Temporal Fair Division of Indivisible Items

Edith Elkind
Northwestern University
Evanston, Illinois, United States
edith.elkind@northwestern.edu

Alexander Lam
Hong Kong Polytechnic University
Hong Kong, Hong Kong SAR
alexander-a.lam@polyu.edu.hk

Mohamad Latifian
University of Edinburgh
Edinburgh, United Kingdom
mohamad.latifian@ed.ac.uk

Tzeh Yuan Neoh
IHPC, CFAR, A*STAR
Singapore, Singapore
neoh_tzeh_yuan@cfar.a-star.edu.sg

Nicholas Teh
University of Oxford
Oxford, United Kingdom
nicholas.teh@cs.ox.ac.uk

ABSTRACT

We study a fair division model where indivisible items arrive sequentially, and must be allocated immediately and irrevocably. Previous work on online fair division has shown impossibility results for achieving approximate envy-freeness under the assumption that agents have no information about future items. In contrast, we assume that the algorithm has complete knowledge of the future, and aim to ensure that the cumulative allocation at each round satisfies approximate envy-freeness, which we define as *temporal envy-freeness up to one item* (TEF1). We focus on settings where items are exclusively goods or exclusively chores. For goods, while TEF1 allocations may fail to exist, we identify several special cases where they do—two agents, two item types, generalized binary valuations, unimodal preferences—and provide polynomial-time algorithms for these cases. We also prove that determining the existence of a TEF1 allocation is NP-hard. For chores, we obtain analogous results for the special cases, but present a slightly weaker intractability result. We also show that TEF1 is incompatible with Pareto optimality, with the implication that it is intractable tofind a TEF1 allocation that maximizes any p-mean welfare, even for two agents.

KEYWORDS

Fair Division, Envy-Freeness, Temporal Fairness

ACM Reference Format:
Edith Elkind, Alexander Lam, Mohamad Latifian, Tzeh Yuan Neoh, and Nicholas Teh. 2025. Temporal Fair Division of Indivisible Items. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 27 pages.

1 INTRODUCTION

Fair division, a topic at the intersection of economics and computer science, has been extensively studied over the years, with applications ranging from divorce settlements and inheritance disputes to load balancing [23, 57]. Typically, in fair division there is a set of agents and a set of items, and the goal is to obtain a fair allocation of items to agents. In our work, we study a model where these items are *indivisible*, so each must be wholly allocated to an agent. Moreover, the items can provide either positive utility (in which case they are called *goods*) or negative utility (in which case they are called *chores*, or *tasks*). When allocating indivisible items, a desirable and widely-studied fairness notion is *envy-freeness up to one item (EF1)*, a natural relaxation of *envy-freeness (EF)*. In an envy-free allocation, each agent values the bundle of items they receive at least as highly as every other agent's bundle. However, this desideratum is not always achievable for indivisible items (consider two agents and a single item that they both value). In contrast, in an EF1 allocation, the envy that agent A has towards another agent B can be eliminated by removing a single item from B's bundle (in case of goods) or A's bundle (in case of chores).

Most prior research studies fair division in the *offline* setting, assuming that all of the items are immediately available and ready to be allocated. However, there are various applications where the items arrive and need to be allocated on the spot in a sequential manner. For example, when the university administration places an order for lab equipment, or when a company orders new machines for its franchises, the items may arrive over time due to their availabilities and delivery logistics. In case of chores, collaborative project management may require division of tasks over time. For a variety of reasons, arriving items may have to be allocated immediately; there may not be any storage space to keep any unallocated goods, or the central decision maker may desire a non-wasteful allocation in the sense that items or tasks should not sit idle for periods of time.

These applications can be captured by an *online fair division* model, in which items arrive over time and must be immediately and irrevocably allocated, though it is assumed that each item's valuation is not known until its arrival. Prior research has found that a complete EF1 allocation of goods cannot be guaranteed under the online fair division model[1] [19]. However, this result relies on the assumption that the algorithm has no information about the future. In contrast, in our examples, the delivery services could provide the estimated delivery dates, and there may be a pre-planned timeline for the tasks. Motivated by these nuances, our work studies the *informed online fair division setting*, assuming that the algorithm can access the items' valuations and arrival order upfront.

Note that the assumption of complete information about the future trivially leads to a complete EF1 allocation at the *end* of the allocation period: simply treat the instance as an offline problem

[1] In fact, the maximum pairwise envy is $\Omega(\sqrt{t})$ after t rounds in the worst case.

This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

and apply any algorithm known to satisfy EF1 (e.g., [8, 25, 53]). However, this approach ignores the cumulative bundles of the items throughout the allocation period, and, consequently, agents may feel that their partial allocations are unfair for extended periods of time. Inspired by this issue, we propose *temporal EF1 (TEF1)*, an extension of EF1 to the informed online fair division setting which requires that at each round, the cumulative allocation satisfies EF1.

The main focus of our work is on achieving TEF1, so for the informed online fair division of indivisible goods or chores, we aim to answer the following existence and computational questions:

> *Which restricted settings guarantee the existence of a TEF1 allocation, and can we compute such an allocation in polynomial time in these settings? Is it computationally tractable to determine the existence[2] of a TEF1 allocation? In terms of existence and tractability, is TEF1 compatible with natural notions of efficiency?*

1.1 Our Contributions

We outline our paper's answers to these key questions as follows.

In Section 3, we show the existence of TEF1 allocations (for goods or chores) in restricted settings, such as the case of two agents, when there are two types of items, when agents have generalized binary valuations, or when they have unimodal preferences. For each of these cases, we provide an accompanying polynomial-time algorithm. For the allocation of goods, we show that determining whether there exists a TEF1 allocation is NP-hard; whereas for chores, we show that given a partial TEF1 allocation, it is NP-hard to determine if there exists a TEF1 allocation that allocates all remaining chores.

In Section 4, we investigate the compatibility of TEF1 and Pareto-optimality (PO). We show that even in the case of two agents, while a TEF1 allocation is known to exist and can be computed in polynomial time (for both goods and chores), existence is no longer guaranteed if we mandate PO as well. Moreover, we show that in this same setting, determining the existence of TEF1 and PO allocations is NP-hard. Our result also directly implies the computational intractability of determining whether there exists a TEF1 allocation that maximizes any p-mean welfare objective (which subsumes most popular social welfare objectives).

Finally, in Section 5, we consider the special case where the same set of items arrive at each round, and show that even determining whether repeating a particular allocation in two consecutive rounds can result in a TEF1 allocation is NP-hard. We complement this with a polynomial-time algorithm for computing a TEF1 allocation in this case when there are just two rounds.

1.2 Related Work

Our work is closely related to *online fair division*, whereby items arrive over time and must be irrevocably allocated to agents. The key difference is that in the standard online setting, the algorithm has completely no information on future items, whereas we assume complete future information. Moreover, the goal in online fair division models is typically to guarantee a fair allocation to agents at the end of the time horizon, rather than at every round.

[2]Prior work by He et al. [47] has shown that a TEF1 allocation is not guaranteed to exist for goods in the general setting with three or more agents.

As we focus on EF1, papers satisfying envy-based notions in online allocations are particularly relevant. Aleksandrov et al. [2] consider envy-freeness from both ex-ante and ex-post standpoints, giving a best-of-both-worlds style result by designing an algorithm for goods which is envy-free in expectation and guarantees a bounded level of envy-freeness. Additionally, Benadè et al. [19] find that allocating goods uniformly at random leads to maximum pairwise envy which is sublinear in the number of rounds. For further reading, we refer the reader to the surveys by Aleksandrov and Walsh [4] and Amanatidis et al. [6].

There has also been work on online fair division with *partial information* on future items. Benadè et al. [19] study the extent to which approximations of envy-freeness and Pareto efficiency can be simultaneously satisfied under a spectrum of information settings, ranging from identical agents and i.i.d. valuations to zero future information. An emerging line of work on *learning-augmented* online algorithms has an alternate approach to partial future information: the algorithms are aided by (possibly inaccurate) predictions, typically from a machine-learning algorithm. The focus is to design algorithms which perform *consistently* well with accurate predictions, and are *robust* under inaccurate predictions. These predictions could be of each agent's total utility for the entire item set [11, 12], or for a random subset of k incoming items [20].

Unlike the aforementioned papers, our work considers a completely *informed* variant of the online fair division setting, which has been studied by He et al. [47] for the allocation of goods. Similar to our paper, their objective is to ensure that EF1 is satisfied at each round, but they allow agents to swap their bundles. When multiple goods may arrive at each round, our setting also generalizes the *repeated* fair division setting, in which the same set of goods arrives at each round. For this model, Igarashi et al. [48] give results on the existence of allocations which are envy-free and Pareto optimal in the end, with the items in each round being allocated in an EF1 manner. However, they do not analyse whether the *cumulative* allocation at each round can satisfy some fairness constraint, which is the focus of our paper. Caragiannis and Narang [26] also consider a model where the same set of items appear at each round, but each agent gets exactly one item per round.

When the valuations are known upfront for the allocation of chores, the model is similar to thefi eld of work on job scheduling. There have been numerous papers studying fair scheduling, but the fairness is typically represented by an objective function which the algorithm aims to minimize or approximate [17, 49, 64]. On the other hand, there is little work on satisfying envy-based notions in scheduling problems, but Li et al. [52] study the compatibility of EF1 and Pareto optimality in various settings. While we consider separately the cases of goods allocation and chores allocation, to the best of our knowledge, there is no prior work which studies an online fair division model with both goods and chores in the same instance under any information assumption.

Similar temporal models that study concepts of achieving fairness over time have also been recently considered in the social choice literature [5, 27, 36, 38, 39, 51, 54, 58, 59, 70].

In a contemporary and independent work, Cookson et al. [32] consider the same setting of temporal fair division but with a different approach. They only consider goods and prove positive result in the three settings: when there are only two agents, when the identical set of goods appear in each timestep, and when agents have an identical ranking over the items. They consider several fairness notions and seek to achieve different pairs of these notions per-day and overall.

2 PRELIMINARIES

For each positive integer k, let $[k] := \{1, \ldots, k\}$. We consider the problem of fairly allocating indivisible items to agents over multiple rounds. An instance of the *informed online fair division problem* is a tuple $\mathcal{I} = \langle N, T, \{O_t\}_{t \in [T]}, \mathbf{v} = (v_1, \ldots, v_n) \rangle$, where $N = [n]$ is a set of *agents*, T is the number of *rounds*, for each $t \in [T]$ the set O_t consists of items that arrive at round t, with $O = \cup_{t \in [T]} O_t$, and for each $i \in N$ the *valuation function* $v_i : O \to \mathbb{R}$ specifies the values that agent i assigns to items in O.

We assume that agents have additive valuations, i.e., we extend the functions v_i to subsets of O by setting $v_i(S) = \sum_{o \in S} v_i(O)$ for each $S \subseteq O$. We write v instead of v_i when all agents have identical valuation functions. We refer to the vector $\mathbf{v} = (v_1, \ldots, v_n)$ as the *valuation profile*. We define the *cumulative set of items* that arrive in rounds $1, \ldots, t$ by $O^t := \bigcup_{\ell \in [t]} O_\ell$. Note that $O = O^T$.

We consider both *goods allocation*, where $v_i(o) \geq 0$ for each $i \in N$ and $o \in O$, and *chores allocation*, where $v_i(o) \leq 0$ for each $i \in N$ and $o \in O$. For clarity, in the goods setting we use g instead of o and refer to the items as *goods*, while in the chores setting we use c instead of o and refer to the items as *chores*.

An *allocation* $\mathcal{A} = (A_1, \ldots, A_n)$ of items in O to the agents is an ordered partition of O, i.e., $\bigcup_{i \in N} A_i = O$ and $A_i \cap A_j = \varnothing$ for all $i,j \in N$ with $i \neq j$. For $t \in [T]$, $i \in N$ we write $A_i^t = A_i \cap O^t$; then $\mathcal{A}^t = (A_1^t, \ldots, A_n^t)$ is the allocation after round t, with $\mathcal{A} = \mathcal{A}^T$. For $t < T$, we may refer to $\mathcal{A}^t$ as a *partial allocation*.

Our goal is tofi nd an allocation that is fair after each round. The main fairness notion that we consider is *envy-freeness up to one item (EF1)*, a well-studied notion in fair division.

Definition 2.1. In a goods (resp., chores) allocation instance, an allocation $\mathcal{A} = (A_1, \ldots, A_n)$ is said to be EF1 if for each pair of agents $i,j \in N$ there exists a good $g \in A_j$ (resp., chore $c \in A_i$) such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ (resp. $v_i(A_i \setminus \{c\}) \geq v_i(A_j)$).

To capture fairness in a *cumulative* sense, we introduce the notion of *temporal envy-freeness up to one item (TEF1)*, which requires that at every prefix of rounds the cumulative allocation of items that have arrived so far satisfies EF1.

Definition 2.2 (Temporal EF1). For every $t \in [T]$, an allocation $\mathcal{A}^t = (A_1^t, \ldots, A_n^t)$ is said to be *temporally envy-free up to one item (TEF1)* if for each $\ell \in [t]$ the allocation $\mathcal{A}^\ell$ is EF1.

A key distinction between TEF1 and EF1 is that, while the EF1 property only places constraints on thefi nal allocation, TEF1 requires envy-freeness up to one item at every round.

However, He et al. [47, Thm. 4.2] show that for goods TEF1 allocations may fail to exist; they present an example with 3 agents and 23 items, which can be generalized to $n > 3$ agents. For completeness, we include this counterexample along with an intuitive explanation in the full version of the paper. We remark that the construction of He et al. [47] cannot be translated to the chores setting. Indeed, while we conjecture that a non-existence result of this form also holds for chores, this remains an open question.

We assume that the reader is familiar with basic notions of classic complexity theory [61]. All omitted proofs can be found in the full version of the paper [37].

3 ON THE EXISTENCE OF TEF1 ALLOCATIONS

As some instances do not admit TEF1 allocations, ourfi rst goal is to explore if there are restricted classes of instances for which TEF1 allocations are guaranteed to exist. In this section we identify several such settings.

To simplify the presentation, we willfi rst demonstrate that it usually suffices to consider instances where only one item appears at each round (i.e., $T = m$ and $|O_t| = 1$ for all $t \in [T]$). Indeed, any impossibility result for this special setting also holds for the general case, and we will now argue that the converse is true as well.

LEMMA 3.1. *Given an instance $\mathcal{I}$ with $|O| = m$ items, we can construct an instance $\mathcal{I}^{=1}$ with the same set of items and exactly m rounds so that $|O_t| = 1$ for each $t \in [m]$ and if $\mathcal{I}^{=1}$ admits a* TEF1 *allocation, then so does $\mathcal{I}$.*

PROOF. Consider an arbitrary instance $\mathcal{I} = \langle N, T, \{O_t\}_{t \in [T]}, \mathbf{v} = (v_1, \ldots, v_n) \rangle$. Renumber the items in a non-decreasing fashion with respect to the rounds, so that for any two rounds $t, r \in [T]$ with $t < r$ and items $o_j \in O_t, o_{j'} \in O_r$ it holds that $j < j'$. We construct $\mathcal{I}^{=1} = \langle N, m, \{\widetilde{O}_t\}_{t \in [m]}, \mathbf{v} \rangle$ by setting $\widetilde{O}_t = \{o_t\}$ for each $t \in [m]$. Let $\mathcal{A}$ be a TEF1 allocation for $\mathcal{I}^{=1}$. We construct an allocation $\mathcal{B}$ for instance $\mathcal{I}$ by allocating all items in the same way as in $\mathcal{A}$: if $\mathcal{A}$ allocates at item j to agent i in round r, we identify a $t \in [T]$ such that $\sum_{\ell=1}^{t-1} |O_\ell| < r \leq \sum_{\ell=1}^{t} |O_\ell|$ and place j into B_i in round t. To see that $\mathcal{B}$ satisfies TEF1, note that if $\mathcal{B}^t$ violates EF1 for some $t \in [T]$, then for $r = \sum_{\ell=1}^{t} |O_\ell|$ the allocation $\mathcal{A}^r$ satisfies $A_i^r = B_i^t$ for all $i \in N$ and hence violates EF1 as well. □

In what follows, unless specified otherwise, we simplify the notation based on the transformation in the proof of Lemma 3.1: we assume that $|O_t| = 1$ for each $t \in T$ and denote the unique item that arrives in round t by o_t (or g_t, or c_t, if we focus on goods/chores).

3.1 Two Agents

He et al. [47, Thm. 3.4] put forward a polynomial-time algorithm that always outputs a TEF1 allocation for goods when $n = 2$; in particular, this implies that a TEF1 allocation is guaranteed to exist for $n = 2$. We will now extend this result to the case of chores.

Intuitively, in each round the algorithm greedily allocates the unique chore that arrives in that round to an agent that does not envy the other agent in the current (partial) allocation. A counter s keeps track of the last round in which $\mathcal{A}^s$ was envy-free; if for some round $t \in [m]$ the allocation of a chore c_t results in both agents envying each other in $\mathcal{A}^t \setminus \mathcal{A}^s$, then the agents' bundles in $\mathcal{A}^t \setminus \mathcal{A}^s$ are swapped.

THEOREM 3.2. *For $n = 2$, Algorithm 1 returns a* TEF1 *allocation for chores, and runs in polynomial time.*

PROOF. The polynomial runtime of Algorithm 1 is easy to verify: there is only one **for** loop, with a counter that runs from 1 to m, and each operation within the loop runs in $O(m)$ time. Thus, we focus on proving correctness.

Algorithm 1 Returns a TEF1 allocation for chores when $n = 2$

Input: Set of agents $N = \{1, \ldots, n\}$, set of chores $O = \{c_1, \ldots, c_m\}$, and valuation profile $\mathbf{v} = (v_1, v_2)$
Output: TEF1 allocation $\mathcal{A}$ of chores in O to agents in N

1: Initialize $s \leftarrow 0$ and $\mathcal{A}^0 \leftarrow (\varnothing, \varnothing)$
2: **for** $t = 1, 2, \ldots, m$ **do**
3:   **if** $v_1(A_1^{t-1} \setminus A_1^s) \geq v_1(A_2^{t-1} \setminus A_2^s)$ **then**
4:     $\mathcal{A}^t \leftarrow (A_1^{t-1} \cup \{c_t\}, A_2^{t-1})$
5:   **else**
6:     $\mathcal{A}^t \leftarrow (A_1^{t-1}, A_2^{t-1} \cup \{c_t\})$
7:   **end if**
8:   **if** $v_1(A_1^t \setminus A_1^s) < v_1(A_2^t \setminus A_2^s)$ and $v_2(A_2^t \setminus A_2^s) < v_2(A_1^t \setminus A_1^s)$ **then**
9:     $\mathcal{A}^t \leftarrow (A_1^s \cup A_2^t \setminus A_2^s, A_2^s \cup A_1^t \setminus A_1^s)$
10:   **end if**
11:   **if** $v_1(A_1^t \setminus A_1^s) \geq v_1(A_2^t \setminus A_2^s)$ and $v_2(A_2^t \setminus A_2^s) \geq v_2(A_1^t \setminus A_1^s)$ **then**
12:     $s \leftarrow t$
13:   **end if**
14: **end for**
15: **return** $\mathcal{A} = (A_1^m, A_2^m)$

For each $t \in [m]$, we define $r_t < t$ as the latest round before t such that $\mathcal{A}^{r_t}$ is EF. This implies that if $\mathcal{A}^\ell \setminus \mathcal{A}^{r_t}$ is EF1 for all $\ell = r_t, r_t + 1, \ldots, t$, then $\mathcal{A}^t$ is also EF1. Therefore, it suffices to show that $\mathcal{A}^t \setminus \mathcal{A}^{r_t}$ is EF1 for each $t \in [m]$. We will prove this by induction on t.

For $t = 1$, the claim is immediate, as any allocation of a single chore is EF1. Now, suppose that $t > 1$. If $t = r_t + 1$ the allocation $\mathcal{A}^t \setminus \mathcal{A}^{r_t}$ consists of a single chore, so, again, the claim is immediate. Otherwise, $r_{t-1} = r_t$ and by the induction hypothesis it holds that $\mathcal{A}^{t-1} \setminus \mathcal{A}^{r_t}$ is EF1. Let r'_t be the earliest round ahead of r_t such that $\mathcal{A}^{r'_t} \setminus \mathcal{A}^{r_t}$ is EF (if such a round exists). We divide the remainder of the proof into two cases depending on whether a partial bundle swap (as in line 9 of the algorithm) occurs at round r'_t.

Case 1: Round r'_t does not exist or no swap at round r'_t. Suppose without loss of generality that $v_1(A_1^{t-1} \setminus A_1^{r_t}) < v_1(A_2^{t-1} \setminus A_2^{r_t})$, i.e., agent 1 envies agent 2 in $\mathcal{A}^{t-1} \setminus \mathcal{A}^{r_t}$. Then agent 2 does not envy agent 1 (otherwise we would swap the bundles, contradicting the definition of r_t), and consequently receives c_t. If agent 2 envies agent 1 after receiving c_t, this envy can be removed by removing c_t. We also know that $\mathcal{A}^t \setminus \mathcal{A}^{r_t}$ is EF1 w.r.t. agent 1 (who did not receive a chore in round t) because by our inductive assumption, $\mathcal{A}^{t-1} \setminus \mathcal{A}^{r_t}$ is EF1, concluding the proof of this case.

Case 2: Swap occurs at round r'_t. We assume that $r'_t > t$, because if $r'_t = t$, then $\mathcal{A}^t \setminus \mathcal{A}^{r_t}$ is EF and therefore EF1. For each $i \in \{t-1, t\}$, let $\mathcal{B}^i \setminus \mathcal{A}^s$ refer to the algorithm's allocation of the chores $O^i \setminus O^s$ *before* the bundle swap, and suppose without loss of generality that $v_1(B_1^{t-1} \setminus B_1^{r_t}) < v_1(B_2^{t-1} \setminus B_2^{r_t})$. We therefore must have $v_2(B_2^{t-1} \setminus B_2^{r_t}) \geq v_2(B_1^{t-1} \setminus B_1^{r_t})$ to avoid contradicting the definition of r_t. Since agent 2 is not envied by agent 1 in round $t-1$, it receives chore c_t, so we have $\mathcal{B}^t \setminus \mathcal{B}^{r_t} = (B_1^{t-1} \setminus B_1^{r_t}, (B_2^{t-1} \setminus B_2^{r_t}) \cup \{c_t\})$. This means that after the bundle swap is executed, we have $\mathcal{A}^t \setminus \mathcal{A}^{r_t} = ((B_2^{t-1} \setminus B_2^{r_t}) \cup \{c_t\}, B_1^{t-1} \setminus B_1^{r_t})$. Recall that $v_1(B_2^{t-1} \setminus B_2^{r_t}) > v_1(B_1^{t-1} \setminus B_1^{r_t})$, so c_t can be removed from agent 1's bundle to eliminate their envy towards agent 2. Also, by the inductive assumption, there exists a chore $c \in A_2^{t-1} \setminus A_2^{r_t}$ such that $v_2((A_2^{t-1} \setminus A_2^{r_t}) \setminus \{c\}) \geq v_2(A_1^{t-1} \setminus A_1^{r_t})$. Observe that $A_2^t \setminus A_2^{r_t} = A_2^{t-1} \setminus A_2^{r_t}$ and $A_1^t \setminus A_1^{r_t} = (A_1^{t-1} \setminus A_1^{r_t}) \cup \{c_t\}$. Combining this with the inductive assumption, we have that there exists a chore $c \in A_2^t \setminus A_2^{r_t}$ such that

$$\begin{aligned} v_2((A_2^t \setminus A_2^{r_t}) \setminus \{c\}) &= v_2((A_2^{t-1} \setminus A_2^{r_t}) \setminus \{c\}) \\ &\geq v_2(A_1^{t-1} \setminus A_1^{r_t}) \\ &\geq v_2((A_1^{t-1} \setminus A_1^{r_t}) \cup \{c_t\}) \\ &= v_2(A_1^t \setminus A_1^{r_t}). \end{aligned}$$

Therefore, $\mathcal{A}^t \setminus \mathcal{A}^{r_t}$ is EF1 in this case.

We have shown that $\mathcal{A}^t \setminus \mathcal{A}^{r_t}$ is EF1 regardless of whether the allocation has undergone a bundle swap, so by induction, Algorithm 1 returns a TEF1 allocation for chores. □

Next, we consider *temporal envy-freeness up to any item (TEFX)*, the temporal variant of the stronger notion of *envy-freeness up to any item (EFX)*.

Definition 3.3. In a goods (resp., chores) allocation instance, an allocation $\mathcal{A} = (A_1, \ldots, A_n)$ is said to be EFX if for all pairs of agents $i,j \in N$, and all goods $g \in A_j$ (resp., chores $c \in A_i$) we have $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ (resp. $v_i(A_i \setminus \{c\}) \geq v_i(A_j)$).

Definition 3.4 (Temporal EFX). For every $t \in [T]$, an allocation $\mathcal{A}^t = (A_1^t, \ldots, A_n^t)$ is said to be *temporal envy-free up to any item (TEFX)* if for each $\ell \leq t$ the allocation $\mathcal{A}^\ell$ is EFX.

Unfortunately, TEFX allocations (for goods or chores) may not exist, even for two agents with identical valuations, and even when there are only two types of items.

Proposition3.5. *A* TEFX *allocation for goods or chores may not exist, even for $n = 2$ with identical valuations and two types of items.*

3.2 Other Restricted Settings

The next natural question we ask is whether there are other special cases where EF1 allocation is guaranteed to exist. We answer this question affirmatively by demonstrating the existence of EF1 allocations in three special cases, each supported by a polynomial-time algorithm that returns such an allocation.

Two Types of Items. Thefi rst setting we consider is one where items can be divided into two *types*, and each agent values all items of a particular type equally. Formally, let $S_1, S_2 \subseteq O$ be a partition of the set of items, so that $S_1 \cap S_2 = \varnothing$, and $S_1 \cup S_2 = O$. Then, for any $r \in \{1, 2\}$, two items $o, o' \in S_r$, and agent $i \in N$, we have that $v_i(o) = v_i(o')$.

Settings with only two types of items/tasks arise naturally in various applications, such as distributing food and clothing donations from a charity, or allocating cleaning and cooking chores in a household. This preference restriction has been studied for chores in offline settings [9, 43], and we remark that agents may have distinct valuations for up to $2n$ different items, unlike the extensively studied *bi-valued* preferences [33, 42] which involve only two distinct item values.

We show that for this setting, a TEF1 allocation for goods or chores always exists and can be computed in polynomial time. Intuitively, the algorithm treats the two item types independently: items of thefi rst type are allocated in a round-robin manner from agent 1 to n, while items of the second type are allocated in reverse round-robin order from agent n to 1. Then, our result is as follows.

Theorem3.6. *When there are two types of items, a* TEF1 *allocation for goods or chores exists and can be computed in polynomial time.*

Generalized Binary Valuations. The next setting we consider is one where agents have *generalized binary valuations* (also known as *restricted additive valuations* [1, 24]). This class of valuation functions generalizes both identical and binary valuations, which are both widely studied in fair division [46, 62, 66]. Formally, we say that agents have *generalized binary valuations* if for every agent $i \in N$ and item $o_j \in O$, $v_i(o_j) \in \{0, p_j\}$, where $p_j \in \mathbb{R} \setminus \{0\}$.

We show that for this setting, a TEF1 allocation can be computed efficiently, with the following result. We remark that the resulting allocation also satisfies *Pareto-optimality* (Definition 4.1).

Theorem3.7. *When agents have generalized binary valuations, a* TEF1 *allocation for goods or chores exists and can be computed in polynomial time.*

Unimodal Preferences. The last setting that we consider is the class of *unimodal preferences*, which consists of the widely studied *single-peaked* and *single-dipped* preference structures in social choice [7, 22] and cake cutting [21, 68]. We adapt these concepts for the online fair division setting with a single item at each timestep.

Definition 3.8. A valuation profile $\mathbf{v}$ is *single-peaked* if for each agent $i \in N$, there is an item o_{i^*} where for each $j, k \in [m]$ such that $j < k < i^*$, $v_i(o_j) \leq v_i(o_k) \leq v_i(o_{i^*})$, and for each $j, k \in [m]$ such that $i^* < j < k$, $v_i(o_{i^*}) \geq v_i(o_j) \geq v_i(o_k)$.

Definition 3.9. A valuation profile $\mathbf{v}$ is *single-dipped* if for each agent $i \in N$, there is an item o_{i^*} where for each $j, k \in [m]$ such that $j < k < i^*$, $v_i(o_j) \geq v_i(o_k) \geq v_i(o_{i^*})$, and for each $j, k \in [m]$ such that $i^* < j < k$, $v_i(o_{i^*}) \leq v_i(o_j) \leq v_i(o_k)$.

In other words, under single-peaked (resp. single-dipped) valuations, agents have a specific item o_{i^*} that they prefer (resp. dislike) the most, and prefer (resp. dislike) items less as they arrive further away in time from o_{i^*}.

Note that this restricted preference structure is well-defined for the setting of a single item arriving per round, but may not be compatible with a generalization to multiple items per round as described in Lemma 3.1 (unless the items in each round are identically-valued by agents).[3]

Unimodal preferences may arise in settings where agents place higher value on resources at the time surrounding specific events. For example, in disaster relief, the demand for food and essential supplies peaks as a natural disaster approaches, then declines once the immediate crisis passes. Similarly, in project management, the workload for team members intensifies (in terms of required time and effort) as the project nears its deadline, but significantly decreases during thefi nal stages, such as editing and proofreading.

[3]Specifically, in the multiple items per round case, if the bundles of items at each timestep are unimodally valued, the single-item per round transformation of the instance may not necessarily be unimodal.

Unimodal preferences also generalizes other standard preference restrictions studied in fair division and voting models, such as settings where agents have *monotonic valuations* [40] or *identical rankings* [62].

We propose efficient algorithms for computing a TEF1 allocation for goods when agents have single-peaked valuations, and for chores when agents have single-dipped valuations.

Theorem3.10. *When agents have single-peaked valuations, a* TEF1 *allocation for goods exists and can be computed in polynomial time. When agents have single-dipped valuations, a* TEF1 *allocation for chores exists and can be computed in polynomial time.*

We note that while a simple greedy algorithm performs well in the case of single-peaked valuations for goods and single-dipped valuations for chores, it fails in the reverse scenario—single-dipped valuations for goods and single-peaked valuations for chores. This is due to the fact that, in the latter case, the position of the dip or peak becomes critical and significantly complicates the way we allocate the item. We leave the existence of polynomial-time algorithm(s) for the reverse scenario as an open question.

3.3 Hardness Results for TEF1 Allocations

The non-existence of TEF1 goods allocations for $n \geq 3$ prompts us to explore whether we can determine if a given instance admits a TEF1 allocation for goods. Unfortunately, we show that this problem is NP-hard, with the following result.

Theorem3.11. *Given an instance of the temporal fair division problem with goods and $n \geq 3$, determining whether there exists a* TEF1 *allocation is* NP*-hard.*

Proof. We reduce from the 1-in-3-SAT problem which is NP-hard. An instance of this problem consists of a conjunctive normal form formula F with three literals per clause; it is a yes instance if there exists a truth assignment to the variables such that each clause has exactly one `True` literal, and a no instance otherwise.

Consider an instance of 1-in-3-SAT given by the CNF F which contains n variables $\{x_1, \ldots, x_n\}$ and m clauses $\{C_1, \ldots, C_m\}$. We construct an instance $\mathcal{I}$ with three agents and $2n + 2$ goods. For each $i \in [n]$, we introduce two goods t_i, f_i. We also introduce two additional goods s and r. Let the agents' (identical) valuations be defined as follows:

$$v(g) = \begin{cases} 5^{m+n-i} + \sum_{j : x_i \in C_j} 5^{m-j}, & \text{if } g = t_i, \\ 5^{m+n-i} + \sum_{j : \neg x_i \in C_j} 5^{m-j}, & \text{if } g = f_i, \\ \sum_{j \in [m]} 5^{j-1}, & \text{if } g = r, \\ \sum_{i \in [n]} 5^{m+i-1} + 2 \times \sum_{j \in [m]} 5^{j-1}, & \text{if } g = s. \end{cases}$$

Intuitively, for each variable index $i \in [n]$, we associate with it a unique value 5^{m+n-i}. For each clause index $j \in [m]$, we also associate with it a unique value 5^{m-j}. Note that no two indices (regardless of whether its a variable or clause index) share the same value. Then, the value of each good t_i comprises of the unique value associated with i, and the sum over all unique values of clauses C_j which x_i appears as a *positive literal* in; whereas the value of each good f_i comprises of the unique value associated with i, and the sum over all unique values of clauses C_j which x_i appears as a *negative literal* in. We will utilize this in our analysis later.

Then, we have the set of goods $O = \{s, t_1, f_1, t_2, f_2, \dots, t_n, f_n, r\}$. Note that $v(O) = v(s) + v(r) + \sum_{i\in[n]} v(t_i) + \sum_{i\in[n]} v(f_i)$. Also observe that $\sum_{i\in[n]} 5^{m+n-i} = \sum_{i\in[n]} 5^{m+i-1}$. Now, as each clause contains exactly three literals, we have

$$\sum_{i\in[n]} \sum_{j: x_i \in C_j} 5^{m-j} + \sum_{i\in[n]} \sum_{j:\neg x_i \in C_j} 5^{m-j} = 3 \times \sum_{j\in[m]} 5^{j-1}.$$

Then, combining the equations above, we get that

$$v(O) = 3 \times \sum_{i\in[n]} 5^{m+i-1} + 6 \times \sum_{j\in[m]} 5^{j-1}. \tag{1}$$

Let the goods be in the following order:

$$s, t_1, f_1, t_2, f_2, \dots, t_n, f_n, r.$$

Wefi rst prove the following result.

Lemma3.12. *There exists a truth assignment α such that each clause in F has exactly one* True *literal if and only if there exists an allocation $\mathcal{A}$ such that $v(A_1) = v(A_2) = v(A_3)$ for instance $\mathcal{I}$.*

Proof. For the 'if' direction, consider an allocation $\mathcal{A}$ such that $v(A_1) = v(A_2) = v(A_3)$. Then, we have that $O = A_1 \cup A_2 \cup A_3$ and $v(A_1) = v(A_2) = v(A_3) = \frac{1}{3}v(O)$. Since agents have identical valuations, without loss of generality, let $s \in A_1$. Then, since $v(A_1^1) = v(s) = \frac{1}{3}v(O)$, agent 1 should not receive any more goods after s, and each remaining good should go to agent 2 or 3.

Again, without loss of generality, we let $r \in A_2$. Then since $v(A_2) = \frac{1}{3}v(O)$, we have that

$$\begin{aligned} v(A_2 \setminus \{r\}) &= \left(\sum_{i\in[n]} 5^{m+i-1} + 2 \times \sum_{j\in[m]} 5^{j-1} \right) - \sum_{j\in[m]} 5^{j-1} \\ &= \sum_{i\in[n]} 5^{m+i-1} + \sum_{j\in[m]} 5^{j-1}. \end{aligned}$$

Note that this is only possible if for each $i \in [m]$, t_i and f_i are allocated to different agents. The reason is because the only way agent 1 can obtain thefi rst term of the above bundle value (less good r) is if he is allocated exactly one good from each of $\{t_i, f_i\}$ for all $i \in [n]$.

Then, from the goods that exist in bundle A_2, we can construct an assignment α: for each $i \in [n]$, let $x_i =$ True if $t_i \in A_2$ and $x_i =$ False if $f_i \in A_2$. Then, from the second term in the expression of $v(A_1 \setminus \{r\})$ above, we can observe that each clause has exactly one True literal (because the sum is only obtainable if exactly one literal appears in each clause, and our assignment will set each of these literals to True).

For the 'only if' direction, consider a truth assignment α such that each clause in F has exactly one True literal. Then, for each $i \in [n]$, let

$$\ell_i = \begin{cases} t_i & \text{if } x_i = \texttt{True} \text{ under } \alpha, \\ f_i & \text{if } x_i = \texttt{False} \text{ under } \alpha. \end{cases}$$

We construct the allocation $\mathcal{A} = (A_1, A_2, A_3)$ where

$$A_1 = \{s\}, \quad A_2 = \{\ell_1, \dots, \ell_n, r\}, \quad \text{and} \quad A_3 = O \setminus (A_1 \cup A_2).$$

Again, observe that $\sum_{i\in[n]} 5^{m+n-i} = \sum_{i\in[n]} 5^{m+i-1}$. Also note that $v(A_1) = \frac{1}{3}v(O)$. Then, as each clause has exactly one True literal, $v(A_2) = \sum_{i\in[n]} 5^{m+i-1} + 2 \times \sum_{j\in[m]} 5^{j-1}$, and together with (1), we get that $v(A_3) = \frac{2}{3}v(O) - v(A_1) = v(A_1)$ and hence $v(A_1) = v(A_2) = v(A_3)$, as desired. □

Now consider another instance $\mathcal{I}'$ that is similar to $\mathcal{I}$, but with an additional 21 goods $\{g_1, \dots, g_{21}\}$. Let agents' valuations over these new goods be defined as follows:

v	g_1	g_2	g_3	g_4	g_5	g_6	g_7
1	90	80	70	100	100	100	15
2	90	70	80	100	100	100	95
3	80	90	70	100	100	100	25
	g_8	g_9	g_{10}	g_{11}	g_{12}	g_{13}	g_{14}
1	10000	11000	12000	20000	20000	20000	20000
2	10000	11000	12000	20000	20000	20000	20000
3	10000	11000	12000	20000	20000	18500	20000
	g_{15}	g_{16}	g_{17}	g_{18}	g_{19}	g_{20}	g_{21}
1	20000	20000	20000	20000	20000	19010	18005
2	20000	20000	20000	12000	12000	19085	14106
3	20000	20000	20000	20000	20000	19010	19496

Then, we have the set of goods $O' = O \cup \{g_1, \dots, g_{21}\}$.

Let the goods be in the following order:

$$s, t_1, f_1, t_2, f_2, \dots, t_n, f_n, r, g_1, \dots, g_{21}.$$

We now present thefi nal lemma that will give us our result.

Lemma3.13. *If there exists a partial allocation $\mathcal{A}^{2n+2}$ over thefirst $2n+2$ goods such that $v(A_1^{2n+2}) = v(A_2^{2n+2})$, then there exists a* TEF1 *allocation $\mathcal{A}$. Conversely, if there does not exist a partial allocation $\mathcal{A}^{2n+2}$ over thefirst $2n+2$ goods such that $v(A_1^{2n+2}) = v(A_2^{2n+2})$, then there does not exists a* TEF1 *allocation $\mathcal{A}$.*

We use a program as a gadget to verify the lemma (see the full version of the paper), leveraging its output to support its correctness. Specifically, if there exists a partial allocation $\mathcal{A}^{2n+2}$ over thefirst $2n+2$ goods such that $v(A_1^{2n+2}) = v(A_2^{2n+2})$, then our program will show the existence of a TEF1 allocation by returning all such TEF1 allocations. If there does not exist such a partial allocation, our program essentially does an exhaustive search to show that a TEF1 allocation does not exist. This lemma shows that there exists a TEF1 allocation over O' if and only if $v(A_1^{2n+2}) \neq v(A_2^{2n+2})$, and by Claim 3.12, this implies that a TEF1 allocation over O' exists if and only if there is a truth assignment α such that each clause in F has exactly one True literal. □

However, we note that the above approach cannot be extended to show hardness for the setting with chores. Nevertheless, we are able to show a similar, though weaker, intractability result for the case of chores in general. The key difference is that we assume that we can start from any partial TEF1 allocation.

Theorem3.14. *For every $t \in [T]$, given any partial* TEF1 *allocation $\mathcal{A}^t$ for chores, deciding if there exists an allocation $\mathcal{A}$ that is* TEF1 *is* NP*-hard.*

Proof. We reduce from the NP-hard problem Partition. An instance of this problem consists of a multiset S of positive integers; it is a yes-instance if S can be partitioned into two subsets S_1 and S_2 such that the sum of the numbers in S_1 equals the sum of the numbers in S_2, and a no-instance otherwise.

Consider an instance of PARTITION given by a multiset set $S = \{s_1, \dots, s_m\}$ of m positive integers. Then, we construct a set $S' = \{s'_1, \dots s'_m\}$ such that for each $j \in [m]$, $s'_j = s_m - K$ where $K := \max\{s_1, \dots, s_m\} + \varepsilon$ for some small $\varepsilon > 0$. We then scale members of S' such that they sum to -2, i.e., $\sum_{s' \in S'} s' = -2$.

Next, we construct an instance with four agents and $m+4$ chores $O = \{b_1, b_2, b_3, b_4, c_1, \dots, c_m\}$, where agents have the following valuation profile $\mathbf{v}$ for $j \in \{1, \dots, m\}$:

$\mathbf{v}$	b_1	b_2	b_3	b_4	c_1	$\dots$	c_j	$\dots$	c_m
1	(−1)	0	0	0	−1	…	−1	…	−1
2	−1	(−1)	−1	−1	s'_1	…	s'_j	…	s'_m
3	−1	−1	(−1)	−1	s'_1	…	s'_j	…	s'_m
4	0	0	0	(−1)	−1	…	−1	…	−1

Also, suppose we are given the partial allocation $\mathcal{A}^4$ where for each $i \in \{1, 2, 3, 4\}$, chore b_i is allocated to agent i, as illustrated in the table above. Note that the partial allocation $\mathcal{A}^4$ is TEF1.

Wefi rst establish the following two lemmas. Thefi rst lemma states that after chores b_1, b_2, b_3, b_4 are allocated, in order to maintain TEF1, each remaining chore in $\{c_1, \dots, c_m\}$ cannot be allocated to either agent 1 or agent 4. The result is as follows.

LEMMA3.15. *In any* TEF1 *allocation, agents* 1 *and* 4 *cannot be allocated any chore in* $\{c_1, \dots, c_m\}$.

The second lemma states that in any TEF1 allocation, the sum of values that agents 2 and 3 obtain from the chores in $\{c_1, \dots, c_m\}$ that are allocated to them must be equal. We formalize it as follows.

LEMMA3.16. *In any* TEF1 *allocation, let* C_2, C_3 *be the subsets of* $\{c_1, \dots, c_m\}$ *that were allocated to agents* 2 *and* 3 *respectively. Then,* $v_2(C_2) = v_3(C_3)$.

We will now prove that there exists an allocation $\mathcal{A}$ satisfying TEF1 if and only if the set S can be partitioned into two subsets of equal sum.

For the 'if' direction, suppose $S = \{s_1, \dots, s_m\}$ can be partitioned into two subsets S_1, S_2 of equal sum. This means that $S' = \{s'_1, \dots, s'_m\}$ can be correspondingly partitioned into two subsets S'_1, S'_2 of equal sum (of -1 each). Let C_1, C_2 be the partition of chores in $\{c_1, \dots, c_m\}$ with values corresponding to the partitions S'_1, S'_2 respectively. Then we allocate all chores in C_1 to agent 2 and all chores in C_2 to agent 3. By Lemma 3.15, we have that agents 1 and 4 cannot envy any other agent at any round. Also, for any round $t \in [T]$ and $i,j \in \{2, 3\}$ where $i \neq j$, $v_i(A^t_i \setminus \{b_i\}) \geq -1 \geq v_i(A^t_j)$, and for all $i \in \{2, 3\}$ and $k \in \{1, 4\}$, $v_i(A^t_i \setminus \{b_i\}) \geq -1 = v_i(A^t_k)$. Thus, the allocation $\mathcal{A}$ that, for each $i \in \{1, 2, 3, 4\}$, allocates b_i to agent i and for each $j \in \{2, 3\}$, allocates C_j to agent j, is TEF1.

For the 'only if' direction, suppose we have an allocation $\mathcal{A}$ satisfying TEF1. By Lemma 3.15, it must be that any chore in $\{c_1, \dots, c_m\}$ is allocated to either agent 2 or 3. Let C_2, C_3 be the subsets of chores in $\{c_1, \dots, c_m\}$ that are allocated to agents 2 and 3 respectively, under $\mathcal{A}$. Then, by Lemma 3.16, we have that $v_2(C_2) = v_3(C_3)$. By replacing the chores with their corresponding values, we get a partition of S' into two subsets of equal sums, which in turn gives us a partition of S into two subsets of equal sum. □

4 COMPATIBILITY OF TEF1 AND EFFICIENCY

In traditional fair division, many papers have focused on the existence and computation of fair and efficient allocations for goods or chores, with a particular emphasis on simultaneously achieving EF1 and *Pareto-optimality (PO)* [16, 25]. In this section, we explore the compatibility between TEF1 and PO. We begin by defining PO as follows.

Definition 4.1 (Pareto-optimality). We say that an allocation $\mathcal{A}$ is *Pareto-optimal (PO)* if there does not exist another allocation $\mathcal{A}'$ such that for all $i \in N$, $v_i(A'_i) \geq v_i(A_i)$, and for some $j \in N$, $v_j(A'_j) > v_j(A_j)$. If such an allocation $\mathcal{A}'$ exists, we say that $\mathcal{A}'$ *Pareto-dominates* $\mathcal{A}$.

Observe that for any $\mathcal{A}$ that is PO, any partial allocation $\mathcal{A}^t$ for $t \leq [T]$ is necessarily PO as well. We demonstrate that PO is incompatible with TEF1 in this setting, even under very strong assumptions (of two agents and two types of items), as illustrated by the following result.

PROPOSITION4.2. *For any* $n \geq 2$*, a* TEF1 *and* PO *allocation for goods or chores may not exist, even when there are two types of items.*

Despite this non-existence result, one may still wish to obtain a TEF1 and PO outcome when the instance admits one. However, the following results show that this is not computationally tractable.

THEOREM4.3. *Determining whether there exists a* TEF1 *allocation that is* PO *for goods is* NP*-hard, even when* $n = 2$.

THEOREM4.4. *Determining whether there exists a* TEF1 *allocation that is* PO *for chores is* NP*-hard, even when* $n = 2$.

The proof of the above result essentially imply that even determining whether an instance admits a TEF1 and *utilitarian-maximizing* (i.e., sum of agents' utilities) allocation is computationally intractable, since a utilitarian-welfare maximizing allocation is necessarily PO. In fact, for the case of goods, we can make a stronger statement relating to the general class of *p-mean welfares*, defined as follows.[4]

Definition 4.5. Given $p \in (-\infty, 1]$ and an allocation $\mathcal{A} = (A_1, \dots, A_n)$ of goods, the *p-mean welfare* is $\left(\frac{1}{n} \sum_{i \in N} v_i(A_i)^p\right)^{1/p}$.

In the context of fair division, p-mean welfare has been traditionally and well-studied for the setting with goods [13, 28], although it has recently been explored for chores as well [34]. Importantly, p-means welfare captures a spectrum of commonly studied fairness objectives in fair division. For instance, setting $p = 1$ (resp. $p = -\infty$) would correspond to the utilitarian (resp. egalitarian) welfare. Setting $p \to 0$ corresponds to maximizing the geometric mean, which is also known as the Nash welfare [25].

Then, from our construction in the proof of Theorem 4.3 (for goods), we have that an allocation is TEF1 and PO if and only if it also maximizes the p-mean welfare, for all $p \in (-\infty, 1]$, thereby giving us the following corollary.

COROLLARY4.6. *For all* $p \in (-\infty, 1]$*, determining whether there exists a* TEF1 *allocation that maximizes* p*-mean welfare is* NP*-hard, even when* $n = 2$.

[4] Note that we cannot say the same for chores as when agents' valuations are negative, the p-mean welfare may be ill-defined.

5 MULTIPLE ITEMS PER ROUND

We now revisit the setting where multiple items may arrive at each round. While Lemma 3.1 reduces this case to the setting where a single item arrives per round, there are restricted variants of our problem that are not preserved by this reduction. We will now consider two such variants: $T = 2$ and repeated allocation.

We begin by showing that when there are two rounds, a TEF1 allocation can be computed efficiently.

Theorem5.1. *When $T = 2$, a* TEF1 *allocation for goods or chores exists and can be computed in polynomial time.*

For the remainder of Section 5, we consider the *repeated* setting (also studied by Igarashi et al. [48] and Caragiannis and Narang [26]), where the sets $O_1, \dots, O_T$ are identical. Formally, for each $t \in T$ we have $O_t = \{o^t_1, \dots, o^t_k\}$, and $v_i(o^t_j) = v_i(o^r_j)$ for all $t, r \in [T]$ and all $i \in N$, $j \in [k]$. Note that this property of the instance is not preserved by our reduction from many items per round to a singe item per round.

In general, it remains an open question whether a TEF1 allocation exists for this setting. However, we can show that, perhaps surprisingly, it is NP-hard to determine whether there exists a TEF1 allocation that allocates the items in the *same way* at *every round.* We say that an allocation $\mathcal{A}$ is *repetitive* if for each $i \in N$, $j \in [k]$ and all $t, r \in [T]$ we have $o^t_j \in A^t_i \setminus A^{t-1}_i$ if and only if $o^r_j \in A^r_i \setminus A^{r-1}_i$. Then we have the following result.

Theorem5.2. *Determining whether there exists a repetitive allocation $\mathcal{A} = (A_1, \dots, A_n)$ that is* TEF1 *is* NP*-complete both for goods and for chores. The hardness result holds even if $T = 2$ and agents have identical valuations.*

Proof. It is immediate that this problem is in NP: we can guess a repetitive allocation, and check whether it is TEF1. Both for goods and for chores, we reduce from the NP-hard problem Multiway Number Partitioning [45]. An instance of this problem is given by a positive integer κ and a multiset $S = \{s_1, \dots, s_\mu\}$ of μ non-negative integers whose sum is κW; it is a yes-instance if S can be partitioned into κ subsets such that the sum of integers in each subset is W, and a no-instance otherwise.

Consider an instance of MultiwayN umber Partitioning given by a positive integer κ and a multiset $S = \{s_1, \dots, s_\mu\}$ of μ non-negative integers that sum up to κW.

Wefi rst prove the result for goods. We construct an instance with $\kappa+1$ agents and $\mu+1$ goods in each round: $O_1 = \{g^1_1, \dots, g^1_{\mu+1}\}$ and $O_2 = \{g^2_1, \dots, g^2_{\mu+1}\}$. The agents have an identical valuation function v defined as follows: $v(g^1_j) = v(g^2_j) = s_j$ if $j \in [\mu]$, and $v(g^1_{\mu+1}) = v(g^2_{\mu+1}) = 2W$. We will now prove that there exists a repetitive TEF1 allocation $\mathcal{A}$ if and only if the set S can be partitioned into κ subsets with equal sums (of W each).

For the 'if' direction, consider a κ-way partition $\mathcal{P} = \{P_1, \dots, P_\kappa\}$ of S with $\sum_{s \in P_i} s = W$ for each $i \in [\kappa]$. We construct allocations $\mathcal{A}^1$ and $\mathcal{A}^2$ by allocating the goods corresponding to the elements of subset P_i to agent i for $i \in [\kappa]$; the goods $g^1_{\mu+1}$ and $g^2_{\mu+1}$ are allocated to agent $\kappa + 1$. Then, in $\mathcal{A}^1$, for each agent $i \in [\kappa]$ we have $v(A^1_i) = \sum_{s \in P_i} s = W$, and $v(A^1_{\kappa+1}) = v(g^1_{\mu+1}) = 2W$. It is easy to verify that $\mathcal{A}^1$ is EF1: no agent $i \in [\kappa]$ envies another agent $j \in [\kappa] \setminus \{i\}$, as they have the same bundle value, and agent i's envy towards agent $\kappa + 1$ can be removed by simply dropping $g^1_{\mu+1}$ from $A^1_{\kappa+1}$. Also, agent $\kappa + 1$ does not envy thefirst κ agents: she values her bundle at $2W$ and the bundles of $i \in [\kappa]$ at W.

Moreover in $\mathcal{A}^2$ each agent $i \in [\kappa]$ values the bundles $A^2_1, \dots, A^2_\kappa$ at $2W$ and hence does not envy any of thefirst κ agents; her envy towards $\kappa + 1$ can be eliminated by dropping $g^1_{\mu+1}$ from $A^2_{\kappa+1}$. On the other hand, agent $\kappa + 1$ values her bundle at $4W$ and all other bundles at $2W$, so she does not envy thefirst κ agents.

For the 'only if' direction, suppose we have a repetitive allocation $\mathcal{A}^2$ that satisfies TEF1. Since agents have identical valuation functions, we can assume without loss of generality that agent $\kappa+1$ receives goods $g^1_{\mu+1}$ and $g^2_{\mu+1}$ in rounds 1 and 2. Then for agent $i \in [\kappa]$ not to envy $\kappa+1$ in $\mathcal{A}^2$ after we drop one item from A^2_{k+1}, it has to be the case that $v_i(A^2_i) \geq 2W$. As this holds for all $i \in [\kappa]$ and $\sum_{j \in [\mu]} s_j = \kappa W$, this is only possible if there is a κ-way partition of S such that each subset sums up to W.

The proof for chores is similar, and can be found in the full version of the paper. ◻

6 CONCLUSION

In this work, we studied the informed online fair division of indivisible items, with the goal of achieving TEF1 allocations. For both goods and chores, we showed the existence of TEF1 allocations in four special cases and provided polynomial-time algorithms for each case. Additionally, we showed that determining whether a TEF1 allocation exists for goods is NP-hard, and presented a similar, though slightly weaker, intractability result for chores. We further established the incompatibility between TEF1 and PO, which extends to an incompatibility with p-mean welfare. Finally, we explored the special case of multiple items arriving at each round.

Numerous potential directions remain for future work, including revisiting variants of the standard fair division model. Examples include studying the existence (and polynomial-time computability) of allocations satisfying a temporal variant of the weaker *proportionality up to one item* property (as defined by Conitzer et al. [30]), which would be implied by EF1; studying group fairness considerations in the temporal setting [3, 10, 18, 31, 50, 63]; considering the more general class of *submodular* valuations [44, 56, 67, 69]; considering the *house allocation* model where each agent gets a single item [29, 41], which was partially explored by Micheel and Wilczynski [55]; or even looking at more general settings with additional size constraints [14, 15, 35]. Another promising direction is to examine the number of approximate TEF1 allocations that exist in order to identify additional special cases [60, 65]. It would also be interesting to extend our results, which hold for the cases of goods and chores separately, to the more general case of mixed manna(see, e.g., [8]). In fact, with an appropriate modification of the instance, we can extend Theorem 3.2 to show that a TEF1 allocation exists in the mixed manna setting when there are two agents (see the full version of the paper).

ACKNOWLEDGMENTS

Most of this work was done when Edith Elkind was at the University of Oxford. Edith Elkind and Mohamad Latifian were supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grants EP/X038548/1 and EP/Y003624/1. We would like to thank the anonymous reviewers.

REFERENCES

[1] Hannaneh Akrami, Rojin Rezvan, and Masoud Seddighin. 2022. An EF2X allocation protocol for restricted additive valuations. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 17–23.

[2] Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. 2015. Online Fair Division: Analysing a Food Bank Problem. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*. 2540–2546.

[3] Martin Aleksandrov and Toby Walsh. 2018. Group Envy Freeness and Group Pareto Efficiency in Fair Division with Indivisible Items. In *Proceedings of the 41st German Conference on AI (KI)*. 57–72.

[4] Martin Aleksandrov and Toby Walsh. 2020. Online fair division: A survey. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 13557–13562.

[5] Shiri Alouf-Heffetz, Laurent Bulteau, Edith Elkind, Nimrod Talmon, and Nicholas Teh. 2022. Better Collective Decisions via Uncertainty Reduction. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 24–30.

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

[7] Kenneth J. Arrow. 2012. *Social Choice and Individual Values.* Yale University Press.

[8] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. 2019. Fair Allocation of Indivisible Goods and Chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 53–59.

[9] Haris Aziz, Jeremy Lindsay, Angus Ritossa, and Mashbat Suzuki. 2023. Fair Allocation of Two Types of Chores. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 143–151.

[10] Haris Aziz and Simon Rey. 2020. Almost Group Envy-free Allocation of Indivisible Goods and Chores. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*. 39–45.

[11] Siddhartha Banerjee, Vasilis Gkatzelis, Artur Gorokh, and Billy Jin. 2022. Online Nash Social Welfare Maximization with Predictions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1–19.

[12] Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha, and Nisarg Shah. 2023. Proportionally Fair Online Allocation of Public Goods with Predictions. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*. 20–28.

[13] Siddharth Barman, Umang Bhaskar, Anand Krishna, and Ranjani G. Sundaram. 2020. Tight Approximation Algorithms for p-Mean Welfare Under Subadditive Valuations. In *Proceedings of the 28th Annual European Symposium on Algorithms (ESA)*. 11:1–11:17.

[14] Siddharth Barman, Arindam Khan, Sudarshan Shyam, and K.V.N. Sreenivas. 2023. Finding Fair Allocations under Budget Constraints. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*. 5481–5489.

[15] Siddharth Barman, Arindam Khan, Sudarshan Shyam, and K. V. N. Sreenivas. 2023. Guaranteeing Envy-Freeness under Generalized Assignment Constraints. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 242–269.

[16] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018. Finding Fair and Efficient Allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. 557–574.

[17] Sanjoy K Baruah. 1995. Fairness in periodic real-time scheduling. In *Proceedings 16th IEEE Real-Time Systems Symposium*. IEEE, 200–209.

[18] Nawal Benabbou, Mithun Chakraborty, Edith Elkind, and Yair Zick. 2019. Fairness Towards Groups of Agents in the Allocation of Indivisible Items. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 95–101.

[19] Gerdus Benadè, Aleksandr M Kazachkov, Ariel D Procaccia, Alexandros Psomas, and David Zeng. 2024. Fair and efficient online allocations. *Operations Research* 72, 4 (2024), 1438–1452.

[20] Ziyad Benomar and Vianney Perchet. 2024. Non-clairvoyant Scheduling with Partial Predictions. *arXiv preprint arXiv:2405.01013* (2024).

[21] Bhavook Bhardwaj, Rajnish Kumar, and Josue Ortega. 2020. Fairness and efficiency in cake-cutting with single-peaked preferences. *Economics Letters* 190 (2020), 109064.

[22] Duncan Black. 1948. On the Rationale of Group Decision-making. *Journal of Political Economy* 56, 1 (1948), 23–34.

[23] Steven J. Brams and Alan D. Taylor. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution.* Cambridge University Press.

[24] Franklin Camacho, Rigoberto Fonseca-Delgado, Ramón Pino Pérez, and Guido Tapia. 2023. Generalized binary utility functions and fair allocations. *Mathematical Social Sciences* 121 (2023), 50–60.

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

[26] Ioannis Caragiannis and Shivika Narang. 2024. Repeatedly matching items to agents fairly and efficiently. *Theoretical Computer Science* 981 (2024), 114246.

[27] Nikhil Chandak, Shashwat Goel, and Dominik Peters. 2024. Proportional Aggregation of Preferences for Sequential Decision Making. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*. 9573–9581.

[28] Bhaskar Ray Chaudhury, Jugal Garg, and Ruta Mehta. 2021. Fair and Efficient Allocations under Subadditive Valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5269–5276.

[29] Davin Choo, Yan Hao Ling, Warut Suksompong, Nicholas Teh, and Jian Zhang. 2024. Envy-Free House Allocation with Minimum Subsidy. *Operations Research Letters* 54 (2024), 107103.

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

[31] Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Vaughan. 2019. Group Fairness for the Allocation of Indivisible Goods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 1853–1860.

[32] Benjamin Cookson, Soroush Ebadian, and Nisarg Shah. 2024. Temporal Fair Division. *arXiv preprint arXiv:2410.23416* (2024).

[33] Soroush Ebadian, Dominik Peters, and Nisarg Shah. 2022. How to Fairly Allocate Easy and Difficult Chores. In *21st International Conference on Autonomous Agents and Multiagent Systems*.

[34] Owen Eckart, Alexandros Psomas, and Paritosh Verma. 2024. On the Fairness of Normalized p-Means for Allocating Goods and Chores. *arXiv preprint arXiv:2402.14996* (2024).

[35] Edith Elkind, Ayumi Igarashi, and Nicholas Teh. 2024. Fair Division of Chores with Budget Constraints. In *Proceedings of the 17th International Symposium on Algorithmic Game Theory (SAGT)*. 55–71.

[36] Edith Elkind, Sonja Kraiczy, and Nicholas Teh. 2022. Fairness in Temporal Slot Assignment. In *Proceedings of the 15th International Symposium on Algorithmic Game Theory (SAGT)*. 490–507.

[37] Edith Elkind, Alexander Lam, Mohamad Latifian, Tzeh Yuan Neoh, and Nicholas Teh. 2024. *Temporal Fair Division of Indivisible Items.* Technical Report. arXiv:2410.14593

[38] Edith Elkind, Tzeh Yuan Neoh, and Nicholas Teh. 2024. Temporal Elections: Welfare, Strategyproofness, and Proportionality. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*. 3292–3299.

[39] Edith Elkind, Svetlana Obraztsova, and Nicholas Teh. 2024. Temporal Fairness in Multiwinner Voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*. 22633–22640.

[40] Edith Elkind, Svetlana Obraztsova, and Nicholas Teh. 2024. Verifying Proportionality in Temporal Multiwinner Voting. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 2246–2248.

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

[42] Jugal Garg, Aniket Murhekar, and John Qin. 2022. Fair and efficient allocations of chores under bivalued preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 5043–5050.

[43] Jugal Garg, Aniket Murhekar, and John Qin. 2024. Weighted EF1 and PO Allocations with Few Types of Agents or Chores. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI)*. 2799–2806.

[44] Mohammad Ghodsi, MohammadTaghi HajiAghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2022. Fair allocation of indivisible goods: Beyond additive valuations. *Artificial Intelligence* 303 (2022), 103633.

[45] Ronald Graham. 1969. Bounds on Multiprocessing Timing Anomalies. *SIAM J. Appl. Math.* 17, 2 (1969), 416–429.

[46] Daniel Halpern, Ariel D. Procaccia, Alexandros Psomas, and Nisarg Shah. 2020. Fair Division with Binary Valuations: One Rule to Rule Them All. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*. 370–383.

[47] Jiafan He, Ariel Procaccia, Alexandros Psomas, and David Zeng. 2019. Achieving a Fairer Future by Changing the Past. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 343–349.

[48] Ayumi Igarashi, Martin Lackner, Oliviero Nardi, and Arianna Novaro. 2024. Repeated Fair Allocation of Indivisible Items. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*. 9781–9789.

[49] Sungjin Im and Benjamin Moseley. 2020. Fair scheduling via iterative quasi-uniform sampling. *SIAM J. Comput.* 49, 3 (2020), 658–680.

[50] Maria Kyropoulou, Warut Suksompong, and Alexandros Voudouris. 2019. Almost Envy-Freeness in Group Resource Allocation. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*. 400–406.

[51] Martin Lackner. 2020. Perpetual Voting: Fairness in Long-Term Decision Making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. 2103–2110.

[52] Bo Li, Minming Li, and Ruilong Zhang. 2021. Fair scheduling for time-dependent resources. *Advances in Neural Information Processing Systems* 34 (2021), 21744–21756.

[53] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*. 125–131.

[54] Andrew Mackenzie and Vilmos Komornik. 2023. Fairly Taking Turns. *Games and Economic Behavior* 142 (2023), 743–764.

[55] Karl Jochen Micheel and Anaëlle Wilczynski. 2024. Fairness in Repeated House Allocation. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2390–2392.

[56] Luisa Montanari, Ulrike Schmidt-Kraepelin, Warut Suksompong, and Nicholas Teh. 2024. Weighted Envy-Freeness for Submodular Valuations. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*. 9865–9873.

[57] Hervé Moulin. 2003. *Fair Division and Collective Welfare*. MIT Press.

[58] Tzeh Yuan Neoh and Nicholas Teh. 2024. Welfare Maximization in Perpetual Voting (Student Abstract). In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*. 23597–23599.

[59] Tzeh Yuan Neoh and Nicholas Teh. 2025. Strategic Manipulation in Temporal Voting with Undesirable Candidates (Student Abstract). In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*.

[60] Tzeh Yuan Neoh and Nicholas Teh. 2025. Understanding EFX Allocations: Counting and Variants. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*.

[61] Christos H. Papadimitriou. 2007. *Computational complexity*. Academic Internet Publ.

[62] Benjamin Plaut and Tim Roughgarden. 2020. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics* 34, 2 (2020), 1039–1068.

[63] Jonathan Scarlett, Nicholas Teh, and Yair Zick. 2023. For One and All: Individual and Group Fairness in the Allocation of Indivisible Goods. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2466–2468.

[64] Uwe Schwiegelshohn and Ramin Yahyapour. 2000. Fairness in parallel job scheduling. *Journal of Scheduling* 3, 5 (2000), 297–320.

[65] Warut Suksompong. 2020. On the number of almost envy-free allocations. *Discrete Applied Mathematics* 284 (2020), 606–610.

[66] Warut Suksompong and Nicholas Teh. 2022. On Maximum Weighted Nash Welfare for Binary Valuations. *Mathematical Social Sciences* 117 (2022), 101–108.

[67] Warut Suksompong and Nicholas Teh. 2023. Weighted Fair Division with Matroid-Rank Valuations: Monotonicity and Strategyproofness. *Mathematical Social Sciences* 126 (2023), 48–59.

[68] William Thomson. 1994. Resource-monotonic solutions to the problem of fair division when preferences are single-peaked. *Social Choice and Welfare* 11, 3 (1994), 205–223.

[69] Gilad Ben Uziahu and Uriel Feige. 2023. On Fair Allocation of Indivisible Goods to Submodular Agents. *arXiv preprint arXiv:2303.12444* (2023).

[70] Valentin Zech, Niclas Boehmer, Edith Elkind, and Nicholas Teh. 2024. Multiwinner Temporal Voting with Aversion to Change. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*.