

Repeated Fair Allocation of Indivisible Items

Ayumi Igarashi¹, Martin Lackner², Oliviero Nardi², Arianna Novaro³

¹The University of Tokyo

²DBAI, TU Wien

³CES, University of Paris 1 Panthéon-Sorbonne

igarashi@mist.i.u-tokyo.ac.jp, lackner@dbai.tuwien.ac.at, oliviero.nardi@tuwien.ac.at, arianna.novaro@univ-paris1.fr

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, there always exists a sequence of allocations that is proportional and Pareto-optimal. On the other hand, 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. Finally, in case that the number of repetitions can be chosen freely, we show that envy-free and Pareto-optimal allocations are achievable for any number of agents.

1 Introduction

In a variety of real-life scenarios, a group of agents has to divide a set of items among themselves. These items can be desirable (goods) or undesirable (chores) and agents have (heterogeneous) preferences concerning them. In the case of *goods*, we can think, for instance, of employees having to share access to some common infrastructure, like computing facilities. In the case of *chores*, we can think of roommates having to split household duties or teams having to split admin tasks. We may even 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 undesirable (negative) for others.

The examples above are all instances of problems of fair allocation of indivisible items (see the recent survey by

Amanatidis et al. (2022)). One of the challenges of fair division problems is that often, given the agents' preferences, it is impossible to find an allocation of the items which is both fair (e.g., no agent envies another agent's bundle) 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 occur *repeatedly*, 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, computing facilities may be needed daily or weekly by the same teams of employees, and household chores need to be done regularly.

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 be guaranteed when taking a global perspective on the overall sequence of repeated allocations? Can we additionally achieve fairness and efficiency guarantees at each individual repetition?

Contribution and outline. Our main contribution is the definition of a new (offline) 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). We consider two general cases: one where the number of repetitions is predetermined and given as part of the input and one where the number of repetitions can be chosen freely based on the instance. In the sequel, we refer to these as the *fixed* and *variable* cases, respectively.

In Section 3, we study our model for a fixed number of rounds (k) with n agents. We show that: if the number of rounds is a multiple of n , we can always guarantee the existence of a sequence of allocations that is envy-free overall in (Proposition 4), and a sequence of allocations that is proportional and Pareto-optimal (Theorem 7). This is essentially optimal in two respects. First, for any number $n \geq 2$ of

agents and any fixed number k of rounds which is not a multiple of n , a sequence of allocations that is proportional overall may not exist (Proposition 5). Moreover, for any number $n > 2$ of agents and any fixed number k of rounds, there is an instance of chore allocation where an envy-free and Pareto-optimal sequence of allocations may not exist (Theorem 6).

In Section 4, we will stay in the fixed- k setting but focus our model on the case of two agents. This scenario is fundamental in fair division, with numerous applications, including inheritance division, house-chore division, and divorce settlements (Brams and Fishburn 2000; Brams, Kilgour, and Klamler 2014; Kilgour and Vetschera 2018; Igarashi and Yokoyama 2023). Here, we show that 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 14). Moreover, for two rounds, we can strengthen the per-round fairness guarantee to envy-freeness up to one item (EF1) (Corollary 12). 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 15). These results turn out to be the best we can hope for. In Proposition 10, 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.

In Section 5, we investigate the model with a variable number of rounds. Within this model, we can in fact achieve overall and per-round fairness guarantees for the case of n agents. Specifically, we establish the existence of a sequence of allocations that satisfies envy-freeness and Pareto-optimality overall, while also meeting the per-round PROP[1,1] criterion (a weakening of proportionality). For scenarios involving goods only or chores only, we can achieve even per-round PROP1. To show this, we establish the connection between our model and the divisible and probabilistic fair division problems. An overview of our results can be found in Table 1. All the omitted proofs are presented in the full version of our paper (Igarashi et al. 2023).

Related work. Aziz et al. (2022) analyze fairness concepts for the allocation of indivisible goods and chores. Based on some of these results, Igarashi and Yokoyama (2023) 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 (Meir et al. 2021; Payan and Zick 2022), as well as the allocation of students to courses under capacity constraints (Othman, Sandholm, and Budish 2010). These articles focus however on “one-shot” (non-repeated) problems.

Repetitions have been studied in the context of matchings (Hosseini, Larson, and Cohen 2015; Gollapudi, Kollias, and Plaut 2020; Caragiannis and Narang 2022). This line of work uses models similar to ours. However, results do not carry over due to differences between matchings and arbitrary multi-unit assignments (as we consider). Repeated fair allocation is also related to probabilistic fair division (Budish et al. 2013; Aziz et al. 2023). We make this connection

precise in Section 5.

Various settings fall under the umbrella of *dynamic* or *online* fair decision-making (Aleksandrov and Walsh 2020; Kohler et al. 2014; Kash, Procaccia, and Shah 2014; Freeman et al. 2018; Benade et al. 2018; Zeng and Psomas 2020). Guo, Conitzer, and Reeves (2009) and Cavallo (2008) focus on a repeated setting, where a *single* item must be allocated in every round. We work in an *offline* setting, where agents have static and heterogeneous preferences over multiple items instead of demands, and the sets of agents and items is fixed. Balan, Richards, and Luke (2011) also study repeated allocations of items, but they focus on the average of utilities received by the agents for an allocation sequence.

In the context of elections, a closely related framework is that of *perpetual voting*, introduced by Lackner (2020), 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 analyze repeated instances of participatory budgeting problems (Lackner, Maly, and Rey 2021). Freeman, Zahedi, and Conitzer (2017) 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, Freeman, and Shah (2017).

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 assume that the allocation must be *exhaustive*: 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 $i, j \in N$, $\pi_i \cap \pi_j = \emptyset$. We write $[k]$ to denote $\{1, \dots, k\}$. Finally, given a positive integer ℓ , we denote by $\ell\mathbb{N}$ the set of positive integers multiples of ℓ .

Utilities. Each agent $i \in N$ is associated with a (dis)utility function $u_i : I \rightarrow \mathbb{R}$, 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 $\mathbf{u} = (u_1, \dots, u_n)$ the profile of utilities for the agents.

Fairness. We introduce several fairness concepts, as defined by Aziz et al. (2022) and by Amanatidis et al. (2022).

condition		result	reference
#agents (n)	#rounds (k)		
		fairness guarantee	
$n \geq 2$	$k \in n\mathbb{N}$	EF overall	Prop. 4
$n \geq 2$	$k \notin n\mathbb{N}$	PROP overall	Prop. 5
$n > 2$	$k \in \mathbb{N}$	EF+PO overall	Thm. 6
$\star n \geq 2$	$k \in n\mathbb{N}$	PROP+PO overall	Thm. 7
$\star n \geq 2$	variable k	EF+PO overall and per-round PROP[1,1] (PROP1 for only goods/chores)	Thm. 16
$n = 2$	$k > 2$	EF+PO overall and per-round EF1	Prop. 10
$\star n = 2$	$k = 2$	EF+PO overall and per-round EF1	Cor. 12
$\star n = 2$	even	EF+PO overall and per-round weak EF1	Cor. 14
$n = 2$	even	EF overall and per-round EF1	Thm. 15

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 (fixed) multiple of n . The main positive results are highlighted with \star .

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.

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

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 desiring it). 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 $u_i(\pi_i \setminus \{o\}) \geq u_i(\pi_j \setminus \{o\})$.

Yet another concept is that of *proportionality* of an allocation, where each agent receives their due share of utility.

Definition 3 (PROP). For agents N , items I , and profile \mathbf{u} , an allocation π satisfies proportionality if for each agent $i \in N$ we have $u_i(\pi_i) \geq u_i(I)/n$.

Note that envy-freeness implies proportionality when assuming additive utilities—see, e.g., Aziz et al. (2022, Prop. 1).

Finally, in a similar spirit to EF1, we can weaken proportionality to the notion of *proportionality up to one item* (PROP1) as well as its relaxed version (PROP[1,1]), to guarantee that agents receive a share of utility close to their proportional fair share. The notion of PROP1 has been proposed by Conitzer, Freeman, and Shah (2017) and later extended to the mixed setting by Aziz et al. (2022).

Definition 4. An allocation π is said to satisfy

- PROP1 if for each $i \in N$, $u_i((\pi_i \setminus X) \cup Y) \geq u_i(I)/n$ for some $X \subseteq \pi_i$ and $Y \subseteq I \setminus \pi_i$ with $|X \cup Y| \leq 1$.
- PROP[1,1] if for each $i \in N$, $u_i((\pi_i \setminus X) \cup Y) \geq u_i(I)/n$ for some $X \subseteq \pi_i$ and $Y \subseteq I \setminus \pi_i$ with $|X|, |Y| \leq 1$.

We note that PROP1 and PROP[1,1] coincide with each other when all items have non-negative values, since removing an item from an envious agent’s bundle does not increase

his or her utility. The same relation holds when all items have non-positive values. Note that Shoshan, Hazon, and Segal-Halevi (2023) recently introduced EF[1,1], which is an analogous version of our PROP[1,1] for envy-freeness.

Efficiency. Alongside fairness, it is often desirable to distribute the items as efficiently as possible. One way to capture efficiency is via 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 $i \in N$, $u_i(\rho_i) \geq u_i(\pi_i)$ 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 $\bar{\pi}^{(k)} = (\pi^1, \pi^2, \dots, \pi^k)$ 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 $I^1 = \{o_1^1, \dots, o_m^1\}$, \dots , $I^k = \{o_1^k, \dots, o_m^k\}$. Then, each 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 $\pi_i^{\cup k} = \pi_i^1 \cup \dots \cup \pi_i^k$ 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 omit the term “overall”). 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 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 property. 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 $\bar{\pi}^{(k)} = (\pi^1, \dots, \pi^k)$ is *per-round* EF1 if every allocation π^j for all $j \in [k]$ 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 a utility profile over two agents and two items where $u_1(o_1) = 4$, $u_1(o_2) = 5$, $u_2(o_1) = 3$ and $u_2(o_2) = 9$. Now consider a four-round allocation where agent 1 gets both items in the first and second round, 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 . \triangle

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.*

We can see that the converse does not necessarily hold. Indeed, consider a two-agent, two-round scenario with one objective good, where we give 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.

Finally, note that our *overall* results would also hold in the framework of one-shot allocations with k clones of each item. However, the latter model cannot capture the *per-round* results. Thus, our model cannot be reduced to the standard fair division setting with cloned items.

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. Start from an arbitrary initial allocation. Then, let each agent give their current bundle to the next agent in the next allocation (agent n gives theirs to agent 1). As each

agent i receives each bundle exactly k/n times, allocation $\pi^{\cup k}$ is envy-free. Thus, $\bar{\pi}^{(k)}$ is envy-free overall. \square

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

Proposition 5. *For every $n \geq 2$, 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.*

Intuitively, in a profile where every agent has utility 1 for every item, except for a desirable item o^* for which all agents have high utility $k(n-1)(m-1) + 1$, there must be some agent i receiving o^* less times than k/n (as k is not divisible by n).

If we also require Pareto-optimality, we get the following.

Theorem 6. *For all $n > 2$ and $k \in \mathbb{N}$, there exist a set of items I and a profile \mathbf{u} such that no allocation is envy-free and Pareto-optimal overall. The set I can be chosen to contain only two objective 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 a profile with a *small chore* s and a *big chore* b , such that every agent $i \in N$ has $u_i(s) = -1$, while $u_n(b) = -k/n$ and $u_j(b) = -k$ for every $j \in N \setminus \{n\}$.

One can show that in every overall EF sequence of allocations, all agents must receive s and b the same amount of times (namely k/n times). However, 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 agent 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). However, agent 2, e.g., now envies agent 1. \square

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 Pareto-optimal sequence of allocations.

Theorem 7. *For every $n \geq 2$ 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 toward 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. \square

Note that the existence guarantee in Theorem 7 is tight because of Proposition 5. Moreover, we can define an integer linear program (ILP) to compute such a proportional and PO allocation (Figure 1). Based on this, in the full version (Igarashi et al. 2023), we present some fixed-parameter tractability results.

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

Figure 1: An ILP for finding a proportional and PO allocation sequence for n agents. In this ILP, we maximize 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 Proposition 5.

4 The Case of Two Agents

In this section, we consider the case of repeated fair division among two agents. Several authors have focused on this special but important case (see, e.g., the papers by Brams, Kilgour, and Klamler (2012), Brams, Kilgour, and Klamler (2014), Kilgour and Vetschera (2018), Aziz et al. (2022), Shoshan, Hazon, and Segal-Halevi (2023), and the references therein). This captures some practically relevant problems, for example, house chore division among couples (Igarashi and Yokoyama 2023), or divorce settlements (Brams and Taylor 1996).

A first intuitive idea for the repeated setting could be to apply twice the *Generalized Adjusted Winner* algorithm introduced by Aziz et al. (2022), as in the one-shot setting it returns efficiently a PO and EF1 allocation of the goods

and chores, by alternating the roles of the winner and the loser. However, this approach may fail (see the full version (Igarashi et al. 2023)).

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

Now, we can show the following.

Theorem 9. *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. From the algorithm in the proof of Theorem 7 and from Proposition 8 we get an envy-free and Pareto-optimal allocation. Since $n = 2$, there are $(k + 1)^m$ possible allocations (up to symmetry-breaking), as each agent receives each of the m items either 0, 1, \dots , or k times. Thus, we obtain a run time of $O(m \cdot (k + 1)^m)$. \square

Theorem 9 provides strong overall fairness and efficiency guarantees, but it does not ensure per-round fairness. Indeed, an envy-free and PO allocation that is per-round EF1 may not exist, even if all items are goods or chores.

Proposition 10. *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 (sketch). First, observe that the number of rounds k needs to be even (Proposition 5). Next, consider the following utility profile over two items, where $u_1(o_1) = u_2(o_1) = 1$, $u_1(o_2) = 3$ and $u_2(o_2) = 2$. One can show that, for $k > 2$, in any per-round EF1 and overall EF allocation sequence, both agents should receive each item $k/2$ times. However, any such allocation 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. \square

Nevertheless, for two agents and two rounds, i.e., $n = k = 2$, it is always possible to achieve PO and EF as well as per-round EF1. Indeed, in this case, we can always transform an allocation that is PO and EF overall (which is guaranteed to exist for two agents) to an allocation that is per-round EF1 while preserving the PO and EF guarantees. A similar idea (i.e., exchanging items among two agents while preserving some initial property) was also used by Shoshan, Hazon, and Segal-Halevi (2023).

Theorem 11. *Suppose $k = n = 2$. Given an allocation sequence that is Pareto-optimal and envy-free, an allocation sequence which is Pareto-optimal, envy-free, and per-round EF1 can be computed in polynomial time.*

Proof (sketch). Consider an overall PO and EF allocation sequence $\bar{\pi}^{(2)}$ for two agents. Let $I_1 = \pi_1^1 \cap \pi_1^2$ (and similarly for I_2) be the items assigned to agent 1 (resp. 2) in both rounds. Moreover, let $O = I \setminus (I_1 \cup I_2)$ 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. Hence, all items in O are objective goods or chores. Let $O^+ \subseteq O$ be the objective goods in O , and let $O^- \subseteq O$ be the objective chores in O . Clearly, $O^+ \cap O^- = \emptyset$.

Update the allocation $\bar{\pi}^{(2)}$ such that $\pi_1^1 = I_1 \cup O^-$, $\pi_2^1 = I_2 \cup O^+$, $\pi_1^2 = I_1 \cup O^+$ and $\pi_2^2 = I_2 \cup O^-$. We then refine this sequence via the following procedure:

1. For each o in O , do the following:
 - (a) If $\bar{\pi}^{(2)}$ is per-round EF1, break the loop.
 - (b) Else, transfer item o as follows: 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 .
2. Return $\bar{\pi}^{(2)}$.




Figure 2: Exchanges of goods and chores between agents in Step 1(b) of the algorithm.

In other words, we start by assigning all chores in O to agent 1 in the first round (and to agent 2 in the second round), and conversely for the goods. Then, we progressively move the chores from agent 1 to agent 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 runs in polynomial time: it remains to show that it is correct.

First, observe that during the execution of the algorithm, $\bar{\pi}^{(2)}$ remains an allocation that is PO and EF overall, since we never change the total amount of times an agent receives an item w.r.t. the initial allocation. Thus, if the initial allocation is per-round EF1, we are done. We assume in the following that the initial allocation is not per-round EF1, which implies that at least one agent envies another in some round.

By overall EF, $2u_1(I_1) + u_1(O) \geq 2u_1(I_2) + u_1(O)$, which implies $u_1(I_1) \geq u_1(I_2)$. Similarly, we can show $u_2(I_2) \geq u_2(I_1)$. This implies $u_1(I_1 \cup O^+) \geq u_1(I_2 \cup O^-)$ and $u_2(I_2 \cup O^+) \geq u_2(I_1 \cup O^-)$ since all items in O^+ are objective goods and all items in O^- are objective chores.

Since $u_1(I_1 \cup O^+) \geq u_1(I_2 \cup O^-)$, if the algorithm moves all the objective chores in O^- from agent 1 to agent 2 and all the objective goods in O^+ from agent 2 to agent 1 in the first round, then agent 1 does not envy agent 2 with respect to the first round. Moreover, her utility in the first round never decreases, so that once agent 1 becomes envy-free in the first round, she remains envy-free in the latter steps. Similarly, if the algorithm moves all the objective chores in O^- from agent 2 to agent 1 and all the objective goods in O^+ from agent 1 to agent 2 in the second round, agent 2 does not

envy agent 1 with respect to the second round. In addition, once agent 2 becomes envy-free with respect to the second round, she remains envy-free in the latter steps.

Thus, there exists an item o such that after transferring o , agent 1 does not envy agent 2 in the first round, and agent 2 does not envy agent 1 in the second round. Let o be the first such item. By a case-analysis, one can show that at least one of the two allocations (produced by the algorithm) obtained before and after transferring o is per-round EF1. Thus, the algorithm correctly finds a desired allocation. \square

Combining the above with Theorem 9, we obtain the following corollary.

Corollary 12. *If $k = n = 2$, then an allocation which is Pareto-optimal, envy-free, and per-round EF1 always exists. Moreover, it can be computed in time $O(m \cdot 3^m)$.*

As we have seen in Proposition 10, 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 the following (weaker) version of EF1, where the envy of i toward j can be eliminated by either giving a good of j to i , or imposing a chore of i on j .

Definition 6 (Weak EF1). *An allocation π is weak EF1 if for all $i, j \in N$, either $u_i(\pi_i) \geq u_i(\pi_j)$ or there is an item $o \in \pi_i \cup \pi_j$ such that $u_i(\pi_i \cup \{o\}) \geq u_i(\pi_j \setminus \{o\})$ or $u_i(\pi_i \setminus \{o\}) \geq u_i(\pi_j \cup \{o\})$.*

Again, to prove the following theorem, 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 13. *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 14. *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, and can be computed in time $O(m \cdot (k+1)^m)$.*

In order to obtain Corollary 14, it suffices to prove Theorem 13, since we can apply Theorem 9 to get an allocation sequence $\bar{\pi}^{(k)}$ that is PO and envy-free overall when $k \in 2\mathbb{N}$.

To do so, suppose that we are given a PO and EF k -round allocation. Looking into each individual round, there may be an allocation π^ℓ that is not fair. Similarly to the proof of Theorem 11, 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 the full version (Igarashi et al. 2023).

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 15. *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 $k \in \mathbb{N}$ rounds (EF overall is possible in this case by Proposition 4). We leave it as an interesting open problem for future work.

5 Variable Number of Rounds

In the previous sections, we have assumed that k (the number of rounds) is predetermined. In this section, we study the case when k can be *variable*.

We have seen that, whenever k is fixed, overall EF and PO become incompatible in general (Theorem 6). This raises the question of whether this is possible if k is not predetermined. More precisely, we ask the following: given a utility profile, is there a number k of rounds for which there is a sequence of allocations $\bar{\pi}^{(k)}$ that satisfies desirable overall and per-round guarantees? We answer this question affirmatively.

Theorem 16. *If for all $i \in N$ and $o \in I$ we have $u_i(o) \in \mathbb{Q}$, then there exists $k \in \mathbb{N}$ and a $\bar{\pi}^{(k)}$ such that $\bar{\pi}^{(k)}$ is envy-free, Pareto-optimal, and per-round PROP[1,1].*

Recall that PROP[1,1] and PROP1 coincide for the case of non-negative value items and for the case of non-positive value items; thus, in such cases, the above theorem holds with a per-round guarantee of PROP1.

To show Theorem 16, we establish a connection between our setting and the divisible (Bogomolnaia et al. 2017) and probabilistic settings (Aziz et al. 2023; Budish et al. 2013). More precisely, we utilize a proof technique similar to that of Aziz et al. (2023), who used a decomposition lemma by Budish et al. (2013) for divisible item allocations to construct a randomized allocation that satisfies desirable efficiency and fairness notions when all items are goods. We prove that there exists a randomized allocation satisfying similar desirable guarantees for the mixed case of goods and chores and translate it into our repeated setting.

Before we proceed, we provide some preliminary definitions. Consider a set of n agents $N = [n]$ and of items I . A *randomized allocation* is a set of ordered pairs $(p^t, \pi^t)_{t \in [h]}$, such that for every $t \in [h]$, π^t is an allocation implemented with probability $p^t \in [0, 1]$, where $\sum_{t \in [h]} p_t = 1$; the allocations $\pi^1, \pi^2, \dots, \pi^h$ are referred to as the *support* of a randomized allocation.

Given a randomized allocation $(p^t, \pi^t)_{t \in [h]}$ with $p^t \in \mathbb{Q}$ for each $t \in [h]$, we define its *repeated translation* as follows. Since all p^t are rational numbers, there exists a $k \in \mathbb{N}$ such that, for all $t \in [h]$, $p_t = \ell_t/k$ (for some $\ell_t \in \mathbb{N}$). In other words, all fractions defined by $(p_t)_{t \in [h]}$ are expressed in terms of the same denominator, k . The repeated translation of $(p^t, \pi^t)_{t \in [h]}$ is the allocation sequence $\bar{\pi}^{(k)}$ of form:

$$\bar{\pi}^{(k)} = \underbrace{(\pi^1, \dots, \pi^1)}_{\ell_1 \text{ times}}, \underbrace{(\pi^2, \dots, \pi^2)}_{\ell_2 \text{ times}}, \dots, \underbrace{(\pi^h, \dots, \pi^h)}_{\ell_h \text{ times}}.$$

In other words, for each $t \in [h]$, π^t appears exactly ℓ_t times (where $p_t = \ell_t/k$). We can now prove the theorem.

Proof (sketch) of Theorem 16. A *fractional allocation* is a vector $\mathbf{x} = (x_1, \dots, x_n)$ where each $x_i = (x_{i,o})_{o \in I} \in \mathbb{Q}_+^m$

and, for all $o \in I$, $\sum_{i \in N} x_{i,o} = 1$. Given a utility vector \mathbf{u} , the utility of agent i for allocation \mathbf{x} is defined as $u_i(\mathbf{x}_i) = \sum_{o \in I} x_{i,o} u_i(o)$. The notions of envy-freeness and Pareto-optimality (w.r.t. all other fractional allocations) naturally translate to this setting. Next, a randomized allocation $(p^t, \pi^t)_{t \in [h]}$ is said to *implement* a fractional allocation \mathbf{x} if

$$x_{i,o} = \sum_{t=1}^h p_t \mathbf{1}[o \in \pi_i^t] \quad \text{for each } i \in N \text{ and } o \in I.$$

Here, $\mathbf{1}[o \in \pi_i^t]$ is an indicator function which takes value 1 if $o \in \pi_i^t$, and 0 otherwise.

When all agents' utilities are rational-valued, a PO and EF fractional allocation always exists (Bogomolnaia et al. 2017; Chaudhury et al. 2022). Moreover, using the decomposition lemma of Budish et al. (2013), we can show that any PO and EF fractional allocation \mathbf{x} admits a randomized allocation $(p^t, \pi^t)_{t \in [h]}$ that implements \mathbf{x} with these properties:

1. $p_1, p_2, \dots, p_h \in \mathbb{Q}_+$,
2. $\pi_i^t \subseteq \{o \in I \mid x_{i,o} > 0\}$ for each $i \in N$ and $t \in [h]$, and
3. π^t is PROP[1,1] for each $t \in [h]$.

See the full version (Igarashi et al. 2023) for details.

Finally, one can show that the repeated translation of $(p^t, \pi^t)_{t \in [h]}$ is PO and EF overall. Moreover, since π^t is PROP[1,1] for all $t \in [h]$, this translation is per-round PROP[1,1]. \square

Note that due to Theorem 6, we know that this result is impossible for a fixed number of rounds, and thus highlights the difference between fixed and variable k .

6 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 some algorithms guaranteeing overall envy-freeness and Pareto-optimality (as well as per-round approximate envy-freeness) for any even number of rounds.

As some of our algorithms require 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.

Our n -agent algorithm yields slightly weaker guarantees (proportionality and Pareto-optimality), which are still an improvement over the one-shot setting. It remains for future work to determine whether this result can be strengthened by additional per-round guarantees, as for the 2-agent case.

When the number of rounds k can be chosen freely, we have shown that (for any number of agents) an envy-free, Pareto-optimal and per-round PROP[1,1] allocation always exists. We have done so by establishing a connection between our setting and the probabilistic and divisible settings. However, our approach gives no guarantee on the number of rounds it requires. As future work, it would be interesting to investigate the complexity of finding the smallest k for which an envy-free and Pareto-optimal allocation exists.

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): P31890. Oliviero Nardi was supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement number 101034440, and by the Vienna Science and Technology Fund (WWTF) through project ICT19-065. Ayumi Igarashi was supported by JST PRESTO under grant number JPMJPR20C1.

References

- Aleksandrov, M.; and Walsh, T. 2020. Online Fair Division: A Survey. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-2020)*, 13557–13562.
- Amanatidis, G.; Aziz, H.; Birmpas, G.; Filos-Ratsikas, A.; Li, B.; Moulin, H.; Voudouris, A. A.; and Wu, X. 2022. Fair Division of Indivisible Goods: Recent Progress and Open Questions. arXiv:2208.08782.
- Aziz, H.; Caragiannis, I.; Igarashi, A.; and Walsh, T. 2022. Fair Allocation of Indivisible Goods and Chores. *Autonomous Agents and Multi-Agent Systems*, 36(1): 1–21.
- Aziz, H.; Freeman, R.; Shah, N.; and Vaish, R. 2023. Best of Both Worlds: Ex-Ante and Ex-Post Fairness in Resource Allocation. *Operations Research*. Published online in Articles in Advance.
- Balan, G.; Richards, D.; and Luke, S. 2011. Long-Term Fairness with Bounded Worst-Case Losses. *Autonomous Agents and Multi-Agent Systems*, 22(1): 43–63.
- Benade, G.; Kazachkov, A. M.; Procaccia, A. D.; and Psoomas, C.-A. 2018. How to Make Envy Vanish Over Time. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC-2018)*, 593–610.
- Bogomolnaia, A.; Moulin, H.; Sandomirskyi, F.; and Yanovskaya, E. 2017. Competitive Division of a Mixed Manna. *Econometrica*, 85(6): 1847–1871.
- Brams, S. J.; and Fishburn, P. C. 2000. Fair Division of Indivisible Items Between Two People with Identical Preferences: Envy-Freeness, Pareto-Optimality, and Equity. *Social Choice and Welfare*, 17(2): 247–267.
- Brams, S. J.; Kilgour, D. M.; and Klamler, C. 2012. The Undercut Procedure: An Algorithm for the Envy-Free Division of Indivisible Items. *Social Choice and Welfare*, 39: 615–631.
- Brams, S. J.; Kilgour, M.; and Klamler, C. 2014. Two-Person Fair Division of Indivisible Items: An Efficient, Envy-Free Algorithm. *Notices of the AMS*, 61(2): 130–141.
- Brams, S. J.; and Taylor, A. D. 1996. A Procedure for Divorce Settlements. *Mediation Quarterly*, 13(3): 191–205.
- Budish, E.; Che, Y.-K.; Kojima, F.; and Milgrom, P. 2013. Designing Random Allocation Mechanisms: Theory and Applications. *American Economic Review*, 103(2): 585–623.
- Caragiannis, I.; and Narang, S. 2022. Repeatedly Matching Items to Agents Fairly and Efficiently. arXiv:2207.01589.
- Cavallo, R. 2008. Efficiency and Redistribution in Dynamic Mechanism Design. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, 220–229.
- Chaudhury, B. R.; Garg, J.; McGlaughlin, P.; and Mehta, R. 2022. A Complementary Pivot Algorithm for Competitive Allocation of a Mixed Manna. *Mathematics of Operations Research*. Published online in Articles in Advance.
- Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair Public Decision Making. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC-2017)*, 629–646.
- Freeman, R.; Zahedi, S. M.; and Conitzer, V. 2017. Fair Social Choice in Dynamic Settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-2017)*, 4580–4587.
- Freeman, R.; Zahedi, S. M.; Conitzer, V.; and Lee, B. C. 2018. Dynamic Proportional Sharing: A Game-Theoretic Approach. In *Proceedings of the ACM on Measurement and Analysis of Computing Systems (MACS-2018)*.
- Gollapudi, S.; Kollias, K.; and Plaut, B. 2020. Almost Envy-Free Repeated Matching in Two-Sided Markets. In *Web and Internet Economics: 16th International Conference, WINE 2020, Beijing, China, December 7–11, 2020, Proceedings*, 3–16.
- Guo, M.; Conitzer, V.; and Reeves, D. M. 2009. Competitive Repeated Allocation without Payments. In *International Workshop on Internet and Network Economics*, 244–255.
- Hosseini, H.; Larson, K.; and Cohen, R. 2015. Matching with Dynamic Ordinal Preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-2015)*.
- Igarashi, A.; Lackner, M.; Nardi, O.; and Novaro, A. 2023. Repeated Fair Allocation of Indivisible Items. arXiv:2304.01644.
- Igarashi, A.; and Yokoyama, T. 2023. Kajibuntan: A House Chore Division App. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, 16449–16451.
- Kash, I.; Procaccia, A. D.; and Shah, N. 2014. No Agent Left Behind: Dynamic Fair Division of Multiple Resources. *Journal of Artificial Intelligence Research*, 51: 579–603.
- Kilgour, D. M.; and Vetschera, R. 2018. Two-Player Fair Division of Indivisible Items: Comparison of Algorithms. *European Journal of Operational Research*, 271(2): 620–631.
- Kohler, T.; Steghöfer, J.-P.; Busquets, D.; and Pitt, J. 2014. 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)*, 1–10.
- Lackner, M. 2020. Perpetual Voting: Fairness in Long-Term Decision Making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-2020)*, 2103–2110.
- Lackner, M.; Maly, J.; and Rey, S. 2021. Fairness in Long-Term Participatory Budgeting. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI-2021)*, 299–305.

- Meir, R.; Lang, J.; Lesca, J.; Mattei, N.; and Kaminsky, N. 2021. A Market-Inspired Bidding Scheme for Peer Review Paper Assignment. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI-2021)*, 4776–4784.
- Othman, A.; Sandholm, T.; and Budish, E. 2010. Finding Approximate Competitive Equilibria: Efficient and Fair Course Allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2010)*, 873–880.
- Payan, J.; and Zick, Y. 2022. I Will Have Order! Optimizing Orders for Fair Reviewer Assignment. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI-2022)*, 440–446.
- Shoshan, H.; Hazon, N.; and Segal-Halevi, E. 2023. Efficient Nearly-Fair Division with Capacity Constraints. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2023)*, 206–214.
- Zeng, D.; and Psomas, A. 2020. Fairness-Efficiency Trade-offs in Dynamic Fair Division. In *Proceedings of the 2020 ACM Conference on Economics and Computation (EC-2020)*, 911–912.