of two agents, and prove our main positive result: if the number of rounds is even, we can always find a sequence of allocations which is envy-free and Pareto-optimal overall, as well as per-round weak envy-free up to one item (Corollary 16). Moreover, for two rounds, we can strengthen the per-round fairness guarantee to envy-freeness up to one item (EF1) (Corollary 14). At the cost of sacrificing the efficiency requirement, we also show that we can always find a sequence of allocations which is envy-free and per-round EF1 in polynomial time (Theorem 17). These results turn out to be the best one can hope for. In Proposition 12, we show that there is an instance with two agents and $k > 2$ rounds where no sequence of allocations is envy-free and Pareto-optimal overall, as well as per-round EF1.

Finally, we conclude and give interesting directions for future work in Section 5. An overview of our results can be found in Table 1.

Related work. Aziz et al. [3] analyse fairness concepts for the allocation of indivisible goods and chores. Based on some of these results, Igarashi and Yokoyama [14] have also developed an app to help couples divide household chores fairly. Another relevant application is that of the fair allocation of papers to reviewers in peer-reviewed conferences [20, 22], as well as the allocation of students to courses under capacity constraints [21]. These works however all focus on "one-shot" problems, and not on repeated allocations.

Various settings fall under the umbrella of *dynamic* or *online* fair decision-making [1, 16]. Both Kash et al. [15] and Freeman et al. [12] study fair division problems where resources from a common pool are dynamically redistributed to a set of agents: the former assumes that the *agents* may join the process over time, while the latter assumes that their *demands* may vary over time. Benade et al. [5] and Zeng and Psomas [23] study a setting where items may arrive over time, and they aim at decreasing envy among agents as time goes by. Finally, Guo et al. [13] and Cavallo [9] focus on a repeated setting, where a *single* item must be allocated at every round. In our case agents have static and heterogeneous preferences, instead of demands, over multiple items, and the sets of agents and items is fixed. Balan et al. [4] also study the repeated allocation of items (i.e., courses assigned to professors), but they focus on the average of utilities received by the agents for a sequence of allocations.

In the context of elections, a closely related framework is that of *perpetual voting*, introduced by Lackner [17], where the agents participate in repeated elections to select a winning candidate and classical fairness axioms (as well as new ones) are introduced to evaluate aggregation rules with respect to sequences of elections. A similar approach has also been taken to analyse repeated instances of participatory budgeting problems [18]. Freeman et al. [11] consider a setting where at each round one alternative is selected, and agents' preferences may vary over time. A related model of simultaneous decisions, closer to fair division, has been studied by Conitzer et al. [10].

Finally, we can see our work as belonging to a recent trend of research in computational social choice where a set of profiles needs to each be assigned an outcome—see the recent overview paper by Boehmer and Niedermeier [6].

2 The Model

In this section, we present the model used throughout the paper. Furthermore, we recall some familiar concepts from the theory of fair division, and adapt them to our scenario.

We denote by N a finite set of n *agents*, who have to be assigned a set of m *items* in the finite set I. An allocation $\pi \subseteq N \times I$ consists of agent-item pairs (i, o), indicating that agent i is assigned item o. We denote by $\pi_i = \{o \in I : (i, o) \in \pi\}$ the set of items that an agent i receives in allocation π. We will assume in our paper that the allocation must be *exhaustive*, that is, all items must be assigned to some agent (and no two agents may receive the same item). Thus, $\bigcup_{i \in N} \pi_i = I$ and, for all distinct $i, j \in N$, $\pi_i \cap \pi_j = \emptyset$. As is customary, we write $[k]$ to denote $\{1, \ldots, k\}$. Finally, given a positive integer ℓ, we denote by $\ell\mathbb{N}$ the set of positive integers that are multiples of ℓ.

Utilities. Each agent $i \in N$ is associated with a (dis)utility function $u_i : I \to \mathbb{R}$, which indicates how much they like or dislike each item. Namely, we consider a setting where each agent may view each item as a good, a chore, or a null item. In particular, we say that an item $o \in I$ is an *objective good* (resp. *chore* or *null*) if, for all $i \in N$, $u_i(o) > 0$ (resp. $u_i(o) < 0$ or $u_i(o) = 0$). Otherwise, we say that o is a *subjective* item.

We focus on additive utility functions and with a slight abuse of notation we write $u_i(S) = \sum_{o \in S} u_i(o)$ for the utility that agent i gets from set $S \subseteq I$. Thus, the utility of agent i for an allocation π is given by $u_i(\pi_i)$. We denote by $\boldsymbol{u} = (u_1, \ldots, u_n)$ the profile of utilities for the agents.

Fairness. We introduce several fairness concepts, as defined by Aziz et al. [3] and by Amanatidis et al. [2].

A classical notion of fairness is that of *envy-freeness*: an allocation is envy-free if no agent finds that the bundle given to someone else is better than the one they received themselves. Formally, we have:

Definition 1 (EF)**.** *For agents N, items I, and profile $\mathbf{u}$, an allocation π is* envy-free (EF) *if for any $i, j \in N$, $u_i(\pi_i) \geqslant u_i(\pi_j)$.*

It is easy to see that this notion is too strong and cannot always be achieved (consider the case of one objective good and two agents). Thus, it has then been relaxed to *envy-freeness up to one item*, which has been further generalized to take into account both goods and chores.

Definition 2 (EF1)**.** *For agents N, items I, and profile $\mathbf{u}$, an allocation π is* envy-free up to one item (EF1) *if for any $i, j \in N$, either π is envy-free, or there is $o \in \pi_i \cup \pi_j$ such that $u_i(\pi_i \setminus \{o\}) \geqslant u_i(\pi_j \setminus \{o\})$.*

We also consider the following (weaker) version of EF1, where the envy of i towards j can be eliminated by either giving a good of j to i, or imposing a chore of i on j.

Definition 3 (Weak EF1)**.** *An allocation π is* weak EF1 *if for all $i, j \in N$, either $u_i(\pi_i) \geqslant u_i(\pi_j)$ or there is an item $o \in \pi_i \cup \pi_j$ such that $u_i(\pi_i \cup \{o\}) \geqslant u_i(\pi_j \setminus \{o\})$ or $u_i(\pi_i \setminus \{o\}) \geqslant u_i(\pi_j \cup \{o\})$.*

Yet another concept is that of *proportionality* of an allocation, meaning that each agent receives their due share of (dis)utility.

Definition 4 (PROP)**.** *For agents N, items I, and profile $\mathbf{u}$, an allocation π satisfies* proportionality *if for each agent $i \in N$ we have $u_i(\pi_i) \geqslant {}^{u_i(I)}/_n$.*

It is easy to see that envy-freeness implies proportionality when assuming additive utilities (see, e.g., Aziz et al. [3, Proposition 1]).

Efficiency. Alongside fairness, it is often desirable to distribute the items as efficiently as possible. One way to capture efficiency is the notion of *Pareto-optimality*, meaning that no improvement to the current allocation can be made without hurting some agent.

Definition 5 (PO)**.** *For agents N, items I, and profile $\mathbf{u}$, an allocation π is* Pareto-optimal (PO) *if there is no other allocation ρ such that for all $i \in N$, $u_i(\rho_i) \geqslant u_i(\pi_i)$ and for some $j \in N$ it holds that $u_j(\rho_j) > u_j(\pi_j)$. If such a ρ exists, we say that it* Pareto-dominates π.

Repeated setting. In our paper, we will be interested in *repeated* allocations of the items to the agents, i.e., sequences of allocations. We denote by $\bar{\pi}^{(k)} = (\pi^1, \pi^2, \ldots, \pi^k)$ the repeated allocation of the m items in I to the n agents in N over k time periods (or rounds). More formally, we will consider k copies of the set I, such that $I^1 = \{o_1^1, \ldots, o_m^1\}$, $\ldots, I^k = \{o_1^k, \ldots, o_m^k\}$. Then, each allocation π^ℓ corresponds to an allocation of the items in I^ℓ to the agents in N. For all agents $i \in N$, all items $o \in I$ and all $\ell \in [k]$, we let the utility be unchanged: i.e., $u_i(o) = u_i(o^\ell)$. When clear from context, we will drop the superscript ℓ from the items.

As a first approach, we will assess fairness over time by considering the global set of items that each agent has received across the k rounds. We denote by $\pi^{\cup k}$ the allocation of $k \cdot m$ (copies of the) items to the n agents, where $\pi_i^{\cup k} = \pi_i^1 \cup \cdots \cup \pi_i^k$ for each $i \in N$. Namely, we consider the allocation $\pi^{\cup k}$ where each agent gets the bundle of (the copies of) items that they have received across all the k time periods in $\bar{\pi}^{(k)}$. We say that a sequence of allocations $\bar{\pi}^{(k)}$ for some k satisfies an axiom *overall*, if $\pi^{\cup k}$ satisfies it. When clear from context, we will leave the term "overall" implicit. Similarly, we say that $\bar{\pi}^{(k)}$ Pareto-dominates $\bar{\rho}^{(k)}$ overall if $\pi^{\cup k}$ Pareto-dominates $\rho^{\cup k}$, and that $\bar{\pi}^{(k)}$ is Pareto-optimal overall if no $\bar{\rho}^{(k)}$ dominates it.

Since envy-freeness implies proportionality, we immediately get:

Proposition 1. *For additive utilities, if $\bar{\pi}^{(k)}$ is envy-free overall, then it is proportional overall.*

As a second approach, we will assess fairness over time by checking whether each repetition satisfies some desirable criteria. For an axiom A, we say that a sequence of allocations $\bar{\pi}^{(k)}$ (for some k) satisfies *per-round* A, if for every $j \in [k]$, the allocation π^j satisfies the axiom. For example, an allocation $\bar{\pi}^{(k)} = (\pi^1, \ldots, \pi^k)$ is *per-round* EF1 if every allocation π^j for $j \in [k]$ is EF1.

Proposition 2. *If $\bar{\pi}^{(k)}$ is Pareto-optimal overall, then it is per-round Pareto-optimal.*

The converse direction does not hold, namely, per-round Pareto-optimality may not imply overall Pareto-optimality, as the following example shows.

Example 1. Consider the following utility profile over two agents and two items:

	o_1	o_2
1	4	5
2	3	9

Now consider a four-round allocation where agent 1 gets both items in the first and second rounds, and agent 2 gets both items in the remaining two rounds. It is easy to verify that such an allocation sequence is per-round PO. However, it is not PO overall. Indeed, this gives utilities 18 for agent 1 and 24 for agent 2. Instead, an allocation sequence where agent 1 always gets o_1 and gets o_2 once (thus, agent 2 takes o_2 thrice and no item in one round) yields utilities 21 to agent 1 and 27 to agent 2. Observe that we could obtain a similar example where all items are chores by just multiplying all utilities by -1. △

On the other hand, for envy-freeness and proportionality, we get the following.

Proposition 3. *For additive utilities, if $\bar{\pi}^{(k)}$ is per-round envy-free (resp. proportional), then it is envy-free (resp. proportional) overall.*

It is easy to see that the converse does not necessarily hold. Indeed, consider a two-agent, two-round scenario with one objective good, where we assign the good to one agent in the first round, and to the other agent in the second round. This is envy-free (resp. proportional) overall, but not per-round.

3 The case of n agents

In this section, we study the possibility of finding fair and efficient sequences of allocations for the general case of any number of agents.

We start by looking at envy-freeness. First of all, observe that, whenever k is a multiple of n, then we can always guarantee an EF allocation.

Proposition 4. *If $k \in n\mathbb{N}$, there exists a sequence $\bar{\pi}^{(k)}$ which satisfies envy-freeness overall.*

Proof. Consider an arbitrary initial allocation π^1 of the items in I to the agents in N. Now, consider the allocations $\pi^2, \ldots, \pi^k$ such that for all $i \in N$ and $j \in \{1, \ldots, k-1\}$, we have $\pi_i^{j+1} = \pi_{i-1}^j$ and $\pi_1^{j+1} = \pi_n^j$. Observe that across the k iterations, each agent i has received each of the bundles composing π^1 exactly k/n times, thus $\pi_i^{\cup k}$ consists of k/n (copies of) all the items in I for any agent i. Thus, $\pi^{\cup k}$ is envy-free, as the bundles for each of the agents are equivalent, and therefore the sequence $\bar{\pi}^{(k)}$ is envy-free overall. □

Note that this is not always possible if k is not a multiple of n, irrespectively of m.

Proposition 5. *For every $n \geqslant 2$, every $k \in \mathbb{N} \setminus n\mathbb{N}$ and every $m \in \mathbb{N}$, an allocation that is proportional overall is not guaranteed to exist, even if the items are all goods or all chores.*

Proof. Consider the following profile of utilities, where every agent has utility 1 for all items, except for a special item o^*, for which all agents have utility $k(m-1)+1$. Consider any allocation sequence $\bar{\pi}^{(k)}$. Since k is not divisible by n, it is impossible to give o^* an equal amount of times to all agents (i.e., k/n times). Thus, there must be some agent i receiving o^* less times than k/n. Let $s_i < k/n$ be the number of times she receives o^*. Then,

$$\begin{aligned} u_i(\pi_i^{\cup k}) &\leqslant s_i(k(m-1)+1) + k(m-1) \\ &\leqslant (k/n - 1)(k(m-1)+1) + k(m-1) \\ &= k/n(k(m-1)+1) - 1 \\ &< k/n(k(m-1)+1) + k/n(m-1) \\ &= k/n \cdot u_i(I). \end{aligned}$$

Hence, i receives less than her proportional fair share. Observe that, by multypling all utilities by -1, we get an analogous counter-example where all items are chores. □

If we additionally require Pareto-optimality, for $n > 2$, we derive the following.

Theorem 6. *If $n > 2$, for every $k \in \mathbb{N}$, an allocation that is envy-free and Pareto-optimal overall is not guaranteed to exist, even if the items are all chores.*

Proof (Sketch). First, observe that, by Propositions 1 and 5, if k is not a multiple of n, we are done. Now let k be a multiple of n. Consider the following profile, where s and b stand for a *small chore* and a *big chore*, respectively:

	s	b
1	-1	$-k$
2	-1	$-k$
$\vdots$	$\vdots$	$\vdots$
$n-1$	-1	$-k$
n	-1	$-k/n$

We claim that in every sequence of allocations that is EF overall, all agents must receive s and b the same amount of times (namely k/n times). The proof of this claim can be found in Appendix A. Next, observe that any such allocation always has a Pareto improvement. Indeed, suppose that agent 1 gives one of her allocations of b to agent n and agent n gives back to 1 all of her allocations of s. Then, agent 1 will be happier (she increased her utility by $k - k/n$), and all other agents will be equally happy (for $i \in \{2, \ldots, n-1\}$ nothing changes, while n gains k/n but also loses k/n utility). However, agent 2, e.g., now envies 1. □

We do not have an analogous statement where all items are goods. However, in Appendix A we have a concrete example where no envy-free and PO allocation exists with $n = k = 3$ (Example 3). Furthermore, we note that it is not clear whether the following (weaker) statement holds: *Given a utility profile, is there a sequence of allocations $\bar{\pi}^{(k)}$ (of arbitrary length k) that is EF and PO overall? If so, how can we compute $\bar{\pi}^{(k)}$?* We leave this as an interesting open question.

In light of Theorem 6, envy-freeness seems too strong of a requirement even in the repeated setting. However, we can at least always find a proportional and PO sequence of allocations.

Theorem 7. *For every $n \geqslant 2$ and $k \in n\mathbb{N}$, an allocation sequence that is proportional and Pareto-optimal overall always exists, and can be computed in time $O(n^{mk} \cdot m)$.*

Proof. Consider any proportional sequence of allocations $\bar{\pi}^{(k)}$ (which must exist, by Propositions 1 and 4). Clearly, any Pareto-improvement over $\bar{\pi}^{(k)}$ must also be proportional. We get the following algorithm:

1. Initialize a variable r as $\bar{\pi}^{(k)}$, as defined previously.
2. Iterate over all possible k-round allocations $\bar{\rho}^{(k)}$, and do:
   - if $\bar{\rho}^{(k)}$ Pareto-dominates r, set r to $\bar{\rho}^{(k)}$.
3. Return r.

This algorithm runs in $O(n^{mk} \cdot m)$. Indeed, every iteration at Step 2 of the algorithm requires time $O(m)$, as we just need to sum up the utilities for the items received by each agent to compare the total utilities of r and $\bar{\rho}^{(k)}$. Furthermore, there are n^{mk} possible allocations, as for any round $j \in [k]$ and every object $o \in I$, we could assign object o in round j to n different agents. It remains to be shown that the algorithm is correct.

Let us call $\bar{\eta}^{(k)}$ the allocation in r returned by the algorithm. As argued above, $\bar{\eta}^{(k)}$ must be proportional overall, since it is either equal to $\bar{\pi}^{(k)}$ or a Pareto-improvement of it. Thus, suppose towards a contradiction that $\bar{\eta}^{(k)}$ is not PO. Then, there must be some Pareto-optimal $\bar{\rho}^{(k)}$ dominating $\bar{\eta}^{(k)}$ that is encountered before $\bar{\eta}^{(k)}$ in the iteration—since, if it is encountered after, the variable r would be updated to $\bar{\rho}^{(k)}$. But since Pareto-dominance is transitive, we know that $\bar{\rho}^{(k)}$ Pareto-dominates all the allocations corresponding to the values that the variable r took before being updated to $\bar{\eta}^{(k)}$, and hence r is updated to $\bar{\rho}^{(k)}$ during iteration. However, this means that r cannot be updated to $\bar{\eta}^{(k)}$, as $\bar{\rho}^{(k)}$ dominates it: a contradiction. This concludes the proof. □

Note that the existence guarantee in Theorem 7 is tight because of Proposition 5.

We now propose two ways to deal with the exponential-time runtime of Theorem 7. First, observe that we can define an integer linear program (ILP) to compute such a proportional and PO allocation (Figure 1). Furthermore, in practical applications, two relevant parameters are the number of agents n and the maximum value of an item $u_{max} = \max_{i \in N, o \in I} u_i(o)$. Invoking a result of Bredereck et al. [8] for one-shot fair division, we can get fixed-parameter tractability with respect to $n + u_{max}$ when each utility $u_i(o)$ is an integer. Here, a problem is said to be fixed-parameter tractable (FPT) with respect to a parameter τ if each instance H of this problem can be solved in time $f(\tau)poly(|H|)$ for some function f.

Theorem 8. *For every $n \geqslant 2$ and $k \in n\mathbb{N}$, finding a proportional and Pareto-optimal allocation sequence can be done in FPT time with respect to $n + m$, and when each utility $u_i(o)$ is an integer in FPT time with respect to $n + u_{max}$.*

$$
\begin{aligned}
\textbf{maximise} \quad & \sum_{i \in N} \sum_{o \in I} u_i(o) \cdot x_o^i & \\
\textbf{subject to}: \quad & x_o^i \in \{0, \dots, k\} & \text{for } o \in I, i \in N \\
& \sum_{i \in N} x_o^i = k & \text{for } o \in I \\
& \sum_{o \in I} u_i(o) \cdot x_o^i \geqslant {}^{k}/_{n} \cdot u_i(I) & \text{for } i \in N
\end{aligned}
$$

Figure 1: An ILP for finding a proportional and PO allocation sequence for n agents. In this ILP, we maximise the social welfare and thus guarantee PO. Variable x_o^i indicates how often agent i receives item o. If $k \notin n\mathbb{N}$, this ILP may be unsatisfiable due to Prop. 5.

Proof. Recall that such a proportional and PO allocation sequence is guaranteed to exist (Theorem 7). Fixed-parameter tractability in $n + m$ follows from the ILP in Figure 1, together with the fact that solving an ILP is FPT in the number of variables (nm) [19]. Next, Bredereck et al. [8] showed that the problem of computing a PO *maxmin allocation* π that maximizes the minimum utility $\min_{i \in N} u_i(\pi_i)$ can be done in FPT time with respect to $n+u_{max}$ when each utility $u_i(o)$ is an integer. Thus, one can compute an allocation $\pi^{\cup k}$ that maximizes the minimum utility $\min_{i \in N} u_i(\pi_i^{\cup k})$ in FPT time with respect to $n + u_{max}$. □

Observe that Theorem 6 leaves open whether, for $n = 2$, an envy-free and PO allocation is always guaranteed to exist. We will tackle this question in the next section.

4 The case of two agents

In this section, we consider a special but important case of the repeated fair division among two agents. A prime example is a house chore division among couples, where the two members repeatedly decide the division of house labor [14].

One first intuitive idea to achieve this would be the following. For two agents and additive utilities, Aziz et al. [3] introduced an efficient algorithm (*Generalized Adjusted Winner*, or GAW) which finds a PO and EF1 allocation for the one-shot fair allocation problem of goods and chores. Essentially, the procedure first assigns all the goods to one of the agents (the *winner*) and all the chores to the other (the *loser*). Then, it moves all items one by one from one agent to the other until EF1 is achieved.

Thus, one could consider using the GAW algorithm twice, by choosing one of the agents as the winner of the first round, and then swapping their roles for the second round: this would yield us efficiently a sequence $\overline{\pi}^{(2)}$ where each allocation is PO and EF1 (as we used the GAW algorithm each time), and hopefully the overall allocation $\pi^{\cup k}$ is PO and EF. However, this approach may fail; see Example 2 in Appendix A.

Despite this, we will show that, whenever the number of rounds is even, we can still always find an allocation that is PO and EF overall. However, we lose the guarantee of an efficient computation. First, we need the following fact:

Proposition 9. *If $n = 2$, for additive utilities, proportionality implies envy-freeness.*

Now, we can show the following.

Theorem 10. *If $n = 2$ and $k \in 2\mathbb{N}$, an allocation sequence that is envy-free and Pareto-optimal overall always exists, and can be computed in time $O(m \cdot (k + 1)^m)$.*

Proof. Since proportionality implies envy-freeness when $n = 2$ (Proposition 9), here the algorithm in the proof of Theorem 7 returns an envy-free and Pareto-optimal allocation. Furthermore, if $n = 2$, there are $(k+1)^m$ possible allocations (up to symmetry-breaking), as each agent receives each of the m items either 0, 1, ..., or k times. Thus, we obtain a runtime of $O(m \cdot (k+1)^m)$. □

We can thus use the ILP in Figure 1 to compute a PO and envy-free allocation sequence. Furthermore, as a straightforward consequence of Theorem 8 and Proposition 9, we get the following.

Proposition 11. *If $n = 2$ and $k \in 2\mathbb{N}$, an allocation which is Pareto-optimal and envy-free can be computed in FPT time with respect to m, and when each utility $u_i(o)$ is an integer in FPT time with respect to u_{max}.*

Thus, in this scenario we can guarantee an envy-free and PO overall allocation. What about the individual rounds? First, observe that an envy-free and PO allocation that is per-round EF1 may not exist, even if all items are goods or chores.

Proposition 12. *If $n = 2$ and $k > 2$, an allocation that is per-round EF1, Pareto-optimal overall, and envy-free is not guaranteed to exist, even when all items are goods or chores.*

Proof. First, observe that the number of rounds k needs to be even, otherwise we know that EF cannot always be achieved (Proposition 5). Next, consider the following utility profile over two items $I = \{o_1, o_2\}$, where $u_1(o_1) = u_2(o_1) = 1$, $u_1(o_2) = 3$ and $u_2(o_2) = 2$.

For $k > 2$, there is no EF and PO allocation sequence which is also per-round EF1. Indeed, in order to be per-round EF1, both agents should receive at least one item in each round: otherwise, the agent receiving no items would envy the other agent for more than one item. Hence, in each round, each agent will receive either o_1 or o_2 (but not both). It is easy to see that, in any sequence of allocations where one of the agents receives o_1 more than $k/2$ times, they will envy the other agent. Hence, we can discard them and only focus on the case where agent 1 (resp. agent 2) receives both items $k/2$ times.

Call this sequence $\bar{\pi}^{(k)}$. We show that $\bar{\pi}^{(k)}$ is not PO. In fact, it is dominated by the sequence where agent 1 gets only o_1 exactly $k/2 - 2$ times, only o_2 exactly $k/2 + 1$ times, and no items once. Indeed, the former sequence $\bar{\pi}^{(k)}$ gives satisfaction $2k$ and $3/2k$ to agent 1 and agent 2 (respectively), while the latter gives satisfaction $2k + 1$ and $3/2k$. Hence, there is no EF and PO allocation that is per-round EF1.

Finally, observe that we can make an analogous counter-example where all items are chores by multiplying all the utilities by -1. □

Luckily, if $k = n = 2$, then this is always possible. To prove this, we show that, in this case, it is always possible to transform an allocation that is PO and EF overall to an allocation that is per-round EF1 while preserving the PO and EF guarantees.

Theorem 13. *Suppose $k = n = 2$. Given an allocation sequence that is Pareto-optimal and envy-free overall, an allocation sequence which is Pareto-optimal, envy-free, and per-round EF1 can be computed in polynomial time.*

Proof. Consider some overall PO and envy-free allocation sequence $\bar{\rho}^{(2)}$ for two agents. Let $I_1 = \rho_1^1 \cap \rho_1^2$ (and similarly for I_2) be the items assigned to agent 1 (resp. 2) in both rounds. Moreover, let $A = I \setminus (I_1 \cup I_2)$ be the items that each agent receives once. Observe that, by PO, we can assume that every subjective item o is always assigned to the agent $i \in \{1, 2\}$ for whom $u_i(o) > u_j(o)$ (where j is the other agent). Furthermore, we can remove objective null items from consideration, as they can be assigned to any agent in any round without changing the outcome.

Hence, all items in A are objective goods or chores. Let $A^+ \subseteq A$ be the objective goods in A, and let $A^- \subseteq A$ be the objective chores in A. Clearly, $A^+ \cap A^- = \emptyset$.

We refine this allocation sequence via the following procedure:

1. Start with an allocation $\bar{\pi}^{(2)}$ where $\pi_1^1 = I_1 \cup A^-$, $\pi_2^1 = I_2 \cup A^+$, $\pi_1^2 = I_1 \cup A^+$ and $\pi_2^2 = I_2 \cup A^-$.

2. For each o in A, do:

   (a) If $\bar{\pi}^{(2)}$ is per-round EF1, break the loop.

   (b) Else: if o is an objective chore, remove o from π_1^1 and π_2^2, and add it to π_2^1 and π_1^2; if o is an objective good, remove o from π_2^1 and π_1^2, and add it to π_1^1 and π_2^2.

3. Return $\bar{\pi}^{(2)}$.

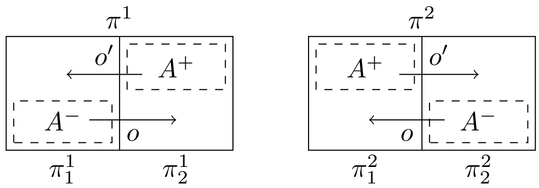

Figure 2: Exchanges of goods and chores between agents in step 2(b) of the algorithm.

In other words, we start by assigning all chores in A to agent 1 in the first round (and to 2 in the second round), and conversely for the goods. Then, we progressively swap the chores from 1 to 2 in the first round, and vice-versa in the second round, and conversely for the goods. We stop when the overall allocation is per-round EF1. This algorithm clearly runs in polynomial time: it remains to show that it is correct.

We begin by showing the following fact, which will be useful throughout the proof:

$$2u_1(I_1) + u_1(A) \geqslant 2u_1(I_2) + u_1(A) \implies \qquad \text{(by envy-freeness)}$$
$$u_1(I_1) \geqslant u_1(I_2)$$

Similarly, we can derive $u_2(I_2) \geqslant u_2(I_1)$. Next, observe that this algorithm returns a PO and EF allocation, since we never change the total amount of times an agent receives an item w.r.t. the initial allocation.

Consider now the initial allocation $\bar{\pi}^{(2)}$, before executing the loop in Step 2. If this is per-round EF1, we are done. Thus, suppose agent 1 envies agent 2 in some round (the case were 2 envies 1 is analogous): then we must have that 1 envies 2 only in the first round. Indeed, suppose towards a contradiction that 1 envies 2 in the second round, i.e., $u_1(I_1 \cup A^+) < u_1(I_2 \cup A^-)$. We get:

$$u_1(I_1 \cup A^-) < u_1(I_1 \cup A^+) < u_1(I_2 \cup A^-) < u_1(I_2 \cup A^+)$$

But $u_1(I_1 \cup A^-) < u_1(I_2 \cup A^+)$ means that agent 1 envies 2 in the first round as well. Furthermore, since the overall allocation is EF, agent 1 can only envy 2 in one (i.e., the first) round; hence, we obtain $u_1(I_1 \cup A^+) \geqslant u_1(I_2 \cup A^-)$.

Thus, in the first round, agent 1 initially envies 2, but by the end of the iteration in Step 2, agent 1 cannot envy 2 anymore, since $u_1(I_1 \cup A^+) \geqslant u_1(I_2 \cup A^-)$. Therefore, there must be an item $o^\star$ such that, after transferring $o^\star$, agent 1 does not envy 2 in the first round anymore. In the following, we assume that $o^\star$ is a chore (the case where $o^\star$ is a good is analogous). That is, there must be two disjoint sets $L, R \subseteq A \setminus \{o^\star\}$ (with $L \cup R \cup \{o^\star\} = A$) and an allocation $\bar{\eta}^{(2)}$ such that:

- In the first round η^1, agent 1 gets $I_1 \cup L \cup \{o^\star\}$ and agent 2 gets $I_2 \cup R$.
- In the second round η^2, agent 1 gets $I_1 \cup R$ and agent 2 gets $I_2 \cup L \cup \{o^\star\}$.
- $u_1(I_1 \cup L \cup \{o^\star\}) < u_1(I_2 \cup R)$ and $u_1(I_1 \cup L) \geqslant u_1(I_2 \cup R \cup \{o^\star\})$.

Crucially, observe that we can suppose w.l.o.g. that after transferring $o^\star$, agent 2 does not envy 1 in the second round anymore, i.e., $u_2(I_2 \cup L) \geqslant u_2(I_1 \cup R \cup \{o^\star\})$. Otherwise, i.e., if after transferring $o^\star$ agent 2 envied agent 1 in the second round, by the same argument as above, we could carry on the iteration until we reach some item $o^\dagger \neq o^\star$ such that, after transferring $o^\dagger$, agent 2 does not envy 1 in the second round anymore. At that point, we would have that agent 1 does not envy 2 in the first round either (by the above assumption about $o^\star$), and thus we could repeat our analysis by taking agent 2 instead of 1 as the initial envious agent, $o^\dagger$ instead of $o^\star$, and obtain a completely symmetrical treatment.

With the above settled, there are four cases to consider:

- **Case 1**: $u_1(I_1 \cup L) \geqslant u_1(I_2 \cup R)$ and $u_2(I_2 \cup L) \geqslant u_2(I_1 \cup R)$. In this case, $\bar{\eta}^{(2)}$ is per-round EF1. Indeed, agent 1 is envious in the first round, so by EF she cannot be envious in the second round. By removing $o^\star$ from her allocation in the first round, we eliminate envy for agent 1, since $u_1(I_1 \cup L) \geqslant u_1(I_2 \cup R)$. Furthermore, agent 2 cannot be envious in the first round (otherwise, since agent 1 already is envious, we could swap their allocations in the first round and increase both agents' utility, contradicting PO). Finally, if agent 2 is envious in the second round, we can eliminate envy by removing $o^\star$ from her allocation, as $u_2(I_2 \cup L) \geqslant u_2(I_1 \cup R)$.

For the remaining cases, let us update $\bar{\eta}^{(2)}$ as follows: assign $o^\star$ to agent 2 in the first round, and to agent 1 in the second round (i.e., we transfer $o^\star$). We show that, in the three remaining cases, this updated sequence is per-round EF1.

First, observe that, since by assumption we have $u_1(I_1 \cup L) \geqslant u_1(I_2 \cup R \cup \{o^\star\})$ and $u_2(I_2 \cup L) \geqslant u_2(I_1 \cup R \cup \{o^\star\})$, agent 1 does not envy 2 in the first round (and agent 2 does not envy 1 in the second round). Next:

- **Case 2**: $u_1(I_1 \cup L) < u_1(I_2 \cup R)$ and $u_2(I_2 \cup L) < u_2(I_1 \cup R)$. This implies that $u_1(I_2 \cup L) < u_1(I_1 \cup R)$ and $u_2(I_1 \cup L) < u_2(I_2 \cup R)$. Here, $\bar{\eta}^{(2)}$ is per-round EF1: if we remove $o^\star$ from the allocation of 1 (resp. 2) in the second round (resp. first) we eliminate envy, since $u_1(I_2 \cup L) < u_1(I_1 \cup R)$ and $u_2(I_1 \cup L) < u_2(I_2 \cup R)$.
- **Case 3**: $u_1(I_1 \cup L) < u_1(I_2 \cup R)$ and $u_2(I_2 \cup L) \geqslant u_2(I_1 \cup R)$. Again, this implies $u_1(I_2 \cup L) < u_1(I_1 \cup R)$. Therefore, if we remove $o^\star$ from the bundle of 1 in the second round, she does not envy 2. It remains to show that agent 2 does is not envious in the first round once we remove $o^\star$ from her bundle. Towards a contradiction, assume that $u_2(I_2 \cup R) < u_2(I_1 \cup L)$. Now, suppose we change the first round as follows:
  - Agent 1 gets assigned $I_2 \cup R$ instead of $I_1 \cup L$ and
  - Agent 2 gets assigned $I_1 \cup L \cup \{o^\star\}$ instead of $I_2 \cup R \cup \{o^\star\}$.

  Observe that, since $u_1(I_1 \cup L) < u_1(I_2 \cup R)$ and $u_2(I_2 \cup R) < u_2(I_1 \cup L)$, this is a Pareto-improvement over the original allocation: contradiction. Thus, $u_2(I_2 \cup R) \geqslant u_2(I_1 \cup L)$, and agent 2 is not envious in the first round if we remove $o^\star$ from her allocation.
- **Case 4**: $u_1(I_1 \cup L) \geqslant u_1(I_2 \cup R)$ and $u_2(I_2 \cup L) < u_2(I_1 \cup R)$. Here, the argument is symmetric to the previous case.

In all cases, we find a per-round EF1 allocation. This concludes the proof. □

Combining the above with Theorem 10 and Proposition 11, we get the following.

Corollary 14. *If $k = n = 2$, then an allocation which is Pareto-optimal and envy-free overall and per-round EF1 always exists. Moreover, it can be computed in time $O(m \cdot 3^m)$, and when each utility $u_i(o)$ is an integer in FPT time with respect to u_{max}.*

As we have seen in Proposition 12, we cannot guarantee per-round EF1 together with PO and EF overall for a more general number of rounds $k \in 2\mathbb{N}$ with $k > 2$. Nevertheless, these properties become compatible if we relax a per-round fairness requirement to weak EF1. Again, to do this, we show that it is always possible to transform an allocation that is PO and EF overall to an allocation that is per-round weak EF1 while preserving the PO and EF guarantees.

Theorem 15. *Suppose $n = 2$. Given a k-round allocation that is Pareto-optimal and envy-free, an allocation which is Pareto-optimal, envy-free, and per-round weak EF1 can be computed in polynomial time.*

Corollary 16. *If $n = 2$ and $k \in 2\mathbb{N}$, then an allocation which is Pareto-optimal, envy-free, and per-round weak EF1 always exists. Moreover, it can be computed in time $O(m \cdot (k+1)^m)$, in FPT time with respect to m, and when each utility $u_i(o)$ is an integer in FPT time with respect to u_{max}.*

Note that in order to obtain Corollary 16, it suffices to prove Theorem 15, since we can apply Theorem 10 and Proposition 11 to obtain an allocation $\bar{\pi}^{(k)} = (\pi^1, \pi^2, \ldots, \pi^k)$ that is PO and envy-free overall when $k \in 2\mathbb{N}$.

Now, suppose that we are given a Pareto-optimal and envy-free k-round allocation. Looking into each individual round, there may be an allocation π^i that is not fair. Similarly to the proof of Theorem 13, we thus repeatedly transfer an item between *envious* rounds and *envy-free* rounds while preserving the property that the overall allocation $\bar{\pi}^{(k)}$ is PO and EF over the course of the algorithm. We can show that this process terminates in polynomial time and eventually yields a per-round weak EF1 allocation. A formal description of our algorithm is presented in Algorithm 1 in Appendix A.

Finally, we show that, if we do not require PO, an allocation that is EF and per-round EF1 can always be found (when k is even) in polynomial time.

Theorem 17. *If $n = 2$ and $k \in 2\mathbb{N}$, then an allocation which is envy-free overall and per-round EF1 always exists, and can be computed in polynomial time.*

We do not know whether EF and per-round EF1 can be simultaneously achieved for $n > 2$ agents and $k \in n\mathbb{N}$ rounds (EF overall is possible in this case by Proposition 4). We leave it as an interesting open problem for future work.

5 Conclusion

We have seen that in our model of repeated allocations the (necessary) trade-off between fairness and efficiency is much more favorable than in the standard setting without repetitions. In the case of two agents, we presented an algorithm guaranteeing envy-freeness and Pareto-optimality for an even number of rounds. As this algorithm requires exponential time, it would be of interest to study the computational complexity of related decision problems, investigating whether and where polynomial-time results are obtainable.

The n-agent algorithm yields slightly weaker guarantees (proportionality and Pareto efficiency), which are still an improvement over the setting without repetition. It remains for future work to determine whether this result can be strengthened by additional per-round guarantees, as we have in the 2-agent case.

Acknowledgments

This paper developed from ideas discussed at the Dagstuhl Seminar 22271 *Algorithms for Participatory Democracy*. Martin Lackner was supported by the Austrian Science Fund (FWF), research grant P31890. Oliviero Nardi was supported by the European Union's Horizon 2020 research and innovation programme, grant 101034440.

References

[1] Martin Aleksandrov and Toby Walsh. Online fair division: A survey. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-2020)*, pages 13557–13562, 2020.

[2] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. Fair division of indivisible goods: A survey. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI-2022)*, pages 5385–5393, 2022.

[3] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi-Agent Systems*, 36(1):1–21, 2022.

[4] Gabriel Balan, Dana Richards, and Sean Luke. Long-term fairness with bounded worst-case losses. *Autonomous Agents and Multi-Agent Systems*, 22(1):43–63, 2011.

[5] Gerdus Benade, Aleksandr M Kazachkov, Ariel D. Procaccia, and Christos-Alexandros Psomas. How to make envy vanish over time. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC-2018)*, pages 593–610, 2018.

[6] Niclas Boehmer and Rolf Niedermeier. Broadening the research agenda for computational social choice: Multiple preference profiles and multiple solutions. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2021)*, 2021.

[7] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[8] Robert Bredereck, Andrzej Kaczmarczyk, Dšan Knop, and Rolf Niedermeier. High-multiplicity fair allocation: Lenstra empowered by N-fold integer programming. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC-2019)*, pages 505–523, 2019.

[9] Ruggiero Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 220–229, 2008.

[10] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC-2017)*, pages 629–646, 2017.

[11] Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. Fair social choice in dynamic settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-2017)*, pages 4580–4587, 2017.

[12] Rupert Freeman, Seyed Majid Zahedi, Vincent Conitzer, and Benjamin C. Lee. Dynamic proportional sharing: A game-theoretic approach. In *Proceedings of the ACM on Measurement and Analysis of Computing Systems (MACS-2018)*, 2018.

[13] Mingyu Guo, Vincent Conitzer, and Daniel M. Reeves. Competitive repeated allocation without payments. In *International Workshop on Internet and Network Economics*, pages 244–255, 2009.

[14] Ayumi Igarashi and Tomohiko Yokoyama. Kajibuntan: A house chore division app. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, Forthcoming.

[15] Ian Kash, Ariel D. Procaccia, and Nisarg Shah. No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research*, 51:579–603, 2014.

[16] Thomas Kohler, Jan-Philipp Steghöfer, Dídac Busquets, and Jeremy Pitt. The value of fairness: Trade-offs in repeated dynamic resource allocation. In *2014 IEEE Eighth International Conference on Self-Adaptive and Self-Organizing Systems (SASO-2014)*, pages 1–10, 2014.

[17] Martin Lackner. Perpetual voting: Fairness in long-term decision making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-2020)*, pages 2103–2110, 2020.

[18] Martin Lackner, Jan Maly, and Simon Rey. Fairness in long-term participatory budgeting. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI-2021)*, pages 299–305, 2021.

[19] Hendrik W. Lenstra Jr. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983.

[20] Reshef Meir, Jérôme Lang, Julien Lesca, Nicholas Mattei, and Natan Kaminsky. A market-inspired bidding scheme for peer review paper assignment. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI-2021)*, pages 4776–4784, 2021.

[21] Abraham Othman, Tuomas Sandholm, and Eric Budish. Finding approximate competitive equilibria: efficient and fair course allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2010)*, pages 873–880, 2010.

[22] Justin Payan and Yair Zick. I will have order! Optimizing orders for fair reviewer assignment. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI-2022)*, pages 440–446, 2022.

[23] David Zeng and Alexandros Psomas. Fairness-efficiency tradeoffs in dynamic fair division. In *Proceedings of the 2020 ACM Conference on Economics and Computation (EC-2020)*, pages 911–912, 2020.

Ayumi Igarashi
The University of Tokyo
Tokyo, Japan
Email: `igarashi@mist.i.u-tokyo.ac.jp`

Martin Lackner
DBAI, TU Wien
Vienna, Austria
Email: lackner@dbai.tuwien.ac.at

Oliviero Nardi
DBAI, TU Wien
Vienna, Austria
Email: oliviero.nardi@tuwien.ac.at

Arianna Novaro
CES, University of Paris 1 (Panthéon-Sorbonne)
Paris, France
Email: arianna.novaro@univ-paris1.fr

A Omitted Examples and Proofs

Example 2. Consider the following profile for two agents $N = \{1, 2\}$ expressing their utilities over three goods $I = \{o_1, o_2, o_3\}$:

	o_1	o_2	o_3
1	4.5	3	7
2	9	5	10

Consider now a two-round sequence $\bar{\pi}^{(2)} = (\pi^1, \pi^2)$ where we apply the GAW algorithm by choosing agent 1 as the loser in the first round and agent 2 as the loser in the second round. We then get allocation π^1 with $\pi^1_1 = \{o_3\}$ and $\pi^1_2 = \{o_1, o_2\}$, and an allocation π^2 with $\pi^2_1 = \{o_2, o_3\}$ and $\pi^2_2 = \{o_1\}$. Both π_1 and π_2 are, by themselves, EF1 and PO (as we used GAW each time). However, $\bar{\pi}^{(2)}$ is not EF, since 2 envies 1. By multiplying all the utilities by -1, we obtain an analogous example for the case where all items are chores. △

Theorem 6. *If $n > 2$, for every $k \in \mathbb{N}$, an allocation that is envy-free and Pareto-optimal overall is not guaranteed to exist, even if the items are all chores.*

Proof. First, observe that, by Propositions 1 and 5, if k is not a multiple of n, we are done. Now let k be a multiple of n. Consider the following profile, where s and b stand for a *small chore* and a *big chore*, respectively:

	s	b
1	-1	$-k$
2	-1	$-k$
$\vdots$	$\vdots$	$\vdots$
$n-1$	-1	$-k$
n	-1	$-k/n$

Let $p = k/n$. We claim that in every sequence of allocations that is EF overall, all agents must receive s and b the same amount of times (namely p times). Before proving this claim, observe that any such allocation always has a Pareto improvement. Indeed, suppose that agent 1 gives one of her allocations of b to agent n and agent n gives back to 1 all of her allocations of s. Then, agent 1 will be happier (she increased her utility by $k - k/n$), and all other agents will be equally happy (for $i \in \{2, \dots, n-1\}$ nothing changes, while n gains k/n but also loses k/n utility). However, agent 2, e.g., now envies 1.

Let us now prove the claim that all agents receive both items p times. Let s_i and b_i be the number of times where agent i receive items s and b, respectively. We will first prove that, for every $i, j \in [n-1]$, $s_i = s_j$ and $b_i = b_j$ must hold. Consider agents 1 and 2, and suppose towards a contradiction that $b_1 < b_2$. To guarantee envy-freeness, we must have $-s_1 - kb_1 = -s_2 - kb_2$, which means that $s_1 - s_2 = k(b_2 - b_1)$. Since $s_1, s_2, b_1, b_2 \in \{0, \dots, k\}$ and $b_1 < b_2$, this holds if and only if $s_1 = k$ and $s_2 = 0$ and $b_2 = b_1 + 1$. Furthermore, note that $s_1 = k$ implies that for every other agent $j \neq 1$, $s_j = 0$. Thus, to guarantee envy-freeness, we must have that, for all agents $i \in \{3, \dots, n-1\}$, $b_i = b_2 = b_1 + 1$ holds as well (as $s_i = 0 = s_2$).

Now, for agent 2 not to envy agent n, we must have $-b_2 k \geqslant -b_n k$ and thus $b_2 \leqslant b_n$. For agent n not to envy agent 2, we obtain similarly that $b_2 \geqslant b_n$. Thus $b_1 + 1 = b_2 = \cdots = b_n$, which implies:

$$\begin{aligned}
&\sum_{i\in[n]} b_i = k,\\
\implies &b_1 + (n-1)(b_1+1) = k,\\
\implies &nb_1 + (n-1) = k,\\
\implies &p = b_1 + \frac{(n-1)}{n}.
\end{aligned}$$

Since $p, n, b_1 \in \mathbb{N}$ and $n > 2$, we finally have our contradiction: thus, $b_1 \geqslant b_2$. By a completely symmetrical argument we get that $b_1 = b_2$ must hold (which implies $s_1 = s_2$ by envy-freeness). Furthermore, we can repeat this argument for every pair of agents i and j with $i, j \in [n-1]$. Thus, every such pair of agents must have $b_i = b_j$ and $s_i = s_j$.

We are ready to prove that all agents must receive both items the same amount of times, namely p times. First, for agent 1 not to envy agent n, we must have:

$$\begin{aligned}
&-s_1 - kb_1 \geqslant -s_n - kb_n,\\
\implies &-s_1 - kb_1 \geqslant -(k-(n-1)s_1) - k(k-(n-1)b_1),\\
\implies &s_1 - \frac{k}{n} \leqslant k(\frac{k}{n} - b_1),\\
\implies &s_1 - p \leqslant np(p - b_1).
\end{aligned}$$

With a similar line of reasoning, we can show that, for agent n not to envy agent 1, we must have that $s_1 - p \geqslant p(p-b_1)$. Additionally, we know that $(n-1)s_1 \leqslant k$ (as we cannot assign s to the agents in $[n-1]$ more than k times), which gives $s_1 \leqslant p\frac{n}{n-1}$. We can now show the following (given $p \geqslant 1$):

$$\begin{aligned}
&s_1 - p \geqslant p(p-b_1),\\
\iff &(p-b_1) \leqslant \frac{s_1}{p} - 1 \leqslant \frac{n}{n-1} - 1 = \frac{1}{n-1} \leqslant \frac{1}{2}.
\end{aligned}$$

Since $p - b_1 \in \mathbb{Z}$, $p > b_1$ would imply $1 \leqslant p - b_1 \leqslant \frac{1}{2}$, a contradiction: hence, $p \leqslant b_1$. By the above, we also know that $p(p-b_1) \leqslant s_1 - p \leqslant np(p-b_1)$, and thus:

$$\begin{aligned}
&p(p-b_1) \leqslant np(p-b_1),\\
\implies &(p-b_1) \leqslant n(p-b_1),
\end{aligned}$$

which implies $p \geqslant b_1$. Hence, $b_1 = p = {}^{k}/_{n}$. Finally:

$$p(p-b_1) = 0 \leqslant s_1 - p \leqslant np(p-b_1) = 0.$$

This gives $s_1 = b_1 = p = {}^{k}/_{n}$, which in turn yields $s_1 = b_1 = \cdots = s_n = b_n = {}^{k}/_{n}$. This concludes the proof. □

Example 3. Consider the following profile for three agents $N = \{1, 2, 3\}$ expressing their utilities over two goods $I = \{o_1, o_2\}$:

	o_1	o_2
1	1	2
2	1	2
3	1	1

Assume $k = 3$. One can verify that, in order to be achieve envy-freeness, all agents must receive both goods exactly once. This gives total utilities 3, 3 and 2 for agents 1, 2 and 3, respectively. However, any such allocation sequence is dominated by a sequence where agent 1 gets o_2 twice, agent 2 gets both items once, and agent 2 receives o_1 twice. Indeed, this gives total utilities 4, 3 and 2 for agents 1, 2 and 3, respectively. However, here agent 2 envies agent 1. Thus, no envy-free and PO allocation exists for $k = 3$. △

Proposition 9. *If $n = 2$, for additive utilities, proportionality implies envy-freeness.*

Proof. Let $N = \{1, 2\}$ and consider some proportional allocation $\bar{\pi}^{(k)}$. Then we get:

$$\begin{aligned} u_1(\pi_1^{\cup k}) \geqslant {}^{k}\!/_{2} \cdot u_1(I) \implies \\ k \cdot u_1(I) - u_1(\pi_2^{\cup k}) \geqslant {}^{k}\!/_{2} \cdot u_1(I) \implies \\ u_1(\pi_2^{\cup k}) \leqslant {}^{k}\!/_{2} \cdot u_1(I) \leqslant u_1(\pi_1^{\cup k}) \end{aligned}$$

Thus, 1 does not envy 2. A symmetrical argument shows that 2 does not envy 1. Hence $\bar{\pi}^{(k)}$ is envy-free overall. □

Theorem 15. *Suppose $n = 2$. Given a k-round allocation that is Pareto-optimal and envy-free, an allocation which is Pareto-optimal, envy-free, and per-round weak EF1 can be computed in polynomial time.*

We prove that given an allocation that is PO and EF overall, Algorithm 1 transforms it into a per-round weak EF1 allocation while keeping the PO and EF properties. Specifically, our adjustment procedure is divided into two phases: one that makes $\bar{\pi}^{(k)}$ weak EF1 for agent 1 (Lines 3 – 13), and another that makes $\bar{\pi}^{(k)}$ weak EF1 for agent 2 (Lines 15 – 25). In the first phase, it identifies π^j and π^i where agent 1 envies the other at π^j but she does not envy at π^i (Line 4 and Line 6). It then finds either a good o in $\pi_1^i \setminus \pi_1^j$ or a chore in $\pi_1^j \setminus \pi_1^i$ (Line 7) and transfers the item between π_1^i and π_1^j without changing the overall allocation. That is, the algorithm transfers item o from π_1^i to π_2^i and from π_2^j to π_1^j if it is a good (Line 9); it transfers item o from π_2^i to π_1^i and from π_1^j to π_2^j if it is a chore (Line 11). In the second phase, we swap the roles of the two agents and apply the same procedures.

The following lemmas ensure that a beneficial transfer is possible between the envious round π^j and the envy-free round π^i while keeping the PO and EF property.

Lemma 18. *Let $\bar{\pi}^{(k)}$ be a k-round allocation. Take any pair of distinct rounds $i, j \in [k]$. Suppose that agent 1 envies agent 2 at π^j but does not envy agent 2 at π^i. Then, there exits an item o such that*

- *$o \in \pi_1^i \setminus \pi_1^j$ and $u_1(o) > 0$, or*
- *$o \in \pi_1^j \setminus \pi_1^i$ and $u_1(o) < 0$.*

Proof. Since agent 1 envies agent 2 in the j-th round and does not envy in the i-th round, we have $u_1(\pi_1^j) < u_1(\pi_2^j)$ and $u_1(\pi_1^i) \geqslant u_1(\pi_2^i)$. Thus, $u_1(\pi_1^i) > \frac{u_1(I)}{2} > u_1(\pi_1^j)$, which means that there exits an item $o \in \pi_1^i \setminus \pi_1^j$ with $u_1(o) > 0$, or there exists an item $o \in \pi_1^j \setminus \pi_1^i$ with $u_1(o) < 0$. □

Lemma 19. *Suppose $n = 2$ and $N = \{1, 2\}$. Let $\bar{\pi}^{(k)}$ be a k-round allocation that is PO and EF. Take any pair of distinct rounds $i, j \in [k]$. Consider the following k-round allocation $(\hat{\pi}^1, \ldots, \hat{\pi}^k)$ where $\hat{\pi}_t = \pi_t$ for $t \neq i, j$ and there exists $o \in \hat{\pi}_1^i$ such that*

- *$\hat{\pi}_1^i = \pi_1^i \setminus \{o\}$ and $\hat{\pi}_2^i = \pi_2^i \cup \{o\}$, and*

- $\hat{\pi}_1^j = \pi_1^j \cup \{o\}$, *and* $\hat{\pi}_2^j = \pi_2^j \setminus \{o\}$.

Then, $(\hat{\pi}^1, \ldots, \hat{\pi}^k)$ *is both PO and EF overall.*

Proof. The claim clearly holds since the two k-round allocations $\bar{\pi}^{(k)}$ and $(\hat{\pi}^1, \ldots, \hat{\pi}^k)$ assign the same set of items to each agent. □

Algorithm 1: Algorithm for computing a PO and EF allocation that is per-round weak EF1

1 **Input**: A PO and EF allocation $\bar{\pi}^{(k)} = (\pi^1, \ldots, \pi^k)$;
2 Let E_1 denote the set of indices i such that π^i is not envy-free for agent 1. Let $F_1 = [k] \setminus E_1$.;
`// Envy-adjustment phase for agent 1`
3 **while** *there exists* $j \in E_1$ *such that* π^j *is not weak EF1 for agent* 1 **do**
4   Take such $j \in E_1$;
5   **while** π^j *is not weak EF1 for agent* 1 **do**
6   Let i be a round with minimum index $i \in F_1$;
7   Take item o such that $o \in \pi_1^i \setminus \pi_1^j$ with $u_1(o) > 0$, or $o \in \pi_1^j \setminus \pi_1^i$ with $u_1(o) < 0$ (the existence of o is guaranteed by Lemma 18);
8   **if** $u_1(o) > 0$ **then**
9   Set $\pi_1^i = \pi_1^i \setminus \{o\}$, $\pi_2^i = \pi_2^i \cup \{o\}$, $\pi_1^j = \pi_1^j \cup \{o\}$, and $\pi_2^j = \pi_2^j \setminus \{o\}$;
10   **else**
11   Set $\pi_1^j = \pi_1^j \setminus \{o\}$, $\pi_2^j = \pi_2^j \cup \{o\}$, $\pi_1^i = \pi_1^i \cup \{o\}$, and $\pi_2^i = \pi_2^i \setminus \{o\}$;
12   **if** π^i *is not envy-free for agent* 1 **then**
13   Set $F_1 = F_1 \setminus \{i\}$ and $E_1 = E_1 \cup \{i\}$;
`// Envy-adjustment phase for agent 2`
14 Let E_2 denote the set of indices i such that π^i is not envy-free for agent 2. (Note that $E_2 \subseteq F_1$ at this point.) Let $F_2 = [k] \setminus E_2$.;
15 **while** *there exists* $j \in E_2$ *such that* π^j *is not weak EF1 for agent* 2 **do**
16   Take such $j \in E_2$;
17   **while** π^j *is not weak EF1 for agent* 2 **do**
18   Let i be a round with minimum index $i \in F_2$;
19   Take item o such that $o \in \pi_2^i \setminus \pi_2^j$ with $u_2(o) > 0$, or $o \in \pi_2^j \setminus \pi_2^i$ with $u_2(o) < 0$ (the existence of o is guaranteed by Lemma 18);
20   **if** $u_2(o) > 0$ **then**
21   Set $\pi_2^i = \pi_2^i \setminus \{o\}$, $\pi_1^i = \pi_1^i \cup \{o\}$, $\pi_2^j = \pi_2^j \cup \{o\}$, and $\pi_1^j = \pi_1^j \setminus \{o\}$;
22   **else**
23   Set $\pi_2^j = \pi_2^j \setminus \{o\}$, $\pi_1^j = \pi_1^j \cup \{o\}$, $\pi_2^i = \pi_2^i \cup \{o\}$, and $\pi_1^i = \pi_1^i \setminus \{o\}$;
24   **if** π^i *is not envy-free for agent* 2 **then**
25   Set $F_2 = F_2 \setminus \{i\}$ and $E_2 = E_2 \cup \{i\}$;

We are now ready to prove Theorem 15.

Proof of Theorem 15. In the following, we refer to the first **while** loop (Lines 3 – 13) as the *first phase* and the second **while** loop (Lines 15 – 25) as the *second phase.*

To see that the first phase is well-defined, we show that during the execution of the first phase, F_1 is nonempty and corresponds to the rounds where agent 1 is envy-free. Indeed, consider the allocation π^j just after Line 11; since π^j is not weak EF1 before the transfer operation in Lines 9 and 11, agent 1 remains envious at π^j just after Line 11, and hence j continues to belong to E_1. Thus, during the execution of the first phase, E_1 corresponds to the set of the rounds where agent 1 envies the other while F_1 corresponds to the set of rounds where agent 1 is envy-free. Further, by Lemma 19, the algorithm keeps the property that the allocation $\bar{\pi}^{(k)}$ is PO and EF; by Pareto-optimality of $\bar{\pi}^{(k)}$, each π^i for $i \in [k]$ is Pareto-optimal within the round. Thus, at most one agent envies the other at each π^i. Therefore, during the execution of the first phase, $\bar{\pi}^{(k)}$ remains envy-free for agent 1, which implies that there exists a round i such that π^i is envy-free for agent 1, namely, F_1 is nonempty. By Lemma 18, there exists an item o satisfying the condition in Line 7. A similar argument shows that the second phase is also well-defined.

Next, we show that the first phase terminates in polynomial time and the allocation $\bar{\pi}^{(k)}$ just after the first phase is per-round weak EF1 for agent 1. To see this, consider π^i defined in Line 6. After the transfer operation in Lines 9 and 11, π^i does not violate weak EF1 from the viewpoint of agent 1; indeed, agent 1 may envy the other at π^i after the transfer operation, but since π^i is envy-free for agent 1 before, the envy can be eliminated by either stealing the good o ($u_1(o) > 0$) from the other agent, or transferring the chore o ($u_1(o) < 0$) to the other agent. Now, consider the allocation π^j that is not weak EF1 for agent 1. Clearly, at each iteration of the **while** loop in Lines 5 – 13, the number of goods o ($u_1(o) > 0$) increases while the number of chores o ($u_1(o) < 0$) decreases for agent 1. Thus, π^j becomes weak EF1 for agent 1 in $O(m)$ iterations. Thus, we conclude that the first phase terminates in polynomial time and the allocation $\bar{\pi}^{(k)}$ just after the first phase is per-round weak EF1 for agent 1.

A similar argument shows that the second phase terminates in polynomial time and the final allocation is per-round weak EF1 for agent 2. It remains to show that $\bar{\pi}^{(k)}$ remains weak EF1 for agent 1 during the second phase.

Consider an arbitrary iteration in the second phase and allocations π^i and π^j defined in Lines 16 and 18, respectively. Let o be an item chosen in Line 19. Assume that π^i and π^j are weak EF1 for agent 1 before the transfer operation in Lines 21 and 23. We show that both π^i and π^j remain weak EF1 for agent 1 just after the transfer operation. First, consider π^i. Since after the swap π^i remains Pareto-optimal within the round, we have:

- If $u_2(o) > 0$, then $u_1(o) > 0$.
- If $u_2(o) < 0$, then $u_1(o) < 0$.

Thus, agent 1's utility does not decrease after the transfer operation. Thus, π^i remains weak EF1 for agent 1.

Next, consider π^j. As we have observed before, agent 2 remains envious at π^j just after the transfer operation in Lines 21 and 23 since π^j is not weak EF1 for agent 2 before the transfer operation. This means that by Pareto-optimality of π, π^j is still envy-free for agent 1 after the transfer operation. Thus, π^j is weak EF1 for agent 1.

It is not difficult to see that Algorithm 1 runs in polynomial time. □

Theorem 17. *If $n = 2$ and $k \in 2\mathbb{N}$, then an allocation which is envy-free overall and per-round EF1 always exists, and can be computed in polynomial time.*

Proof. Consider the following procedure. Create a sequence of items $P = (o_1, o_2, \dots, o_m)$. For $j = 1, 2, \dots, m$, let $I_j = \{o_1, o_2, \dots, o_j\}$ and $\bar{I}_j = \{o_{j+1}, o_{j+2}, \dots, o_m\}$. Let $I_0 = \emptyset$ and $\bar{I}_0 = I_m$. Observe that we have either

- $u_1(I_0) \leqslant u_1(\bar{I}_0)$ and $u_1(I_m) \geqslant u_1(\bar{I}_m)$, or
- $u_1(I_0) \geqslant u_1(\bar{I}_0)$ and $u_1(I_m) \leqslant u_1(\bar{I}_m)$.

Thus, there exists an index $j \in \{1, 2, \ldots, m\}$ where the preference of agent 1 switches when o_j moves from the right to the left bundle, namely,

(i) $u_1(I_{j-1}) \leqslant u_1(\bar{I}_{j-1})$ and $u_1(I_j) \geqslant u_1(\bar{I}_j)$, or

(ii) $u_1(I_{j-1}) \geqslant u_1(\bar{I}_{j-1})$ and $u_1(I_j) \leqslant u_1(\bar{I}_j)$.

Let $L = I_{j-1}$ and $R = \bar{I}_j$. Assume without loss of generality that $u_1(L) \geqslant u_1(R)$.

First, suppose that we are in the first case (i), i.e., $u_1(I_{j-1}) \leqslant u_1(\bar{I}_{j-1})$ and $u_1(I_j) \geqslant u_1(\bar{I}_j)$. In other words, this means that agent 1 weakly prefer a bundle with o_j: $u_1(L) \leqslant u_1(R \cup \{o_j\})$ and $u_1(L \cup \{o_j\}) \geqslant u_1(R)$. Consider the following cases.

- Suppose that agent 2 weakly prefers L to the remaining items, or R to the remaining items, i.e., $u_2(L) \geqslant u_2(R \cup \{o_j\})$ or $u_2(R) \geqslant u_2(L \cup \{o_j\})$. Then, there is an envy-free allocation and hence the sequence that repeats this allocation k times is a desired solution. Indeed, if $u_2(L) \geqslant u_2(R \cup \{o_j\})$, then the allocation that allocates R together with o_j to agent 1 and L to agent 2 is an EF allocation. If $u_2(R) \geqslant u_2(L \cup \{o_j\})$, then the allocation that allocates L together with o_j to agent 1 and R to agent 2 is an EF allocation.

- Suppose that agent 2 weakly prefers $R \cup \{o_j\}$ to L, and $L \cup \{o_j\}$ to R, namely, $u_2(R \cup \{o_j\}) \geqslant u_2(L)$ and $u_2(L \cup \{o_j\}) \geqslant u_2(R)$. If $u_2(R) \leqslant u_2(L)$, this means that both agents weakly prefer L to R; thus, a sequence that repeatedly swaps two bundles L and $R \cup \{o_j\}$ among the two agents is a per-round EF1 and EF. On the other hand, if $u_2(L) \leqslant u_2(R)$, this means that the two agents have a different preference among L and R: while agent 1 weakly prefers L to R, agent 2 weakly prefers R to L. Thus, allocate L to agent 1 and R to agent 2 and alternate between assigning item o_j to agent 1 and to agent 2. Clearly, such a sequence is per-round EF1. To see that it is EF, observe that we have $u_1(L \cup \{o_j\}) \geqslant u_1(R \cup \{o_j\})$ and $u_1(L) \geqslant u_1(R)$ and thus agent 1 does not envy agent 2 at the k-round allocation. Similarly, we have $u_2(R \cup \{o_j\}) \geqslant u_2(L \cup \{o_j\})$ and $u_2(R) \geqslant u_2(L)$ and thus agent 2 does not envy agent 1 at the k-round allocation.

Next, suppose that we are in the second case (ii), i.e., $u_1(I_{j-1}) \geqslant u_1(\bar{I}_{j-1})$ and $u_1(I_j) \leqslant u_1(\bar{I}_j)$. In other words, agent 1 weakly prefer a bundle without o_j: $u_1(L) \geqslant u_1(R \cup \{o_j\})$ and $u_1(L \cup \{o_j\}) \leqslant u_1(R)$. Consider the following cases.

- Suppose that agent 2 weakly prefers $R \cup \{o_j\}$ to L, or $L \cup \{o_j\}$ to R, namely, $u_2(R \cup \{o_j\}) \geqslant u_2(L)$ or $u_2(L \cup \{o_j\}) \geqslant u_2(R)$. Then, we show that there is an envy-free allocation and hence the sequence that repeats this allocation k times is a desired solution. Indeed, if $u_2(R \cup \{o_j\}) \geqslant u_2(L)$, then the allocation giving L to agent 1 and $R \cup \{o_j\})$ to agent 2 is an EF allocation. Similarly, if $u_2(L \cup \{o_j\}) \geqslant u_2(R)$, then the allocation giving R to agent 1 and $L \cup \{o_j\})$ to agent 2 is an EF allocation.

- Suppose that agent 2 weakly prefers L to $R \cup \{o_j\}$, and R to $L \cup \{o_j\}$. If $u_2(L) \geqslant u_2(R)$, this means that both agents weakly prefer L to R; thus, a sequence that repeatedly swaps two bundles $L \cup \{o_j\}$ and R among the two agents is a per-round EF1 and EF.

  If $u_2(R) \geqslant u_2(L)$, then this means that the two agents have a different preference among L and R: while agent 1 weakly prefers L to R, agent 2 weakly prefers R to L. Thus, allocate L to agent 1 and R to agent 2 and alternate between assigning item o_j to agent 1 and to agent 2. Such a sequence is per-round EF1. To see that it is EF, observe that we have $u_1(L) \geqslant u_1(R)$ and $u_1(L \cup \{o_j\}) \geqslant u_1(R \cup \{o_j\})$ and thus agent 1 does not envy agent 2. Similarly, since $u_2(R \cup \{o_j\}) \geqslant u_2(L \cup \{o_j\})$ and $u_2(R) \geqslant u_2(L)$, agent 2 does not envy agent 1.

It is immediate to see that the above procedure can be implemented in polynomial time. This concludes the proof. □