

# Temporal Fair Division

Benjamin Cookson, Soroush Ebadian, Nisarg Shah

University of Toronto  
{bcookson,soroush,nisarg}@cs.toronto.edu

## Abstract

We study temporal fair division, whereby a set of agents are allocated a (possibly different) set of goods on each day for a period of days. We study this setting, as well as a number of its special cases formed by the restrictions to two agents, same goods on each day, identical preferences, or combinations thereof, and chart out the landscape of achieving two types of fairness guarantees simultaneously: fairness on each day (per day) and fairness over time (up to each day, or the weaker version, overall).

In the most general setting, we prove that there always exists an allocation that is stochastically-dominant envy-free up to one good (SD-EF1) per day and proportional up to one good (PROP1) overall, and when all the agents have identical preferences, we show that SD-EF1 per day and SD-EF1 overall can be guaranteed. For the case of two agents, we prove that SD-EF1 per day and EF1 up to each day can be guaranteed using an envy balancing technique. We provide counterexamples for other combinations that establish our results as among the best guarantees possible, but also leave open some tantalizing questions.

## 1 Introduction

How to divide a set of goods amongst a set of agents *fairly* has been an enigma for centuries. There has been remarkable progress on this question in the last decade (Amanatidis et al. 2022). In the most prominent model, there is a set of  $n$  agents  $N$ , each having an (additive) valuation over a set of goods  $M$ . The goal is to find an allocation  $A = (A_1, \dots, A_n)$  which partitions  $M$  into pairwise-disjoint bundles, one allocated to each agent  $i \in N$ .

This one-shot model fails to capture numerous real-world fair division scenarios in which goods are divided over time, e.g., food bank deliveries (Lee et al. 2019), resource allocation in data centers (Ghods et al. 2013), allocation of advertising slots (Mehta et al. 2007), nurse shift scheduling (Miller, Pierskalla, and Rath 1976), and organ transplants (Bertsimas, Farias, and Trichakis 2013). Compared to the one-shot setting, fair division of goods over time has received relatively little attention.

Inspired by this, there has been a flurry of recent works that consider *online fair division*, where agents or goods ar-

rive over time and the principal needs to make allocations in an online fashion in the absence of any information regarding future arrivals (Kash, Procaccia, and Shah 2014; Benadè et al. 2024). The limits of feasible fairness guarantees have been explored under various adversary models (Zeng and Psomas 2020; Benadè et al. 2024).

However, in practice it is rarely the case that we have absolutely no information about the future. Significantly better guarantees have been established when even partial information about the future is available, either in the form of distributional knowledge (Bogomolnaia, Moulin, and Sandomirskiy 2022) or machine-generated predictions (Gkatzelis, Psomas, and Tan 2021; Banerjee et al. 2022, 2023). But this work has left a very basic question wide open: *How fair can we be if we had full information about the future?*

To address this, we introduce the model of *temporal fair division*, where a set of agents  $N$  are allocated a set of goods  $M_t$  on day  $t$ , over a period of days  $t \in \{1, \dots, k\}$ , and the agents' valuations over the whole set of goods  $M = \cup_{t=1}^k M_t$  are available upfront. At first glance, it may seem that this is just a traditional fair division problem where the set of goods  $M$  needs to be divided amongst the set of agents  $N$ . The twist, however, is that in temporal fair division, agents anticipate fairness to prevail not solely at the end of the entire time horizon, but also at or within various interim time intervals. For example, the principal may be confident, based on their knowledge of the future, the allocation will eventually turn out to be fair, but that may not be assurance enough to the agents.

This leads us to seek *temporal fairness notions* in our temporal fair division setting. Specifically, we take prominent fairness notions from one-shot fair division, and seek them on three temporal scales:

- (1) *Per day*: The allocation of the set of goods  $M_t$  on each day  $t$  should be fair.
- (2a) *Overall*: The allocation of the whole set of goods  $M$  in the end should be fair.
- (2b) *Up to each day*: The allocation of the set of goods  $\cup_{r=1}^t M_r$  up to each day  $t$  should be fair.

Clearly, up to each day fairness (2b) is stronger than overall fairness (2a). Solely achieving per day fairness (1) or overall fairness (2a) can be reduced to one-shot fair division.

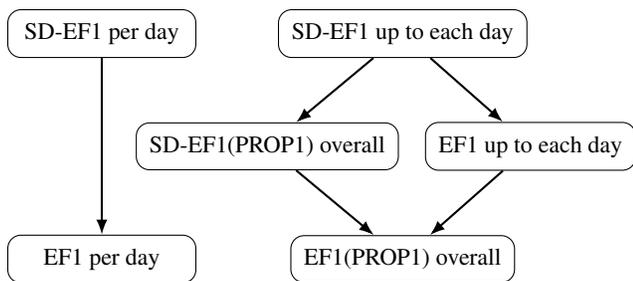


Figure 1: Hierarchy of temporal fairness notions.

Hence, we seek per day and overall fairness simultaneously (1+2a), or per day and up to each day fairness simultaneously (1+2b), or solely up to each day fairness (2b). Our main research question is to...

*...explore the limits of temporal fairness that can be guaranteed in temporal fair division.*

## Our Results & Techniques

We chart out the landscape of the aforementioned temporal fair division model in a general setting, with  $n$  agents having additive, heterogeneous preferences. Further, we identify three relevant restrictions where we can circumvent some of the impossibilities of the general setting, and achieve very strong results. Those are: (1) when there are only two agents; (2) when all agents have the same ordering over the goods; and (3) when an identical set of goods arrives each day.

In these settings, we seek the fairness guarantees of EF1, SD-EF1, PROP1, and SD-PROP1 at the temporal scales of per day, overall, and up to each day. The various temporal fairness definitions are depicted in Figure 1, with arrows indicating logical implications. We discover several surprising results, and develop novel algorithmic tools along the way, which may be of independent interest.

In Section 3, we present an algorithm for finding temporally fair allocations in our most general setting. Specifically, we show how to obtain an allocation that is PROP1 overall, and SD-EF1 per day in polynomial time (Theorem 1).

In Sections 4 to 6, we look at the restricted settings of two agents, identical orderings, and identical days, respectively. In these settings, we provide algorithms that give very strong fairness guarantees that are impossible in the general case (Theorems 3, 7 and 9), while also showing which fairness desiderata are still too strong even after applying these restrictions (Theorems 4, 5 and 8). In these restricted settings, we provide a near complete picture of what is possible, while leaving open some interesting questions, particularly surrounding the powerful notion of “up to each day” fairness.

Most of our results, along with several open questions, are summarized in Table 1.

## Related Work

Our temporal fair division model is related to (but separate from) several fair division models studied in the literature. We include the most relevant ones here, and include a more

in-depth look at the literature in the full version of this paper<sup>1</sup>.

**Repeated fair division.** The repeated fair division model of Igarashi et al. (2024), which is the case of identical days in our more general model, is the most closely related to our work. Some of their results are for two agents (still with identical days). As will be seen in Section 4, this is where the strongest guarantees from Figure 1 of SD-EF1 per day and SD-EF1 up to each day can be achieved simultaneously. This result is the only overlap between our works, and we present it again because we obtain a shorter proof with a much simpler algorithm. The rest of their results with two or more than two agents seek exact fairness guarantees, such as (exact) envy-freeness overall, in limited cases such as when the number of days is a multiple of the number of agents.

**Online fair division.** In the online fair division model, goods arrive one by one and must be irrevocably allocated to an agent upon arrival with no knowledge of agent preferences over the goods to arrive later. Typically, one seeks to maintain a certain level of fairness. Clearly, any online fair division algorithm can be simulated in our temporal fair division model to achieve the same guarantee up to each day. One online fair division paper of particular note to this research is He et al. (2019). This paper introduced the “Informed Model” of online fair division, where irrevocable allocation decisions must be made as goods arrive one at a time in adversarial order, but the allocation algorithm is given all goods and the order they will arrive in advance. The main goal in the informed model is to achieve an allocation that remains EF1 after each good is allocated. Clearly, this is equivalent to achieving EF1 up to each day in Temporal Fair Division when each day only contains a single good. He et al. concludes that it is impossible to allocate the goods in such a way that EF1 is always maintained. By corollary, EF1 up to each day is also infeasible.

**Constrained fair division.** Seeking fair allocations over a set of goods that remain fair when only looking at the goods from a single “day” can also be modeled as a constraint on the space of feasible allocations, and the question becomes whether there is a *constrained allocation* that still achieves the desired fairness guarantee. This model of constrained fair division has also been studied in the literature. Biswas and Barman (2018) study a model with cardinality constraints. Cardinality constraints (partition matroid constraints) have been generalized to matroid constraints, and the existence of an EF1 allocation subject to matroid feasibility constraints is a major open question (Biswas and Barman 2018; Dror, Feldman, and Segal-Halevi 2023). Finally, the bihierarchy framework of Budish et al. (2013) can also be viewed as a method for finding a constrained allocation, which we use in some of our results. Although, our most interesting results deal with sets of constraints that go beyond bihierarchies.

**Temporal fairness in social choice.** While we look at temporal fairness applied to the allocation of indivisible resources, the idea of temporal fairness has been explored in

<sup>1</sup>Full version: <https://arxiv.org/abs/2410.23416>

	Up to each day		Overall			
	SD-EF1	EF1	SD-EF1	SD-PROP1	EF1	PROP1
<b>General Setting</b>						
SD-EF1 Per Day	X	X	X	?	?	✓ (Thm 1)
EF1 Per Day	X	X	X	?	?	✓
$\emptyset$	X	X (He et al. 2019)	✓ (Aziz 2020)	✓	✓	✓
<b>Two Agents</b>						
SD-EF1 Per Day	X	✓ (Thm 3)	X	?	✓	✓
EF1 Per Day	X	✓	X (Thm 5)	?	✓	✓
$\emptyset$	X (Thm 4)	✓	✓	✓	✓	✓
<b>Identical Orderings</b>						
SD-EF1 Per Day	X	?	✓ (Thm 7)	✓	✓	✓
EF1 Per Day	X	?	✓	✓	✓	✓
$\emptyset$	X (Thm 8)	?	✓	✓	✓	✓
<b>Identical Days</b>						
SD-EF1 Per Day	X	?	?	✓ (Thm 9)	?	✓
EF1 Per Day	X	?	?	✓	?	✓
$\emptyset$	X (Thm 8)	?	✓	✓	✓	✓

Table 1: Possibilities, impossibilities, and open questions in temporal fair division. ✓ indicates a possibility result. X indicates an impossibility. ? indicates an open question. Green highlights indicate the main results of this paper; non-highlighted cells are either open, already known, or implied by other results.

other areas of social choice theory. See the work of Elkind, Obraztsova, and Teh (2024) for a detailed look at temporal multi-winner voting, and for a synopsis of other papers that look at fairness over time. Also, Alamdari et al. (2024) present a model of temporal fairness for a very general decision making setting.

## 2 Preliminaries

### Model

For any  $r \in \mathbb{N}$ , define  $[r] \triangleq \{1, 2, \dots, r\}$ . A multiset is a set that allows repetitions.

**Agents, goods, and valuations.** Let  $N = [n]$  be a set of agents who are allocated a set of goods on each day over  $k$  consecutive days. For  $t \in [k]$ , denote by  $M_t$  the set of goods to be allocated on day  $t$ ,  $\overline{M}_t = \cup_{r \in [t]} M_r$  the set of goods up to day  $t$ , and  $M = \overline{M}_k = \cup_{t \in [k]} M_t$  the set of all goods. We can view  $(M_1, \dots, M_k)$  as a partition of  $M$ . Each agent  $i \in N$  has an additive valuation function  $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$ , where  $v_i(\{g\})$  (henceforth, with a slight abuse of notation, written as  $v_i(g)$ ) is her utility for receiving good  $g \in M$  and  $v_i(S) = \sum_{g \in S} v_i(g)$  for all  $S \subseteq M$ . Collectively,  $(N, (M_1, \dots, M_k), \{v_i\}_{i \in N})$  form an instance of *temporal fair division*. An instance with  $k = 1$  is a (regular) fair division instance, so a temporal fair division instance can be viewed as a sequence of fair division instances in which the same agents participate. We will always assume that an algorithm to solve a temporal fair division problem is given the entire instance as input, i.e. it knows the entire set of goods and what day those goods will arrive upfront.

**Preferences.** Define  $\succsim_i$  (resp.,  $\succ_i$ ) as the weak (resp., strict) ordering over the goods in  $M$  induced by  $v_i$ , where, for all  $g, g' \in M$ ,  $g \succsim_i g'$  if and only if  $v_i(g) \geq v_i(g')$  and  $g \succ_i g'$  if and only if  $v_i(g) > v_i(g')$ . For all  $S \subseteq M$  and all  $r \in \mathbb{N}$ , define  $T_i(S, r)$  to be the  $r$  most preferred goods among the goods in  $S$  according to the ordering  $\succsim_i$ ; all ties are broken consistently across  $i, S$ , and  $r$ .<sup>2</sup>

**Allocations.** An allocation  $A = (A_1, \dots, A_n)$  is a partition of  $M$  into  $n$  pairwise-disjoint bundles, where  $A_i$  is the bundle allocated to agent  $i$ . For  $S \subseteq M$ , let  $A_S = (A_{S,1}, \dots, A_{S,n})$  be the allocation of the goods in  $S$  that is induced by  $A$  (i.e., for each good  $g \in M$  and agent  $i \in N$ ,  $g \in A_{S,i}$  if and only if  $g \in S$  and  $g \in A_i$ ). For  $t \in [k]$ , we refer to  $A_{M_t}$  as the allocation on day  $t$  and  $A_{\overline{M}_t}$  as the allocation up to day  $t$ .

**Restrictions.** We study three restrictions of this general setup (and their combinations).

1. *Two agents:*  $|N| = 2$ .
2. *Identical valuations/orderings:* Under identical valuations,  $v_i = v$  for all agents  $i \in N$ . Under identical orderings,  $\succsim_i = \succ$  for all agents  $i \in N$ . Here, we simply write  $v, \succ$ , and  $T(S, r)$ , skipping the agent in the subscript. For this case, our results deal with desiderata which depend only on the orderings; thus, no distinction between

<sup>2</sup>That is, we use an arbitrary global ordering over  $M$  as the tiebreaker to convert the weak ordering  $\succsim_i$  of every agent  $i \in N$  into her strict ordering over  $M$ , and compute all  $T_i(S, r)$ -s according to these strict orderings. Our negative results do not depend on this tie-breaking and positive results hold regardless of it.

valuations and orderings is necessary.<sup>3</sup>

3. *Identical days*: Informally, copies of the same goods are allocated on each day. Formally, for all days  $t, t' \in [k]$ , there is a bijection  $f : M_t \rightarrow M_{t'}$  such that  $v_i(g) = v_i(f(g))$  for all agents  $i \in N$  and goods  $g \in M_t$ .

## Fairness Desiderata

We first introduce the main desiderata we will be studying. Later, we will introduce their temporal extensions. Other notions referred to in specific sections will be introduced therein.

**Definition 1** (Envy-Freeness Up to One Good (EF1)). *An allocation  $A$  of a set of goods  $S$  is envy-free up to one good (EF1) if for all  $i, j \in N$  with  $A_j \neq \emptyset$ , there exists a  $g \in A_j$  such that  $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ , i.e., no agent envies another agent if some good from the latter agent's bundle is removed.*

In addition to EF1, we will also introduce a weaker notion of measuring fairness that does not require directly comparing one agent's bundle to another's.

**Definition 2** (Proportionality Up to One Good (PROPI)). *An allocation  $A$  of a set of goods  $S$  is proportional up to one good (PROPI) if for all  $i \in N$  with  $A_i \neq S$ , there exists a good  $g \in S \setminus A_i$ , such that  $v_i(A_i \cup \{g\}) \geq \frac{1}{n}v_i(S)$ .*

It is well known that EF1 is a stronger notion than PROPI (Conitzer, Freeman, and Shah 2017).

Given a set of goods  $S$  and a good  $g \in S$ , define an agent  $i$ 's top-set with respect to  $S$  as  $H_i(S, g) = \{g' \in S : g' \succ_i g\}$ . When given only a weak ordering  $\succ_i$  over a set of goods  $S$ , we can compare two bundles  $X, Y \subseteq S$  using the *stochastic dominance* (SD) relation:  $X \succ_i^{\text{SD}} Y$  if for all  $g \in S$ ,  $|X \cap H_i(S, g)| \geq |Y \cap H_i(S, g)|$ . That is,  $X$  has at least as many goods weakly preferred to any good as  $Y$  has. It is known that  $X \succ_i^{\text{SD}} Y$  if and only if  $v_i(X) \geq v_i(Y)$  for all (additive) valuations  $v_i$  over  $S$  that would induce  $\succ_i$ . Hence, using the SD comparison in the EF1 definition yields its stronger counterpart, which has also been studied extensively (Aziz et al. 2015; Freeman, Micha, and Shah 2021; Aziz et al. 2023).

**Definition 3** (SD-EF1). *An allocation  $A$  of a set of goods  $S$  is stochastically-dominant envy-free up to one good (SD-EF1) if for all  $i, j \in N$  with  $A_j \neq \emptyset$ , there exists a  $g \in A_j$  such that  $A_i \succ_i^{\text{SD}} A_j \setminus \{g\}$ .*

We also introduce a stochastic dominance extension of PROPI.

**Definition 4** (SD-PROPI). *An allocation  $A$  of a set of goods  $S$  is stochastically-dominant proportional up to one good (SD-PROPI) if, for all  $i \in N$  with  $A_i \neq S$ , there exists a  $g \in S \setminus A_i$  such that  $|(A_i \cup \{g\}) \cap H_i(S, g')| \geq \lceil |H_i(S, g')|/n \rceil$  for all  $g' \in S$ .*

Informally,  $A_i$ , after adding at most one good to it, must contain at least  $\lceil k/n \rceil$  goods among the  $k$  most preferred goods of agent  $i$  in  $S$ , for each  $k \in [|S|]$ .

<sup>3</sup>In other words, our positive results hold even under identical orderings (weaker restriction), while our negative results hold even under identical valuations (stronger restriction).

Just as EF1 implies PROP1, we have that SD-EF1 implies SD-PROP1, and similarly to SD-EF1, if an allocation  $A$  is SD-PROP1 for certain orderings  $\{\succ_i\}_{i \in N}$ , then  $A$  will be PROP1 for any additive valuation functions that induce  $\{\succ_i\}_{i \in N}$ . Both these facts are proven in the full version.

## Temporal Fairness

In a temporal fair division instance given by a set of goods  $M$  partitioned as  $(M_1, \dots, M_k)$  across  $k$  days, we can ask for fairness to hold at different levels of granularity, yielding various temporal extensions of the fairness desiderata introduced above. These extensions also apply to any other type of desiderata (e.g., efficiency).

**Definition 5** (Per Day Fairness). *For desideratum  $X$ , allocation  $A$  satisfies  $X$  per day if  $A_{M_t}$  satisfies  $X$  for all  $t \in [k]$ .*

**Definition 6** (Overall Fairness). *For desideratum  $X$ , allocation  $A$  satisfies  $X$  overall if  $A_M = A$  satisfies  $X$ .*

**Definition 7** (Up To Each Day Fairness). *For desideratum  $X$ , allocation  $A$  satisfies  $X$  up to each day if  $A_{\overline{M}_t}$  satisfies  $X$  for all  $t \in [k]$ .*

Note that 'up to each day' is a strengthening of 'overall', while 'per day' is incomparable to those two. Plugging in our fairness desiderata into these three temporal extensions gives us the hierarchy of fairness guarantees depicted in Figure 1. Because SD-EF1 is achievable for (regular) fair division (e.g., via a simple round-robin procedure (Caragiannis et al. 2019)), SD-EF1 per day and SD-EF1 overall are both individually achievable, implying the same for EF1, PROP1, and SD-PROP1.

## 3 General Preferences

In this section, we present temporal fair division results in the most general setting: an arbitrary set of goods arrive each day, and each agent has arbitrary additive preferences over them. Let us present our main result for this general setting.

**Theorem 1.** *For any temporal fair division instance, an allocation that is SD-EF1 per day and PROP1 overall exists and can be computed in polynomial time.*

We find such an allocation using Algorithm 1. The derivation of Theorem 1 can be divided into three conceptual steps.

**Identical ordering transformation.** Algorithm 1 begins by creating an auxiliary temporal fair division instance as follows. For each day  $t$ , it creates a new instance for that day with a set of goods  $M'_t$  and valuations  $v'$  such that agents have identical orderings but with the same set of utility values as they had previously. More formally, let  $M'_t = \{g'_{t,1}, \dots, g'_{t,|M_t|}\}$ . Then, for each agent  $i \in N$ ,  $v'_i(g'_{t,1}) \geq v'_i(g'_{t,2}) \dots \geq v'_i(g'_{t,|M'_t|})$  (common ordering) and there exists a bijection  $o_{i,t}$  between  $M_t$  and  $M'_t$  such that  $v_i(g) = v'_i(o_{i,t}(g))$  for all  $g \in M_t$  (same utility values). This technique has been used previously for designing algorithms to achieve (approximate) maximin share fairness (MMS) (Bouveret and Lemaître 2014), but we use it with a novel and nontrivial analysis to ensure PROP1.

---

**Algorithm 1: SD-EF1 per day + PROP1 Overall**


---

```

1: // Identical Ordering Transformation
2:  $v'_i = \emptyset$  for all  $i \in N$ 
3: for  $t \in [k]$  do
4:    $M'_t \leftarrow \{g'_{t,1}, \dots, g'_{t,|M_t|}\}$ 
5:   for  $i \in N$  do
6:      $o_{i,t}$  be the goods  $M_t$  in non-increasing order of  $v_i$ 
7:      $v'_i(g'_{t,j}) \leftarrow v_i(o_{i,t}(j))$  for all  $j \in [|M_t|]$ 
8: // Invoking EF1 with Cardinality Constraints algorithm of
   Biswas and Barman (2018)
9: for  $t \in [k]$  do
10:  Partition  $M'_t$  into groups of size  $n$  (last one may be
    smaller than  $n$ ) as follows:  $C_{t,1} \leftarrow \{g'_{t,1}, \dots, g'_{t,n}\}$ ,
     $C_{t,2} \leftarrow \{g'_{t,n+1}, \dots, g'_{t,2n}\}$ ,  $\dots$ ,  $C_{t, \lceil |M'_t|/n \rceil} \leftarrow$ 
     $\{g'_{t, (\lceil |M'_t|/n \rceil - 1)n + 1}, \dots, g'_{t, |M'_t|}\}$ 
11:  $A' \leftarrow \text{BiswasBarmanCC}(\bigcup_{t \in [k]} \bigcup_{j \in [\lceil |M'_t|/n \rceil]} C_{t,j}, v')$ 
12: // Final Allocation with Daily Picking Sequences based on  $A'$ 
13:  $A \leftarrow \emptyset$ 
14: for  $t \in [k]$  do
15:   for  $j \in [|M_t|]$  do
16:      $i \leftarrow$  Agent allocated  $g'_{t,j}$  in  $A'$ 
17:      $g \leftarrow$   $i$ 's favourite unallocated good from  $M_t$ 
18:      $A_i \leftarrow A_i \cup \{g\}$ 
19: return  $A$ 

```

---

**Connection to cardinality constraints.** Algorithm 1 invokes a key subroutine due to Biswas and Barman (2018) that returns an EF1 allocation subject to cardinality constraints summarized below.

**Theorem 2** (Theorem 1 of (Biswas and Barman 2018)). *Given  $p$  disjoint sets of goods  $C_1, \dots, C_p$  and  $n$  agents with heterogeneous additive valuations, there always exists an EF1 allocation  $A$  such that  $\lfloor |C_\ell|/n \rfloor \leq |A_i \cap C_\ell| \leq \lceil |C_\ell|/n \rceil$  for every agent  $i \in N$  and  $\ell \in [p]$ , and such an allocation can be computed in polynomial time.*

Algorithm 1 invokes the algorithm of Theorem 2 on the following instance. Fix a day  $t$ . Recall that agents have identical ranked preferences for  $M'_t$  given as  $g'_{t,1} \succ \dots \succ g'_{t,|M'_t|}$ . Divide  $M'_t$  into groups of size  $n$  in the decreasing order of value, breaking ties arbitrarily and letting the last group have possibly fewer than  $n$  goods: that is, let  $C_{t,1} = \{g'_{t,1}, \dots, g'_{t,n}\}$ ,  $C_{t,2} = \{g'_{t,n+1}, \dots, g'_{t,2n}\}$ , and so on. By Theorem 2, we find an allocation  $A'$  that is an EF1 allocation for  $\bigcup_{t \in [k]} M'_t$  and  $v'$ , and that each agent is allocated at most one good from  $C_{t,j}$  for all  $t$  and  $j$ .

**Final Allocation.** Algorithm 1 then takes the allocation  $A'$  and, for each day  $M'_t = \{g'_{t,1}, \dots, g'_{t,|M'_t|}\}$ , allocates  $M_t$  through a “serial dictatorship” with the picking sequence derived from  $A'$ . First, the agent that is allocated  $g'_{t,1}$  in  $A'$  will pick their most favourite good from  $M_t$  (the original set of goods for day  $t$ ); next, the owner of  $g'_{t,2}$  picks their most favourite good among the remaining goods of  $M_t$ ; and so on. We now prove that the resulting allocation  $A$  is SD-EF1 per day and PROP1 overall.

**Lemma 1.** *Algorithm 1 returns an allocation that is SD-EF1 per day.*

*Proof sketch.* In the picking sequence for each day  $t$ , due to partitioning of the goods  $C_{t,1}, \dots, C_{t, \lceil |M'_t|/n \rceil}$  and the property that  $A'$  satisfies the “at most one per group” requirement, each agent appears exactly once in the first  $n$  positions, once in the next  $n$  positions, and so on. Additionally, each agent appears at most once among the last  $|M_t| \bmod n$  positions. Such a picking order is known as “recursively-balanced”, and is known to yield SD-EF1 (Aziz 2020). A detailed proof appears in the full version.  $\square$

To prove that  $A$  is PROP1 overall, we use the following technical lemma, the proof of which appears in the full version.

**Lemma 2.** *Let  $V$  be a multiset of  $m$  real values, and  $A = \{a_1, \dots, a_k\}$  and  $A' = \{a'_1, \dots, a'_k\}$  be two subsets of  $V$  with equal size such that  $a_j \geq a'_j$  for all  $j \in [k]$ . Let  $g = \max\{x : x \in V \setminus A\}$  and  $g' = \max\{x : x \in V \setminus A'\}$ . Then, there exists a bijection  $z$  from  $A \cup \{g\}$  to  $A' \cup \{g'\}$  such that  $x \geq z(x)$  for all  $x \in A \cup \{g\}$ .*

**Lemma 3.** *Algorithm 1 returns an allocation that is PROP1 overall.*

*Proof.* Fix an agent  $i \in N$ . Take a day  $t \in [k]$ . Rename the goods so that  $A'_i \cap M'_t = \{g'_1, \dots, g'_{|A'_i \cap M'_t|}\}$  are the goods that  $i$  is allocated in  $A'$  in a non-increasing order of  $v'_i$ . Similarly, let  $A_i \cap M_t = \{g_1, \dots, g_{|A_i \cap M_t|}\}$  be the goods  $i$  picks according to the picking sequence in order. That is,  $g_1$  is the good picked corresponding to  $g'_1$ ,  $g_2$  corresponding to  $g'_2$ , and so on.

Towards invoking Lemma 2, a helpful observation is that  $v_i(g_j) \geq v'_i(g'_j)$  for all  $j \in [|A_i \cap M_t|]$ . Suppose  $g'_j$  is the  $r$ -th preferred good among  $M'_t$ . Since  $M_t$  and  $M'_t$  share the same multiset of utility values for  $i$ , and  $g_j$  is the top pick of  $i$  when  $r - 1$  goods are picked,  $g_j$  is at least as good as  $i$ 's  $r$ -th most preferred good among  $M_t$  (and hence,  $M'_t$ ). This argument, combined across all days, implies existence of a bijection  $z_i : A_i \rightarrow A'_i$  such that  $v_i(g) \geq v'_i(z_i(g))$  for all  $g \in A_i$ .

By invoking Lemma 2 with  $A \leftarrow A_i$  and  $A' \leftarrow A'_i$  over the multiset  $V$  being  $i$ 's utility values, we have that

$$v_i(A_i) + \max_{g \notin A_i} v_i(g) \geq v'_i(A'_i) + \max_{g \notin A'_i} v'_i(g)$$

Every EF1 allocation is also PROP1 (Conitzer, Freeman, and Shah 2017), therefore, since  $A'$  is EF1 (Theorem 2), we have that

$$v'_i(A'_i) + \max_{g \notin A'_i} v'_i(g) \geq \frac{1}{n} v'_i(M') = \frac{1}{n} v_i(M),$$

the last equality being true from the way we constructed  $M'$ . Combining the two inequalities above, we get  $v_i(A_i) + \max_{g \notin A_i} v_i(g) \geq \frac{1}{n} v_i(M)$ . Thus,  $A$  is PROP1.  $\square$

It is worth noting that PROP1 is not a monotonic property, i.e., if  $v_i(A_i) \geq v_i(A'_i)$  and  $A'_i$  is PROP1, it is possible that  $A_i$  is not PROP1 (as the best good for  $i$  in  $M \setminus A_i$  could be worth less than the best good in  $M \setminus A'_i$ ). This is why we needed to use a more involved argument in Lemmas 2 and 3.

## 4 Two Agents

In this section, we consider temporal fair division with two agents. We provide a complete picture of temporal fairness notions that can be guaranteed in this case. We establish a strong positive result, then show that it is the best possible by producing counterexamples for stronger desiderata.

### Possibilities

Our main goal in this section is to show that SD-EF1 per day and EF1 up to each day can be achieved for two agents. We begin by introducing an *envy-balancing lemma*, a powerful tool for finding temporal allocations to two agents. We later use this lemma to derive not only the aforementioned guarantee, but also other appealing guarantees.

**Definition 8** (Cancelling Allocations). *We say that allocations  $B$  and  $B'$  of a set of goods  $S$  to two agents cancel out if  $v_i(B_i) + v_i(B'_i) \geq v_i(B_{3-i}) + v_i(B'_{3-i}), \forall i \in [2]$ . In words, they cancel out if hypothetically allocating two copies of each good in  $S$ , one according to  $B$  and the other according to  $B'$ , achieves (exact) envy-freeness.*

**Lemma 4** (Envy-Balancing Lemma). *Suppose that for each day  $t \in [k]$ , we are given two EF1 allocations  $B_t$  and  $B'_t$  of the set of goods  $M_t$  that cancel out. Then, we can compute, in polynomial time, an allocation  $A$  of the set of all goods  $M = \cup_{t \in [k]} M_t$  that is EF1 up to each day and  $A_{M_t} \in \{B_t, B'_t\}$  for each day  $t \in [k]$ .*

Note that the lemma achieves EF1 up to each day while not only retaining the EF1 per day property of the input allocations, but in fact by using exactly one of the two input allocations on each day. Thus, if the per day allocations given as input satisfy properties stronger than EF1, those properties are also retained per day; this is important as we will use this lemma to derive such stronger per day guarantees.

We include the proof of the lemma in the full version. At a high level, it works by realizing that when two allocations cancel out, we know that both agents have at least one of the two allocations where they feel no envy (they like the bundle they were given more than the bundle given to the other agent). With this in mind, we can carefully choose which allocation to assign on each day, in such a way that whenever an agent is feeling too much envy, we can give them their preferred allocation on that day, always keeping envy levels “balanced” after each time-step.

For readers familiar with the “informed” model of online fair division from He et al. (2019), this can be seen as a generalization of their two agent algorithm. While they achieve EF1 up to each day while allocating a single good during each time step, we provide a similar guarantee while allocating batches of goods and simultaneously maintaining fairness over the batches.

We are now ready to show that the pair of allocations required by the envy-balancing lemma — both satisfying EF1 and cancelling each other out — exists and can be computed in polynomial time. In fact, we find a single partition  $(B_{t,1}, B_{t,2})$  of the goods in  $M_t$  such that both  $B_t = (B_{t,1}, B_{t,2})$  and  $B'_t = (B_{t,2}, B_{t,1})$  satisfy the stronger property of SD-EF1.

**Lemma 5.** *Given the preferences of two agents over a set of goods  $M_t$ , one can efficiently compute a partition  $(B_{t,1}, B_{t,2})$  of  $M_t$  such that both  $B_t = (B_{t,1}, B_{t,2})$  and  $B'_t = (B_{t,2}, B_{t,1})$  are SD-EF1 allocations.*

We can plug these allocations into the envy-balancing lemma (Lemma 4) to get our desired main result.

**Theorem 3.** *For temporal fair division with  $n = 2$  agents, an allocation that is SD-EF1 per day and EF1 up to each day exists and can be computed in polynomial time.*

We include both the proofs for Lemma 5 and Theorem 3 in the full version. For the interested reader, we also include some notes on the connection between Lemma 5 and the well-known Bihierarchy Theorem from Budish et al. (2013), as well as how to achieve other fairness notions per day while using the envy-balancing lemma, such as EFX and EF1+PO.

### Impossibilities

We have shown that when there are only two agents, the strong guarantee of EF1 up to each day can be obtained along with the guarantee of SD-EF1 per day. However, one wonders if even stronger guarantees are possible, such as strengthening EF1 up to each day to SD-EF1 up to each day. We find that not only is this strengthening impossible, but even if we relax SD-EF1 up to each day to SD-EF1 overall, it is impossible to achieve alongside EF1 per day. Together, these two impossibility results prove that our guarantee of SD-EF1 per day and EF1 up to each day from Theorem 3 is the strongest possible in the hierarchy shown in Figure 1. We include the counterexamples proving both these claims in the full version.

**Theorem 4.** *For temporal fair division with  $n = 2$  agents, SD-EF1 up to each day cannot be guaranteed.*

**Theorem 5.** *For temporal fair division with  $n = 2$  agents, EF1 per day and SD-EF1 overall cannot be guaranteed simultaneously.*

### Two Agents and Further Restrictions

As the final part of this section, we note that when further restrictions are placed on two-agent instances, we receive even stronger results. Particularly, in the full version, we prove and provide a discussion of the following theorem.

**Theorem 6.** *For temporal fair division with  $n = 2$  agents and identical days, an allocation that is SD-EF1 per day, SD-EF1 up to each day, and SD-EF up to each even day exists and can be computed in polynomial time.*

## 5 Identical Orderings

We next look at instances where agents have identical orderings over all goods in  $M$ . Not only are results in this setting practically useful, as there many real life scenarios where participants agree on the ordinal ranking of goods, but results under identical orderings are also very technically useful. As can be seen from our main result in Section 3, reducing a general setting to one where agents have similar orderings over the goods can lead to fairness guarantees in the original setting. We will show in future sections

that the possibility results we develop here can be applied as black-boxes to achieve strong results in scenarios where agents have heterogeneous orderings.

**Possibilities.** Our main result for the case of identical orderings is the following theorem:

**Theorem 7.** *For temporal fair division with identical orderings, an allocation that is SD-EF1 per day and SD-EF1 overall exists and can be computed in polynomial time.*

We provide a detailed proof of Theorem 7 in the full version. Intuitively, we accomplish this by creating two partitions over the set of goods  $M$ , labeled  $P_1$  and  $P_2$ , which are defined below.

$$P_1 = \left\{ T(M_t, nr) \setminus T(M_t, n(r-1)) : r \in \left[ \left\lceil \frac{|M_t|}{n} \right\rceil \right], t \in [k] \right\},$$

$$P_2 = \left\{ T(M, nr) \setminus T(M, n(r-1)) : r \in \left[ \left\lceil \frac{|M|}{n} \right\rceil \right] \right\}.$$

In words,  $P_2$  splits the entire set of goods  $M$  into the agents' most preferred  $n$  goods, their next most preferred  $n$  goods, etc.  $P_1$  does a similar partitioning over  $M$ , but does the partitioning separately for each day. It is known that when agents have identical orderings, guaranteeing that each agent receives 1 good from their  $n$  favourite goods, 1 good from their next  $n$  favourite goods, etc. will guarantee an SD-EF1 allocation. We prove that it is possible to construct an allocation where each agent gets exactly 1 good from each set in  $P_1$  and from each set in  $P_2$ , thereby ensuring SD-EF1 per day and overall.

To an initiated reader, the problem of finding an allocation that meets the above constraints may be immediately reminiscent of the bihierarchy framework of Budish et al. (2013). We require that from each set of (at most)  $n$  goods out of a family of sets, each agent receives (at most) one good. The sets produced by each desideratum are mutually non-overlapping, forming a "hierarchy", but the sets produced by one desideratum can be overlapping with those produced by the other, resulting in two different hierarchies. However, the problem is that with  $n > 2$  agents, we have a third set of constraints: each good must be assigned to (exactly) one agent. This forms a third hierarchy (which cannot be assimilated into either of the two previous hierarchies), preventing one from applying the bihierarchy framework.

**Impossibilities.** In the case of 2 agents, imposing the additional restriction of identical days allowed for very strong results, making it possible to satisfy SD-EF1 per day and SD-EF1 up to each day. This is unfortunately not the case when the identical days restriction is imposed in addition to identical preferences. We show via a fascinating counterexample in the full version that even when an instance has identical preferences and identical days, it is not always possible to even achieve SD-EF1 up to each day by itself.

**Theorem 8.** *For temporal fair division with identical days and identical preferences, SD-EF1 up to each day cannot be guaranteed.*

## 6 Identical Days

In Section 3, we showed that in the general model, we can achieve SD-EF1 per day and PROP1 overall. In this section,

we will show if we assume the additional restriction that the sets of goods which arrive on each day are identical, then a slightly stronger guarantee can be achieved overall.

**Theorem 9.** *For any temporal fair division instance with identical days, it is possible to find an allocation that is SD-EF1 per day and SD-PROP1 overall in polynomial time.*

To achieve these guarantees, we use an algorithm that is almost identical to Algorithm 1, the algorithm which was used to achieve SD-EF1 per day and PROP1 overall in the general case, but with one major change.

Algorithm 1 first finds an EF1 overall allocation in a reduced version of the problem where agents have identical orderings over all the goods on each day, and uses that to construct an allocation in the original instance that is PROP1 overall. The key insight that can be leveraged to get stronger guarantees in this less general setting is that when we have identical days, we know that the output of the "Identical Ordering Transformation" from Algorithm 1 will result in an instance where all agents have identical orderings over the entire set of goods  $M$ , not just over the goods from each individual day. This stronger guarantee from the Identical Ordering Transformation allows us to use Theorem 7 to find an SD-EF1 per day and SD-EF1 overall allocation over the reduced instance (rather than the algorithm of Biswas and Barman (2018) which only guarantees EF1 overall). We can then use this allocation as the basis for the picking order that Algorithm 1 uses to construct the final allocation over the original instance. We include a detailed proof of Theorem 9 in the full version, including why an SD-EF1 allocation over the identical orderings instance will lead to a SD-PROP1 allocation in the original instance.

## 7 Discussion

In this work, we are able to find possibility and impossibility results that give a picture of what can be achieved in the temporal fair division model. This picture is quite clear when focusing on special cases such as two agents or identical orderings. However, we still leave many questions open for future work. All of the entries in Table 1 marked by a "?" remain open. The most interesting question in the general setting is:

**Open Question:** In temporal fair division, does an allocation that is SD-EF1 (or EF1) per day and EF1 overall always exist?

Other interesting open questions include the existence of EF1 up to each day under identical orderings or identical days, and the existence of (SD-)EF1 per day and (SD-)EF1 overall under identical days.

Finally, it would be an interesting further direction to take a more abstract view of the temporal fair division model. We explore this in the full version of this paper, introducing a generalized model of temporal fair division, where a fair allocation must be found simultaneously over a set of goods, and over each set in a collection of subsets of those goods. It would be very interesting to explore this interpretation of our model further.

## Acknowledgements

This research was partially supported by an NSERC Discovery grant and an NSERC-CSE Research Communities Grant. Researchers funded through the NSERC-CSE Research Communities Grants do not represent the Communications Security Establishment Canada or the Government of Canada. Any research, opinions or positions they produce as part of this initiative do not represent the official views of the Government of Canada.

## References

- Alamdari, P. A.; Klassen, T. Q.; Creager, E.; and McIlraith, S. A. 2024. Remembering to Be Fair: Non-Markovian Fairness in Sequential Decision Making. In *Proceedings of the 41th International Conference on Machine Learning (ICML)*, 906–920.
- Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; and Voudouris, A. A. 2022. Fair Division of Indivisible Goods: A Survey. In *Proceedings of the 31st European Conference on Artificial Intelligence (ECAI)*, 5385–5393.
- Aziz, H. 2020. Simultaneously Achieving Ex-ante and Ex-post Fairness. In *Proceedings of the 21st Conference on Web and Internet Economics (WINE)*, 341–355.
- 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*, 72(4): 1674–1688.
- Aziz, H.; Gaspers, S.; Mackenzie, S.; and Walsh, T. 2015. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227: 71–92.
- Banerjee, S.; Gkatzelis, V.; Gorokh, A.; and Jin, B. 2022. Online nash social welfare maximization with predictions. In *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1–19.
- Banerjee, S.; Gkatzelis, V.; Hossain, S.; Jin, B.; Micha, E.; and Shah, N. 2023. Proportionally Fair Online Allocation of Public Goods with Predictions. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 20–28.
- Benadè, G.; Kazachkov, A. M.; Procaccia, A. D.; Psomas, A.; and Zeng, D. 2024. Fair and efficient online allocations. *Operations Research*, 72(4): 1438–1452.
- Bertsimas, D.; Farias, V. F.; and Trichakis, N. 2013. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1): 73–87.
- Biswas, A.; and Barman, S. 2018. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 91–97.
- Bogomolnaia, A.; Moulin, H.; and Sandomirskiy, F. 2022. On the fair division of a random object. *Management Science*, 68(2): 1174–1194.
- Bouveret, S.; and Lemaître, M. 2014. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 1321–1328.
- 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.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation*, 7(3): Article 12.
- Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair Public Decision Making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 629–646.
- Dror, A.; Feldman, M.; and Segal-Halevi, E. 2023. On fair division under heterogeneous matroid constraints. *Journal of Artificial Intelligence Research*, 76: 567–611.
- Elkind, E.; Obraztsova, S.; and Teh, N. 2024. Temporal Fairness in Multiwinner Voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 22633–22640.
- Freeman, R.; Micha, E.; and Shah, N. 2021. Two-Sided Matching Meets Fair Division. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 203–209.
- Ghods, A.; Zaharia, M.; Shenker, S.; and Stoica, I. 2013. Choosy: Max-min fair sharing for datacenter jobs with constraints. In *Proceedings of the 8th ACM European Conference on Computer Systems*, 365–378.
- Gkatzelis, V.; Psomas, A.; and Tan, X. 2021. Fair and efficient online allocations with normalized valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 5440–5447.
- He, J.; Psomas, C.-A.; Procaccia, A. D.; and Zeng, D. 2019. Achieving a Fairer Future by Changing the Past. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 348–353.
- Igarashi, A.; Lackner, M.; Nardi, O.; and Novaro, A. 2024. Repeated Fair Allocation of Indivisible Items. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 9781–9789.
- 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.
- Lee, M. K.; Kusbit, D.; Kahng, A.; Kim, J. T.; Yuan, X.; Chan, A.; See, D.; Noothigattu, R.; Lee, S.; Psomas, A.; et al. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW): 181:1–181:35.
- Mehta, A.; Saberi, A.; Vazirani, U.; and Vazirani, V. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5): 22:1–22:19.
- Miller, H. E.; Pierskalla, W. P.; and Rath, G. J. 1976. Nurse scheduling using mathematical programming. *Operations Research*, 24(5): 857–870.
- Zeng, D.; and Psomas, A. 2020. Fairness-efficiency trade-offs in dynamic fair division. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 911–912.