Indivisible Goods Allocation in Future-Looking Multi-Round Settings

JACK DESCHLER, GILI RUSAK, ITAI SHAPIRA, SHIRLEY ZHANG

Fair allocation of indivisible goods is a problem with significant real-world applications and numerous variations. In this paper, we focus on an offline, multi-round setting where goods arrive in predetermined batches, and agents discount future valuations of items while "forgetting about the past." This model captures the essence of certain allocation problems, such as resource allocation in food banks, vaccine distribution, and position assignments in organizations. Our main results include: (1) demonstrating that under general settings, it is impossible to guarantee envy-free up to one good (EF1) fairness without adjustments; (2) constructing an algorithm that ensures EF1 for two agents; (3) bounding the number of adjustments needed in the general case; and (4) constructing a fair algorithm for the case of binary valuations. These findings advance the understanding of fair allocation in multi-round settings with future discounting and have potential applications in various real-world domains.

1 Introduction

The fair allocation of a set of indivisible goods is a problem with many variations and real-world applications. In its simplest form, the task is to divide a set of discrete goods that cannot be divided further or separated as fairly as possible to a set of agents, each of whom may value any individual good differently. One notion of "fair" allocation is ensuring that after allocation no agent would prefer to have another agent's bundle of goods. In many cases an "envy-free" allocation (where no agent envies another) is impossible, and so algorithms seek to make allocations that are envy-free up to a certain number of goods, or to minimize the amount that an agent envies another's bundle. Variations on the problem include changing the number of players, having varying numbers of goods arrive at the same time in batches, and allowing for adjustments to previously allocated goods. Our paper proposes a new model, one that assumes goods arrive in a known number of predetermined batches while allowing agents to discount the past, and focuses on when the model can guarantee that an allocation that is envy-free up to one good at the end of each round.

The indivisible goods problem has a multitude of real-world analogs. Food banks play a version of this game, when they allocate food to hungry families, who all may value a can of beans or a loaf of bread differently. The food bank setting is an "uninformed" (sometimes referred to as "online" in the literature) setting, where the number of batches and number of goods in each batch is not known. Compare this setting to rural hospitals receiving COVID-19 vaccines in early 2021, when governments could possibly guarantee the number of vaccines of each type in a guaranteed number of batches to come, but may still want to guarantee some level of fairness at each round of allocation. Different hospitals may have had different valuations for the different types of COVID-19 vaccine, creating an "informed" (or "offline," in some of the literature) allocation problem, where the information about future goods and batches is known at the allocation of each individual batch. The goods in each of these two settings were both exhaustible - once the food was handed out or the vaccines used, they obviously could not be given from one agent to another upon the arrival of a later batch in order to bandage envy caused by that batch. Goods that are inexhaustible, though, can be adjusted later, creating yet another flavor of the allocation problem (He et al. [4]). In this setting, one might think of positions or benefits in a university that are allocated to professors, or in a business that are allocated to partners. A corner office does not simply go away in the same manner that a COVID-19 vaccine does once someone uses it for a year; it can be reallocated at a later date to make an overall allocation more fair once a new batch of "goods" comes in.

Our model focuses on the offline setting. In this model, agents discount future valuations of items, which means they assign less importance to items that will be allocated in future rounds, consequently affecting their preferences and priorities during the allocation process. Furthermore, agents "forget about the past," signifying that they disregard previous allocations in their valuations. Our research aims to identify the limitations of fair mechanisms in these multi-round settings and develop a new algorithm that can guarantee EF1 fairness under various assumptions.

The structure of the paper is organized as follows: (a) Preliminaries, where we introduce the necessary definitions, notation, and concepts used throughout the paper; (b) Impossibility Result, where we prove that no algorithm can guarantee EF1 in multi-round settings with more than two agents; (c) Two-Agent Case, where we focus on the two-agent case, including the Backwards Induction Envy Balancing Algorithm that ensures fair solution; (d) Adjustments, where we explore the upper bound of the number of adjustments needed to achieve fairness; and (e) Binary Valuations, where we delve into the specific case of binary valuations.

1.1 Related Work

In our study of online fair division with repeated rounds and a discount factor, we aim to achieve envy-free up to one item (EF1) allocations. It is crucial to differentiate our model from prior literature, as variations in the studied models substantially impact the allocation mechanisms developed. In this section, we discuss the prior literature that has been studied in these contexts. A thorough comprehension of these differences allows us to identify unique algorithms and mechanisms for our specific setting. Building upon the insights from previous works, our research introduces a novel approach to online fair division by incorporating both future-looking repeated round allocation and discount factors. This distinctive perspective enables the development of tailored mechanisms and algorithms, further advancing the field of online fair division and promoting the creation of fair allocation mechanisms in diverse contexts.

Our work builds upon a foundation laid by He et al. [4], who investigate a model in which T indivisible items arrive sequentially, one at a time. In their study, they focus on guaranteeing EF1 at each round, considering an "accumulated" notion of fairness, where agents are not future-looking and accumulate utility from previous rounds. In contrast, our model explores a future-looking perspective with agents discounting future valuations and "forgetting about the past" in their decision-making process. Additionally, our model incorporates the arrival of batches of items. While the connection may not be immediately apparent or straightforward, in some cases our model can be reduced to the results presented by He et al. [4]. Consequently, some of the findings discussed in this paper can be regarded as extensions or adaptations of their work. For two players in the offline setting, their present an algorithm that guarantees EF1 with no adjustments. They also provide an algorithm for the online setting that uses $O(T^{3/2})$ adjustments. In our work, we note that the impossibility result for $n \geq 3$ players still holds in our setting, and adapt both the two-player offline algorithm and the three-player online algorithm from He et al. [4] to our context.

Our work also relates to other recent work in the space of online fair division. Benade et al. [3] study the problem of minimizing the maximum envy when there are T items arriving in an online fashion and must be allocated upon arrival. This paper shows that there exists a polynomial-time, deterministic algorithm where the ratio of envy over time goes to zero as T goes to infinity. While our problem is not inherently online, it does involve items arriving over time and losing value after they arrive.

Our work also builds on the survey work of Aleksandrov and Walsh [1]. "Online Fair Division: A Survey" is a comprehensive review of recent research on fair division algorithms for online settings. The paper provides an overview of the challenges involved in allocating resources fairly in online environments and reviews various approaches that have been proposed to address them. The authors discuss the theoretical foundations of fair division, including concepts such as envy-freeness, proportionality, and equitability, and examine how these concepts can be applied in online settings. They also review a range of algorithms that have been proposed for online fair division, including probabilistic and combinatorial methods, as well as adaptive and dynamic algorithms. The paper discusses and shows different properties of different online mechanisms. For example, they highlight that an informed algorithm requires no re-allocations to ensure EF1 while an uninformed algorithm requires $\Theta(m)$ re-allocations to guarantee EF1. They also discuss both asymptotic guarantees and intractability results that have been proven in these settings (Aleksandrov and Walsh [1]). In our work, we discuss the informed case and particularly make contributions to the EF1 scenario without adjustments in the general setting. We also make new extensions to the binary valuations case and show that the algorithm with binary valuations is EF1 at every round. To the best of our knowledge, this specific setting not been studied in this literature to date.

2 Preliminaries

Consider a set of n agents $A = \{a_1, \ldots, a_n\}$ and a set O containing all available items (sometimes referred to interchangeably as "goods"). We define a sequence of T batches of goods, with the goods in each batch denoted by $G^1, \ldots, G^T$. At each round t, the batch is of size t_m and individual goods are numbered such that $G^t = \{g^t_1, \ldots, g^t_{t_m}\}$. At the end of the final round, all goods in O have been allocated: $O = \bigcup_{t=1}^{T} G^t$. Each agent a_i possesses a valuation function v_i that represents the value assigned to a subset of items $S \subset O$. For simplicity, we use $v_i(g_j)$ to denote $v_i(\{g_j\})$. We assume that $v_i(S) \geq 0$ for every subset of items S. A valuation function v_i is additive if $v_i(S) = \sum_{g \in S} v_i(g)$. An allocation of goods at round t is represented by a partition $A^t = (A^t_1, \ldots, A^t_n)$, where $A^t_i \subset G^i$ is the bundle of goods allocated to agent i.

Intuitively, agents possess knowledge of the future but discount future valuations, favoring immediate allocation of items. Moreover, agents "forget about the past" in the sense that in round t, all items allocated in round $t' < t$ do not contribute to the agents' utility. We formalize this concept by introducing a discount factor $\delta \in (0, 1]$. The *utility* of agent i at round t is the discounted valuations of all the items they will receive, given by:

$$U_i(A^t_i, ..., A^T_i) := \sum_{t'=t}^{T} \delta^{t'-t} v_i(A^{t'}_i)$$

Our focus is on fair allocations that exhibit envy-free up to one item (EF1) characteristics. Informally, an EF1 allocation is one where, for every instance of one player envying another, the envy could be resolved by removing only one good from the envied player's set. In our setting, we define an allocation $A = (A^1, \ldots, A^T)$ to be *envy-free up to one item at round t* if for any pair of agents i and j, there exists $k \geq t$ and an item $g \in G^k$ such that:

$$U_i(A^t_i, \ldots, A^T_i) \geq U_i(A^t_j, \ldots A^k_j \setminus \{g\} \ldots A^T_j)$$

That is, agent i does not envy player j if we remove item g from player j's allocation. We say that an allocation A is EF1 if it is EF1 at *every* round $t \in [T]$.

The *Round Robin* mechanism is an allocation method in which goods are allocated to agents in a cyclical and pre-determined order, with each agent selecting their most preferred item from the available pool in their turn. This method ensures an EF1 allocation in a single-round settings.

3 The Impossibility Result

We begin our discussion by demonstrating that no mechanism in multi-round settings can guarantee EF1 when more than two agents are involved. This impossibility result underscores the inherent challenges in devising an algorithm that ensures fairness at every round with time discounts especially in real world settings where the number of rounds may be high. Formally, we present the following theorem:

Theorem 1. *For any $n > 2$, assume that items arrive in at least 16 distinct batches ($T \geq 16$). There exist valuations $(v_i)_{i \in [n]}$ and a time discount factor $\delta \in (0, 1)$ such that, for any allocation algorithm, the resulting allocation fails to achieve EF1 at every round.*

The result presented here is primarily an extension of Theorem 4.2 from He et al. [4], with the key differences being that agents in our setting consider future items and "forget" the past, as well as the introduction of batches. The observation most relevant to us is that their result can be extended by "reversing the order" of the items and allowing for batches of more than one item. Adapting their result to our setting requires only minor adjustments, which we have included in full in Appendix A. In doing so, we not only demonstrate the broader applicability of their findings, but also set the stage for the rest of our results and highlight the challenges in designing fair allocation mechanisms for more complex situations involving multiple agents.

4 Two Agents

Given the impossibility result discussed in the previous section, we now turn our attention to the simpler two-player case. In this section, we introduce the Backwards Induction Envy Balancing algorithm (Algorithm 1), which ensures envy-freeness up to one item within our extended framework. Our primary objective is to provide a comprehensive explanation of the algorithm and prove its EF1 guarantee.

The Backwards Induction Envy Balancing algorithm (Algorithm 1) builds upon the "Envy Balancing" algorithm presented by He et al. [4]. Our extended algorithm adapts the envy balancing concept to suit our problem setting. (Algorithm 1) aims to iteratively build an allocation for two agents in an online setting, ensuring that all suffixes of the allocation are EF1 allocations. Each player i has two baskets: permanent basket P_i and a temporary basket D_i. The algorithm initializes both of them to the empty set for each agent and iterates through each round in reverse order, from the last round T to the first round. At each step, the algorithm first checks if agent a_1 is unenvied with respect to the temporary baskets D. If so, it updates the temporary allocation D by applying the Round Robin algorithm to the goods G^t, starting with agent a_1 (we denote by RoundRobin(G, a) the allocation made by the Round Robin mechanism with respect to items G and agent a choosing first). Otherwise, if a_1 is not unenvied, the algorithm updates the temporary allocation D by applying the RoundRobin algorithm to the goods G^t, starting with agent a_2.

After updating the temporary allocations, there are three cases to consider with respect to the temporary baskets. First, if both agents envy each other with respect to the temporary baskets $D = (D_1, D_2)$ the algorithm swaps the temporary baskets of both agents, creating an envy-free allocation with respect to the items in D. Second, if the current allocation is envy-free with respect to D, the algorithm moves the items from each agent's temporary basket to their permanent baskets, effectively transferring the envy-free allocation to the permanent baskets. This second step will also happen *after* the temporary baskets are swapped, as that allocation is now envy-free. Finally, if exactly one agent envies, we move to the next step. Intuitively, these steps ensure that if we swap items at round t to guarantee fairness, we do not affect the envy-free allocation at rounds $t' > t$. If $g^{t'} \in G^{t'}$ remains in a player's temporary baskets, it must be the case that the (temporary) allocation was not envy-free at rounds t to t', and a swap at round t would only change which agent was envious. If the allocation was envy-free, we "fix" the allocation by moving it to the permanent basket.

Finally, after iterating through all rounds, the algorithm constructs the final allocation A^t for each round t using the contents of the permanent and temporary baskets. The algorithm returns a sequence of allocations for all rounds $(A^1, A^2, \ldots, A^T)$.

Algorithm 1 Backwards Induction Envy Balancing

Require: v_1, v_2, δ
$P \leftarrow (\emptyset, \emptyset)$, $D \leftarrow (\emptyset, \emptyset)$ {Initialize permanent and temporary baskets}
for $t = T$ to 1 **do**
  if a_1 is unenvied in D **then**
    $D \leftarrow D \cup \text{RoundRobin}(G^t, a_1)$
  else
    $D \leftarrow D \cup \text{RoundRobin}(G^t, a_2)$
  end if
  if both a_1 and a_2 envy each other in D **then**
    $D \leftarrow (D_2, D_1)$ {Resolve envy by swapping temporary baskets}
  end if
  if D is envy-free **then**
    $P \leftarrow P \cup D$ for both agents i {temporary to permanent}
    $D \leftarrow (\emptyset, \emptyset)$ for both agents i
  end if
end for
Construct the final allocation A^t by $(P \cup D) \cap G^t$ for each round t
return $(A^1, A^2, \ldots, A^T)$

Theorem 2. *Algorithm 1 is EF1 at every round.*

The complete proof can be found in Appendix B.

5 Adjustments

Theorem (1) establishes that no algorithm can guarantee EF1 in every round for more than 2 players. Consequently, in this section, we focus on determining the number of adjustments needed to achieve this guarantee. We provide an upper bound for the number of adjustments required when a total of T items arrive in k batches, for $n \geq 3$ players, by adapting the Double Round Robin algorithm proposed in He et al. [4]. Note that in contrast to notation introduced in the preliminaries, in this section T represents the total number of items instead of the number of rounds. We make this notation change to highlight the way in which changing the number of rounds affects the bound on the number of adjustments needed.

Algorithm 2 Double Round Robin

Require: v_i for each agent a_i

$M \leftarrow \emptyset, S \leftarrow \emptyset$
for $t = T/k$ to 1 **do**
  $S \leftarrow S \cup \{G_t\}$
  if $|S| \geq \sqrt{k} \cdot \sqrt{T}$ **then**
    $M \leftarrow M \cup S$
    $S \leftarrow \emptyset$
  end if
  $A_M \leftarrow \text{RoundRobin}(M, a_1 > ... > a_n)$
  $A_S \leftarrow \text{RoundRobin}(S, a_n > ... > a_1)$
  Let A^t be the combination of allocations A_S and A_M
end for
return $(A^1, A^2, \ldots, A^T)$

At a high level, the algorithm works in the following way. In each round, all the items to be allocated are added to the side pile. Whenever the number of items in the side pile exceeds $\sqrt{k} \cdot \sqrt{T}$, we move the items from the side pile to the main pile (clearing the side pile). After resolving the current state of the piles, the players take turns selecting from the main pile in round robin order, and take turns selecting from the side pile in reverse round robin order. The allocation at the end of each round is the combination of both round robin allocations.

Theorem 3. *Algorithm 2 is an EF1 algorithm that requires $O(T^{3/2}/\sqrt{k})$ adjustments.*

Proof. We first prove that the allocation is EF1 for each round. This part of the proof is identical to that in section 4 of He et al. [4], but we summarize it here for completeness. Observe that when items are allocated in a round robin fashion, for any pair of agents a_i and a_j, a_i will envy a_j by at most one item. Furthermore, if a_i was earlier in the round robin order than a_j, then a_i does not envy a_j at all, as we could map each item in a_j's bundle to an item in a_i's bundle that a_i has a higher value for.

We observe that the main pile and side pile are allocated using reverse order, which implies that a_i only envies a_j in the main pile allocation if a_i does not envy a_j in the side allocation, and vice versa. When a_i envies a_j, we know that a_i only envies a_j by one item, so a_i must only envy a_j by one item total after combining allocations.

We now prove that the algorithm uses at most $O(T^{3/2}/\sqrt{k})$ adjustments. The main pile only changes every $\sqrt{k} \cdot \sqrt{T}$ rounds. Therefore, we can observe that:

$$\begin{aligned}
\text{Total adjustments in main pile} &\leq \sqrt{Tk} + 2\sqrt{Tk} + ... + T \\
&= \left(1 + 2 + ... + \frac{\sqrt{T}}{\sqrt{k}}\right)\sqrt{Tk} \\
&= \frac{1}{2} \cdot \frac{\sqrt{T}}{\sqrt{k}}\left(\frac{\sqrt{T}}{\sqrt{k}} + 1\right) \cdot \sqrt{Tk} \\
&= O(T^{3/2}/\sqrt{k})
\end{aligned}$$

On the other hand, the side pile might be reallocated every round, but, crucially, there are only T/k rounds and at most $\sqrt{Tk}$ items in the side pile in any given round. Therefore:

$$\text{Total adjustments in side pile} \leq \sqrt{Tk} \cdot \frac{T}{k} = O(T^{3/2}/\sqrt{k})$$

This implies that the total number of adjustments is also $O(T^{3/2}/\sqrt{k})$. □

Here, we provide a brief discussion on how the result for our setting compares to the result in [4]. Most notably, our algorithm uses $O(T^{3/2}/\sqrt{k})$ adjustments instead of $O(T^{3/2})$ adjustments because we wait until there are more items in the side pile before moving the items in the side pile to the main pile. Specifically, we move the side pile to the main pile whenever the side pile accumulates at least $\sqrt{k} \cdot \sqrt{T}$ items instead of $\sqrt{T}$ items. We are able to wait for more items to accumulate because there are fewer rounds total, which significantly impacts the number of times the side pile is reallocated. We further observe that because we are allowing readjustments of each pile whenever the pile is changed, the addition of δ does not change our algorithm. Finally, like in [4], our algorithm works whether or not we have knowledge of the future.

6 Binary Valuations

In this section, we focus on the binary valuation case where agents assign a value of either 0 or 1 to each item. In other words, they either value an item (1) or not at all (0). The binary valuation scenario is particularly interesting as it contrasts with the general settings (Theorem 1) where no algorithm can guarantee EF1 at every round.

We investigate an extended version of the BALANCED LIKE algorithm (Aleksandrov et al. [2]), which aims to balance the number of items allocated to agents. In every batch, it assigns the subsequent item exclusively to those agents who value it, resolving any tie by allocating the item to the agent with the lowest cumulative utility. The BALANCED LIKE algorithm is proven to be EF1 for the binary valuations case. Although the core of the algorithm remains the same as in (Aleksandrov et al. [2]), the proof generalizes their argument to adapt to our extended settings.

Algorithm 3 Extended Balanced Like

Require: $(v_i)_{i \in [n]}, \delta$

$D \leftarrow \emptyset$

$A^t = (\emptyset, \ldots, \emptyset)$ for every round $t \in [T]$

for $t = T$ to 1 **do**

  for $i = 1$ to $|G^t|$ **do**

    $D \leftarrow \{a_j \mid v_j(g_i) = 1\}$

    Let $j \in [n]$ with $a_j = \arg\min\{U_j(A_j^{t+1}, ..., A_j^T) \mid a_j \in D\}$

    $A_j^t \leftarrow A_j^t \cup \{g_i\}$

  end for

end for

return $(A^1, A^2, \ldots, A^T)$

Theorem 4. *Algorithm (3) with binary valuations is EF1 at every round.*

The main step of the proof shows that if, for example, agent a_1 gets the item g^t, then any other player k that also values g^t does not envy player a_1 with respect to the allocation of the items from round $t+1$ to T. The proof demonstrates that the valuations of all agents from their own allocations remain approximately balanced throughout all steps, as the name of the algorithm implies.

Proof. To prove this result, it suffices to consider the case where $|G^t| = 1$, i.e., batches of one item. We denote by g^t the item we are allocating at time t. Let $t \in [T]$, we will prove that the allocation $(A^t, \dots, A^T)$ is EF1. We use backward induction to establish this proof.

For $t = T$, only one item is being allocated, and the allocation is clearly EF1. Now let $t < T$. Suppose, without loss of generality, that agent a_1 was selected to receive the item (line 6). Let $k \neq 1$. We will prove that agent a_k does not envy agent a_1 by more than one item.

We consider two cases. If $a_k \notin D$ (as defined in line 5), then $v_k(g^t) = 0$ and the result follows from the induction hypothesis.

Otherwise, $v_k(g^t) = 1$. By the selection rule (line 6), agent 1's valuation of their own allocation from round $t+1$ is not as large as agent k's valuation of their own allocation:

$$U_1(A_1^{t+1}, \dots A_1^T) \leq U_k(A_k^{t+1} \dots A_k^T) \tag{1}$$

Furthermore, we argue that:

$$U_k(A_1^{t+1}, \dots A_1^T) \leq U_1(A_1^{t+1}, \dots A_1^T) \tag{2}$$

This is because $A_1^i \neq \emptyset$ only if $v_1(g^i) = 1$. In this sense, U_1 is the utility function that attains the maximum utility from the allocation $(A_1^{t+1}, \dots A_1^T)$. Formally, if the non-empty allocations among $(A_1^{t+1}, \dots A_1^T)$ are $A_1^{t_1}, \dots A_1^{t_m}$ then:

$$U_1(A_1^{t+1}, \dots A_1^T) = \sum_{i=1}^{m} \delta^{t_i - t} = \sum_{i=1}^{m} \delta^{t_i - t} v_1(g^{t_i}) \geq \sum_{i=1}^{m} \delta^{t_i - t} v_k(g^{t_i}) = U_k(A_1^{t+1}, \dots A_1^T)$$

Combining equations (1) and (2) we find that at the start of round t, agent k does not envy agent 1. Consequently, by the end of round t, agent k does not envy agent 1 by more than one item. □

7 Conclusion

In this paper, we set out to address the problem of fairly allocating indivisible goods in an offline, multi-round setting where agents discount future valuations and forget about past allocations. In our setting, goods arrive in predetermined batches, and the primary goal is to guarantee envy-freeness up to one good (EF1) at the end of each round.

Our key findings include the impossibility of achieving EF1 without adjustments in general settings, the development of an algorithm that ensures EF1 for two agents, a bound on the number of adjustments needed in the general case, and the construction of a fair algorithm for the case of binary valuations. These results provide important insights into the limitations and possibilities for fair allocation in such settings.

The broader implications of our research extend to various real-world applications, such as resource distribution and task allocation, where fairness is crucial. Our findings contribute to a deeper understanding of the underlying dynamics and complexities of multi-round allocation problems.

References

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

[2] Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. Online fair division: Analysing a food bank problem. *arXiv preprint arXiv:1502.07571*, 2015.

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

[4] Jiafan He, Ariel D Procaccia, CA Psomas, and David Zeng. Achieving a fairer future by changing the past. *IJCAI'19*, 2019.

A Proof of Theorem 1

Lemma 5. *Consider the case where $n = 3$. If $\mathscr{F}$ is a mechanism that ensures EF1 at every round, then for every sequence of goods $(G^1, \dots, G^T)$, $\delta \in (0, 1]$ and valuation functions (v_1, v_2, v_3), $\mathscr{F}$ cannot produce an allocation with an envy-cycle.*

Proof. If an envy-cycle exists with respect to $(G^1, \dots, G^T)$ and (v_1, v_2, v_3), we introduce a new batch of items $G^0 = \{g\}$ at the beginning such that $v_i(g) > 0$ for all agents. Then, if $\mathscr{F}$ generates an envy-cycle from round $t = 1$, G^0 cannot be allocated without violating EF1. □

Next, we invoke the following statement by He et al. [4]:

Theorem 6 (He et al. [4] Theorem 4.2). *For the case of $n = 3$ and $\delta = 1$, there exists a sequence of 16 batches $(G^1, \dots, G^{16})$ and valuation functions (v_1, v_2, v_3) such that $\mathscr{F}$ generates an envy-cycle.*

The original claim was presented for the case of accumulated utility without future-looking. The proof can be adjusted for our settings by reversing the order of items. The number of batches needed can be reduced from 22 in the original paper to 16 by combining items $1 - 3$, $8 - 10$, and $17 - 19$ into three three-item batches. The proof can then be generalized to any batch size by adding dummy items of negligible valuation and to every number of agents by adding dummy agents. Finally, for a general discount factor δ, one can adjust the valuations of the the items in G^t by multiplying by δ^{-t}.

B Proof of Theorem 2

proof Theorem 2. Let $t \in [T]$. Then our goal is to show that the allocation $(A^t, \dots, A^T)$ is EF1. Denote the temporary and permanent baskets at the end of round t as D^t and P^t, respectively, and the "candidate" allocation at round $t \leq T$ as $C^t := D^t \cup P^t$. There exists a round $t < s \leq T + 1$ for which the candidate allocation C^s is envy-free (we consider round $T + 1$ to be the empty allocation). By definition, the final allocation at round t, A^t, is equal to $(P_1^t \cup D_1^t, P_2^t \cup D_2^t)$ if no swap occurred in rounds 1 to $t - 1$, or $(P_1^t \cup D_2^t, P_2^t \cup D_1^t)$ otherwise. Due to the choice of s, $P^t = P^s$. Let M_1 and M_2 denote the value of the most valuable items for player 2 and player 1, respectively, between rounds t to s:

$$M_1 := \max\{v_1(g) \mid g \in \cup_{t'=t}^{s-1} A_2^{t'}\} \qquad M_2 := \max\{v_2(g) \mid g \in \cup_{t'=t}^{s-1} A_1^{t'}\} \tag{3}$$

Notice that:

$$\begin{aligned} U_1(A_1^t, \dots A_1^T) &= \sum_{t'=t}^{T} \sum_{g \in A_1^{t'}} \delta^{t'-t} v_1(g) = \sum_{t'=t}^{s-1} \sum_{g \in A_1^{t'}} \delta^{t'-t} v_1(g) + \delta^{s-t} \sum_{t'=s}^{T} \sum_{g \in A_1^{t'}} \delta^{t'-s} v_1(g) \\ &= U_1(A_1^t, .., A_1^{s-1}) + \delta^{s-t} U_1(A_1^s, .., A_1^T) \geq U_1(A_1^t, .., A_1^{s-1}) + \delta^{s-t} U_1(A_2^s, .., A_2^T) \end{aligned}$$

and similarly for the second player. It is therefore suffices, by the definition of EF1, to show that:

$$U_1(A_2^t \dots A_2^{s-1}) - U_1(A_1^t, \dots A_1^{s-1}) \leq M_1 \tag{4}$$

$$U_2(A_1^t \dots A_1^{s-1}) - U_2(A_2^t, \dots A_2^{s-1}) \leq M_2 \tag{5}$$

because combining the above two equations with equation 3, results in

$$U_1(A_1^t, \dots A_1^T) \geq U_1(A_2^t, \dots A_2^T) - M_1$$

(and similarly for the second player). Without loss of generality, it suffices to show (4), and the result follows by symmetry.

Although the final allocation at any time t', $A^{t'}$, is determined only at the end of the for loop (line 2), we can still argue that:

$$\bigcup_{t'=t}^{s-1} A^{t'} \in \{(D_1^{t'}, D_2^{t'}), (D_2^{t'}, D_1^{t'})\}$$

This is because no swap occurred between rounds t to s (otherwise, there would be a candidate allocation that is envy-free before round s). Hence, it suffices to discard the permanent baskets and prove that:

$$|U_1(D_2^t \dots D_2^{s-1}) - U_1(D_1^t, \dots D_1^{s-1})| \leq M_1 \tag{6}$$

We prove (6) by backward induction on t. If $t = s - 1$, the Round Robin algorithm ensures EF1, and in particular, equation (6). Now suppose $t < s - 1$. By the inductive hypothesis, we know that for the same s and some $M_1' \leq M_1$, it holds that

$$|U_1(D_2^{t+1} \dots D_2^{s-1}) - U_1(D_1^{t+1}, \dots D_1^{s-1})| \leq M_1'$$

Consider the following two cases: if $U_1(D_2^{t+1} \dots D_2^{s-1}) - U_1(D_1^{t+1}, \dots D_1^{s-1}) > 0$, a_2 is envied and line 3 is True, and a_1 chooses first. The value of a_1 for the most valuable item in G^t is at most M_1 and so by additivity:

$$U_1(D_2^t \dots D_2^{s-1}) - U_1(D_1^t, \dots D_1^{s-1}) = (U_1(D_2^t) - U_1(D_1^t)) + \delta\left(U_1(D_2^{t+1} \dots D_2^{s-1}) - U_1(D_1^{t+1}, \dots D_1^{s-1})\right)$$

The right hand size is smaller than $0 + \delta M_1' \leq M_1$, and greater than $-M_1 + 0$. In either case, equation (6) holds.

The other case is $U_1(D_2^{t+1} \dots D_2^{s-1}) - U_1(D_1^{t+1}, \dots D_1^{s-1}) \leq 0$. In this case, a_1 does not envy, and a_2 chooses first. In this case, since the maximum envy that could occur at round t is M_1, clearly equation (6) holds.

□