

Round-Robin Beyond Additive Agents: Existence and Fairness of Approximate Equilibria

GEORGIOS AMANATIDIS, University of Essex, UK and Archimedes / Athena RC, Greece

GEORGIOS BIRMPAS, Sapienza University of Rome, Italy

PHILIP LAZOS, Input Output, UK

STEFANO LEONARDI, Sapienza University of Rome, Italy

REBECCA REIFFENHÄUSER, University of Amsterdam, The Netherlands

Fair allocation of indivisible goods has attracted extensive attention over the last two decades, yielding numerous elegant algorithmic results and producing challenging open questions. The problem becomes much harder in the presence of *strategic* agents. Ideally, one would want to design *truthful* mechanisms that produce allocations with fairness guarantees. However, in the standard setting without monetary transfers, it is generally impossible to have truthful mechanisms that provide non-trivial fairness guarantees. Recently, Amanatidis et al. [2021] suggested the study of mechanisms that produce fair allocations in their equilibria. Specifically, when the agents have additive valuation functions, the simple Round-Robin algorithm always has pure Nash equilibria and the corresponding allocations are *envy-free up to one good* (EF1) with respect to the agents' *true valuation functions*. Following this agenda, we show that this outstanding property of the Round-Robin mechanism extends much beyond the above default assumption of additivity. In particular, we prove that for agents with *cancelable* valuation functions (a natural class that contains, e.g., additive and budget-additive functions), this simple mechanism always has equilibria and even its approximate equilibria correspond to approximately EF1 allocations with respect to the agents' true valuation functions. Further, we show that the approximate EF1 fairness of approximate equilibria surprisingly holds for the important class of *submodular* valuation functions as well, even though exact equilibria fail to exist!

CCS Concepts: • **Theory of computation** → **Algorithmic game theory and mechanism design**; *Approximation algorithms analysis*; • **Mathematics of computing** → Discrete mathematics.

Additional Key Words and Phrases: fair division, Nash equilibria, round-robin, envy-freeness up to one good

ACM Reference Format:

Georgios Amanatidis, Georgios Birmas, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. 2023. Round-Robin Beyond Additive Agents: Existence and Fairness of Approximate Equilibria. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23)*, July 9–12, 2023, London, United Kingdom. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3580507.3597796>

1 INTRODUCTION

Fair division refers to the problem of dividing a set of resources among a group of agents in a way that every agent feels they have received a “fair” share. The mathematical study of (a continuous version of) the problem dates back to the work of Banach, Knaster, and Steinhaus [1949], who, in a first attempt to formalize fairness, introduced the notion of *proportionality*, i.e., each of the n agents receives at least $1/n$ -th of the total value from her perspective. Since then, different variants of the problem have been studied in mathematics, economics, political science, and computer science, and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '23, July 9–12, 2023, London, United Kingdom

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0104-7/23/07...\$15.00

<https://doi.org/10.1145/3580507.3597796>

various fairness notions have been defined. The most prominent fairness notion is *envy-freeness* [Foley 1967; Varian 1974], where each agent values her set of resources at least as much as the set of any other agent. When the available resources are *indivisible* items, i.e., items that cannot be split among agents, notions introduced for infinitely divisible resources, like proportionality and envy-freeness are impossible to satisfy, even approximately. In the last two decades fair allocation of indivisible items has attracted extensive attention, especially within the theoretical computer science community, yielding numerous elegant algorithmic results for various new fairness notions tailored to this discrete version of the problem, such as *envy-freeness up to one good* (EF1) [Budish 2011; Lipton et al. 2004], *envy-freeness up to any good* (EFX) [Caragiannis et al. 2019], and *maximin share fairness* (MMS) [Budish 2011]. We refer the interested reader to the surveys of Amanatidis et al. [2022]; Bouveret et al. [2016]; Procaccia [2016].

In this work, we study the problem of fairly allocating indivisible *goods*, i.e., items of non-negative value, to *strategic* agents, i.e., agents who might misreport their private information if they have an incentive to do so. Incentivising strategic agents to truthfully report their valuations is a central goal—and often a notorious challenge—in mechanism design, in general. Specifically in fair division, this seems particularly necessary, since any fairness guarantee on the outcome of a mechanism typically holds with respect to its input, namely the *reported* preferences of the agents rather than their true, private preferences which they may have chosen not to reveal. Without truthfulness, fairness guarantees seem to become meaningless. Unfortunately, when monetary transfers are not allowed, as is the standard assumption in fair division, such *truthful* mechanisms fail to exist for any meaningful notion of fairness, even for simple settings with two agents who have additive valuation functions [Amanatidis et al. 2017a].

As an alternative, Amanatidis et al. [2021] initiated the study of *equilibrium fairness*: when a mechanism always exhibits stable (i.e., pure Nash equilibrium) states, each of which corresponds to a fair allocation with respect to the *true* valuation functions, the need for extracting agents’ true preferences is mitigated. Surprisingly, they show that for the standard case of additive valuation functions, the simple *Round-Robin* routine is such a mechanism with respect to EF1 fairness. Round-Robin takes as input an ordering of the goods for each agent, and then cycles through the agents and allocates the goods one by one, giving to each agent their most preferred available good. For agents with additive valuation functions, Round-Robin is known to produce EF1 allocations (see, e.g., [Markakis 2017]). Note that, without monetary transfers, what distinguishes a mechanism from an algorithm is that its input is the, possibly misreported, agents’ preferences.

To further explore the interplay between incentives and fairness, we take a step back and focus solely on this very simple, yet fundamental, allocation protocol. It should be noted that the Round-Robin algorithm is one of the very few fundamental procedures one can encounter throughout the discrete fair division literature. Its central role is illustrated by various prominent results, besides producing EF1 allocations: it can be modified to produce approximate MMS allocations [Amanatidis et al. 2017b], as well as EF1 allocations for *mixed goods and chores* (i.e., items with negative value) [Aziz et al. 2022]. It produces *envy-free* allocations with high probability when the values are drawn from distributions [Manurangsi and Suksompong 2021], it is used to produce a “nice” initial allocation as a subroutine in the state-of-the-art approximation algorithms for *pairwise maximin share fair* (PMMS) allocations [Kurokawa 2017] and EFX allocations [Amanatidis et al. 2020], it has the lowest communication complexity of any known fair division algorithm, and, most relevant to this work, it is the *only* algorithm for producing fair allocations for more than two agents that, when viewed as a mechanism, is known to even have equilibria [Aziz et al. 2017b].

We investigate the existence and the EF1 guarantees of approximate pure Nash equilibria of the Round-Robin mechanism beyond additive valuation functions, i.e., when the goods already assigned to an agent potentially change how they value the remaining goods. In particular, we are

interested in whether anything can be said about classes that largely generalize additive functions, like *cancelable* functions, i.e., functions where the marginal values with respect to any subset maintain the relative ordering of the goods, and *submodular* functions, i.e., functions capturing the notion of diminishing returns. Although the stability and equilibrium fairness properties of Round-Robin have been visited before [Amanatidis et al. 2021; Aziz et al. 2017b], to the best of our knowledge, we are the first to study the problem for non-additive valuation functions and go beyond exact pure Nash equilibria. Cancelable functions also generalize budget-additive, unit-demand, and multiplicative valuation functions [Berger et al. 2022], and recently have been of interest in the fair division literature as several results can be extended to this class [Akrami et al. 2022; Berger et al. 2022; Caragiannis et al. 2022]. For similar reasons, cancelable functions seem to be a good pairing with Round-Robin as well, at least in the algorithmic setting (see, e.g., Proposition 2.5).

Nevertheless, non-additive functions seem to be massively harder to analyze in our setting and come with various obstacles. First, it is immediately clear that, even without strategic agents, the input of an ordinal mechanism implemented as a simultaneous-move one-shot game, like the Round-Robin mechanism we study here, can no longer capture the complexity of a submodular function (see also the relevant discussion in Our Contributions). As a result, translating this sequential assignment to an estimate on the value of each agent’s *bundle* of goods, is not obvious. Lastly, and this applies to cancelable functions as well, assuming equilibria do exist and enough can be shown about the value of the assigned bundles to establish fairness, there is no reason to expect that any fairness guarantee will hold with respect to the true valuation functions, as the agents may misreport their preferences in an arbitrary fashion.

1.1 Contribution and Technical Considerations

We study the well-known Round-Robin mechanism (Mechanism 1) for the problem of fairly allocating a set of indivisible goods to a set of strategic agents. We explore the existence of approximate equilibria, along with the fairness guarantees that the corresponding allocations provide with respect to the agents’ true valuation functions. Qualitatively, we generalize the surprising connection between the stable states of this simple mechanism and its fairness properties to all approximate equilibria and for valuation functions as general as subadditive cancelable and submodular. In more detail, our main contributions can be summarized as follows:

- We show that the natural generalization of the *bluff profile* of Aziz et al. [2017b] is an exact PNE that always corresponds to an EF1 allocation, when agents have *cancelable* valuation functions (Theorem 3.2 along with Proposition 2.5). Our proof is simple and intuitive and generalizes the results of Aziz et al. [2017b] and Amanatidis et al. [2021].
- For agents with submodular valuation functions, we show that there are instances where no $(3/4 + \epsilon)$ -approximate PNE exists (Proposition 3.4), thus creating a separation between the cancelable and the submodular cases. Nevertheless, we prove that an appropriate generalization of the bluff profile is a $1/2$ -approximate PNE (Theorem 3.7) that also produces an $1/2$ -EF1 allocation with respect to the true valuation functions (Theorem 3.8).
- We provide a unified proof that connects the factor of an approximate PNE with the fairness approximation factor of the respective allocation. In particular, any α -approximate PNE results in a $\alpha/2$ -EF1 allocation for subadditive cancelable agents (Theorem 4.5), and in a $\alpha/3$ -EF1 allocation for submodular agents (Theorem 4.4). We complete the picture by providing lower bounds in both cases (Theorem 4.3 and Proposition 4.8), which demonstrate that our results are almost tight.

As a minor byproduct of our work, Theorem 3.8 and the definition of the bluff profile imply that, given *value oracles* for the submodular functions, we can use Round-Robin as a subroutine to

produce 1/2-EF1 allocations. This also raises the question of whether one should allow a more expressive bid, e.g., a value oracle. Although this is a viable direction, we avoid it here as it comes with a number of issues. Allowing the input to be exponential in the number of goods is already problematic, especially when simplicity and low communication complexity are two appealing traits of the original mechanism. Moreover, extracting orderings from value oracles would essentially result in a mechanism equivalent to ours (if the ordering of an agent depended only on *her* function) or to a sequential game (if the orderings depended on all the functions) which is not what we want to explore here. Note that less information is not necessarily an advantage towards our goal. While the space of equilibria becomes reacher, fairness guarantees are increasingly harder to achieve.

As a final remark, all the algorithmic procedures we consider run in polynomial time, occasionally assuming access to value oracles, e.g., Algorithms 2, 3, 4. Although we do not consider computational complexity questions here, like how do agents compute best responses or how do they reach approximate equilibria, we do consider such questions interesting directions for future work.

1.2 Further Related Work

The problem of fairly allocating indivisible goods to additive agents in the non-strategic setting has been extensively studied; for a recent survey, see Amanatidis et al. [2022]. Although the additivity of the valuation functions is considered a standard assumption, there are many works that explore richer classes of valuation functions. Some prominent examples include the computation of EF1 allocations for agents with general non-decreasing valuation functions [Lipton et al. 2004], EFX allocations (or relaxations of EFX) under agents with cancelable valuation functions [Akrami et al. 2022; Berger et al. 2022; Caragiannis et al. 2022] and subadditive valuation functions [Chaudhury et al. 2021; Plaut and Roughgarden 2020], respectively, and approximate MMS allocations for submodular, XOS, and subadditive agents [Barman and Krishnamurthy 2020; Ghodsi et al. 2022]. Although Round-Robin has been considered for non-additive agents before, see, e.g., [Bouveret and Lang 2014], to the best of our knowledge, only Montanari et al. [2022] study its fairness guarantees for submodular valuation functions and show that it produces allocations which satisfy *marginal* EF1, a relaxation of EF1 introduced by Caragiannis et al. [2019].

Moving to the strategic setting, Caragiannis et al. [2009], Markakis and Psomas [2011] and Amanatidis et al. [2016] were the first to consider the question of whether it is possible to have mechanisms that are truthful and fair at the same time, again assuming additive agents. Amanatidis et al. [2017a] resolved this question for two agents, showing there is no truthful mechanism with fairness guarantees under any meaningful fairness notion. As a result, subsequent papers considered truthful mechanism design under restricted valuation function classes [Babaioff et al. 2021; Halpern et al. 2020]. The stability of Round-Robin was first studied by Aziz et al. [2017b], who proved that it always has PNE by using a special case of retracted result of Bouveret and Lang [2014] (this did not affect the former though; see [Aziz et al. 2017a]). Finally, besides the work of Amanatidis et al. [2021] mentioned earlier, the fairness properties of Round-Robin under strategic agents have recently been studied by Psomas and Verma [2022]. Therein it is shown that Round-Robin, despite being non-truthful, satisfies a relaxation of truthfulness, as it is *not obviously manipulable*.

2 PRELIMINARIES

For $a \in \mathbb{N}$, let $[a]$ denote the set $\{1, 2, \dots, a\}$. We will use $N = [n]$ to denote the set of agents and $M = \{g_1, \dots, g_m\}$ to denote the set of goods. Each agent $i \in N$ has a valuation function $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$ over the subsets of goods. We assume that all v_i are *normalized*, i.e., $v_i(\emptyset) = 0$. We also adopt the shortcut $v_i(T | S)$ for the *marginal value* of a set T with respect to a set S , i.e., $v_i(T | S) = v_i(T \cup S) - v_i(S)$. If $T = \{g\}$, we write $v_i(g | S)$ instead of $v_i(\{g\} | S)$. For each agent $i \in N$, we say that v_i is

- *non-decreasing* (often referred to as *monotone*), if $v_i(S) \leq v_i(T)$ for any $S \subseteq T \subseteq M$.
- *submodular*, if $v_i(g | S) \geq v_i(g | T)$ for any $S \subseteq T \subseteq M$ and $g \notin T$.
- *cancelable*, if $v_i(S \cup \{g\}) > v_i(T \cup \{g\}) \Rightarrow v_i(S) > v_i(T)$ for any $S, T \subseteq M$ and $g \in M \setminus (S \cup T)$.
- *additive*, if $v_i(S \cup T) = v_i(S) + v_i(T)$ for every $S, T \subseteq M$ with $S \cap T = \emptyset$.
- *subadditive*, if $v_i(S \cup T) \leq v_i(S) + v_i(T)$ for every $S, T \subseteq M$.

Throughout this work, we only consider non-decreasing valuation functions, e.g., when we refer to submodular functions, we mean non-decreasing submodular functions. Note that although both submodular and (subadditive) cancelable functions are strict superclasses of additive functions, neither one is a superclass of the other.

We will occasionally need an alternative characterization of submodular functions due to Nemhauser et al. [1978].

THEOREM 2.1 (NEMHAUSER ET AL. [1978]). *A function $v : 2^M \rightarrow \mathbb{R}_{\geq 0}$ is (non-decreasing) submodular if and only if we have $v(T) \leq v(S) + \sum_{i \in T \setminus S} v(i | S)$, for all $S, T \subseteq M$.*

Also, the following lemma summarizes some easy observations about cancelable functions.

LEMMA 2.2. *If $v : 2^M \rightarrow \mathbb{R}_{\geq 0}$ is cancelable, then $v_i(S \cup R) > v_i(T \cup R) \Rightarrow v_i(S) > v_i(T)$ and, thus, $v_i(S) \geq v_i(T) \Rightarrow v_i(S \cup R) \geq v_i(T \cup R)$, for any $S, T, R \subseteq M$, such that $R \subseteq M \setminus S \cup T$. In particular, $v_i(S) = v_i(T) \Rightarrow v_i(S \cup R) = v_i(T \cup R)$.*

Note that, for $S, T \subseteq M$, Lemma 2.2 directly implies that $\arg \max_{g \in T} v(g) \subseteq \arg \max_{g \in T} v(g | S)$.

Despite the fact that the agents have valuation functions, the mechanism we study (Mechanism 1) is *ordinal*, i.e., it only takes as input a *preference ranking* from each agent. Formally, the preference ranking $\succ_i$, which agent i reports, defines a total order on M , i.e., $g \succ_i g'$ implies that good g precedes good g' in agent i ' declared preference ranking.¹ We call the vector of the agents' declared preference rankings, $\succ = (\succ_1, \dots, \succ_n)$, the *reported profile* for the instance. So, while an instance to our problem is an ordered triple $(N, M, \mathbf{v})$, where $\mathbf{v} = (v_1, \dots, v_n)$ is a vector of the agents' valuation functions, the input to Mechanism 1 is $(N, M, \succ)$ instead.

Note that $\succ_i$ may not reflect the actual underlying values, i.e., $g \succ_i g'$ does not necessarily mean that $v_i(g) > v_i(g')$ or, more generally, $v_i(g | S) > v_i(g' | S)$ for a given $S \subseteq M$. This might be due to agent i misreporting her preference ranking, or due to the fact that any single preference ranking is not expressive enough to fully capture all the partial orders induced by a submodular function. Nevertheless, a valuation function v_i does induce a *true preference ranking* $\succ_{i|S}^*$ for each set $S \subseteq M$, which is a partial order, i.e., $g \succ_{i|S}^* g' \Leftrightarrow v_i(g | S) \geq v_i(g' | S)$ for all $g, g' \in M$. We use $\succ_{i|S}^*$ if the corresponding preference ranking is *strict*, i.e., when $g \succ_{i|S}^* g' \wedge g' \succ_{i|S}^* g \Rightarrow g = g'$, for all $g, g' \in M \setminus S$. For additive (and more generally, for cancelable) valuations, we drop S for the notation and simply write $\succ_i^*$ or $\succ_i$. Finally, for a total order $\succ$ on M and a set $T \subseteq M$, we use $\text{top}(\succ, T)$ to denote the “largest” element of T with respect to $\succ$.

2.1 Fairness Notions

A fair division mechanism produces an *allocation* $(A_1, \dots, A_n)$, where A_i is the *bundle* of agent i , which is a partition of M . The latter corresponds to assuming no free disposal, namely all the goods must be allocated.

There are several different notions which attempt to capture which allocations are “fair”. The most prominent such notion in the fair division literature has been *envy-freeness* (EF) [Foley 1967; Varian 1974], which has been the starting point for other relaxed notions, more appropriate for the

¹See the discussion after the statement of Mechanism 1 about why assuming that the reported preference rankings are total (rather than partial) orders is without loss of generality.

indivisible goods setting we study here, as *envy-freeness up to one good* (EF1) [Budish 2011; Lipton et al. 2004] and *envy-freeness up to any good* (EFX) [Caragiannis et al. 2019]. Here we focus on EF1.

Definition 2.3. An allocation $(A_1, \dots, A_n)$ is

- α -*envy-free* (α -EF), if for every $i, j \in N$, $v_i(A_i) \geq \alpha \cdot v_i(A_j)$.
- α -*envy-free up to one good* (α -EF1), if for every pair of agents $i, j \in N$, with $A_j \neq \emptyset$, there exists a good $g \in A_j$, such that $v_i(A_i) \geq \alpha \cdot v_i(A_j \setminus \{g\})$.

When for every agent $j \in N$ with $A_j \neq \emptyset$, we have $v_i(A_i) \geq \alpha \cdot v_i(A_j \setminus \{g\})$ for some good $g \in A_j$, we say that $(A_1, \dots, A_n)$ is α -EF1 *from agent i 's perspective*, even when the allocation is not α -EF1!

2.2 Mechanisms and Equilibria

We are interested in *mechanisms* that produce allocations with EF1 guarantees. When *no payments* are allowed, like in our setting, an allocation mechanism $\mathcal{M}$ is just an allocation algorithm that takes as input the agents' reported preferences. In particular, Round-Robin, the mechanism of interest here, takes as input the reported profile $\succ$ and produces an allocation of all the goods. This distinction in terminology is necessary as the reported input may not be consistent with the actual valuation functions due to the agents' incentives. When the allocation returned by $\mathcal{M}(\succ)$ has some fairness guarantee, e.g., it is 0.5-EF1, we will attribute the same guarantee to the reported profile itself, i.e., we will say that $\succ$ is 0.5-EF1.

We study the fairness guarantees of the (approximate) pure Nash equilibria of Round-Robin. As is common, given a preference profile $\succ = (\succ_1, \dots, \succ_n)$, we write $\succ_{-i}$ to denote $(\succ_1, \dots, \succ_{i-1}, \succ_{i+1}, \dots, \succ_n)$ and, given a preference ranking $\succ'_i$, we use $(\succ'_i, \succ_{-i})$ to denote the profile $(\succ_1, \dots, \succ_{i-1}, \succ'_i, \succ_{i+1}, \dots, \succ_n)$. For the next definition we abuse the notation slightly: given an allocation $(A_1, \dots, A_n)$ produced by $\mathcal{M}(\succ)$, we write $v_i(\mathcal{M}(\succ))$ to denote $v_i(A_i)$; similarly for $\mathcal{M}(\succ'_i, \succ_{-i})$.

Definition 2.4. Let $\mathcal{M}$ be an allocation mechanism and consider a preference profile $\succ = (\succ_1, \dots, \succ_n)$. We say that the total order $\succ_i$ is an α -*approximate best response* to $\succ_{-i}$ if for every total order, i.e., permutation $\succ'_i$ of M , we have $\alpha \cdot v_i(\mathcal{M}(\succ'_i, \succ_{-i})) \leq v_i(\mathcal{M}(\succ))$. The profile $\succ$ is an α -*approximate pure Nash equilibrium* (PNE) if, for each $i \in N$, $\succ_i$ is an α -approximate best response to $\succ_{-i}$.

When $\alpha = 1$, we simply refer to best responses and exact PNE.

2.3 The Round-Robin Mechanism

We state Round-Robin as a mechanism (Mechanism 1) that takes as input a reported profile $(\succ_1, \dots, \succ_n)$. For the sake of presentation, we assume that the agents in each *round* (lines 3–6) are always considered according to their “name”, i.e., agent 1 is considered first, agent 2 second, and so on, instead of having a permutation determining the priority of the agents as an extra argument of the input. This is without loss of generality, as it only requires renaming the agents accordingly. We often refer to the process of allocating a good to an agent (lines 4–6) as a *step* of the mechanism.

Note that there is no need for a tie-breaking rule here, as the reported preference rankings are assumed to be total orders. Equivalently, one could allow for partial orders (either directly or via cardinal bids as it is done in [Amanatidis et al. 2021]) paired with a deterministic tie-breaking rule, e.g., lexicographic tie-breaking, a priori known to the agents.

In the rest of the paper, we will assume that $m = kn$ for some $k \in \mathbb{N}$, for simplicity. Note that this is without loss of generality, as we may introduce at most $n - 1$ dummy goods that have marginal value of 0 with respect to any set for everyone and append them at the end of the reported preference rankings to be allocated during the last steps of the mechanism.

Mechanism 1 Round-Robin($\succ_1, \dots, \succ_n$) // For $i \in N$, $\succ_i$ is the reported preference ranking of agent i .

```

1:  $S = M$ ;  $(A_1, \dots, A_n) = (\emptyset, \dots, \emptyset)$ ;  $k = \lceil m/n \rceil$ 
2: for  $r = 1, \dots, k$  do                                // Each value of  $r$  determines the corresponding round.
3:   for  $i = 1, \dots, n$  do                                // The combination of  $r$  and  $i$  determines the corresponding step.
4:      $g = \text{top}(\succ_i, S)$ 
5:      $A_i = A_i \cup \{g\}$                                 // The current agent receives (what appears to be) her favorite available good.
6:      $S = S \setminus \{g\}$                                 // The good is no longer available.
7: return  $(A_1, \dots, A_n)$ 

```

We have already mentioned that Round-Robin as an algorithm produces EF1 allocations for additive agents, where the input is assumed to be any strict variant $\succ^* = (\succ_{1|\emptyset}^*, \succ_{2|\emptyset}^*, \dots, \succ_{n|\emptyset}^*)$ of the truthful profile $(\succ_{1|\emptyset}^*, \succ_{2|\emptyset}^*, \dots, \succ_{n|\emptyset}^*)$, i.e., the profile where each agent ranks the goods according to their singleton value. This property fully extends to cancelable valuation functions as well.

PROPOSITION 2.5. *Let $\succ^*$ be as described above. When all agents have cancelable valuation functions, the allocation returned by Round-Robin($\succ^*$) is EF1.*

The proof of Proposition 2.5 is rather simple, but as it requires Lemma 3.3 from the next section, we defer it to the full version of our paper [Amantidis et al. 2023]. All missing proofs, due to space constraints, are also included there.

3 EXISTENCE OF APPROXIMATE PNE

At first glance, it is not clear why Mechanism 1 has any pure Nash equilibria, even approximate ones for a constant approximation factor. For additive valuation functions, however, it is known that for any instance we can construct a simple preference profile, called the *bluff profile*, which is an exact PNE. While the proof of this fact, in its full generality, is fragmented over three papers [Amantidis et al. 2021; Aziz et al. 2017b; Bouveret and Lemaître 2016], we give here a simple proof that generalizes the existence of exact PNE to cancelable valuation functions. As we shall see later, extending this result to submodular functions is not possible and even defining a generalization of the bluff profile which is a 0.5-approximate PNE is not straightforward.

3.1 Cancelable valuations

Defining the bluff profile for cancelable agents, we will start from a strict variant of the truthful profile $(\succ_{1|\emptyset}^*, \succ_{2|\emptyset}^*, \dots, \succ_{n|\emptyset}^*)$, i.e., the profile where each agent ranks the goods according to their value (as singletons) in descending order, as we did for Proposition 2.5. Assume that any ties are broken deterministically to get the strict version $\succ^* = (\succ_{1|\emptyset}^*, \succ_{2|\emptyset}^*, \dots, \succ_{n|\emptyset}^*)$. Now, consider Round-Robin($\succ^*$) and let $h_1, h_2, \dots, h_m$ be a renaming of the goods according to the order in which they were allocated and $\succ^b$ be the corresponding total order (i.e., $h_1 \succ^b h_2 \succ^b \dots \succ^b h_m$). The *bluff profile* is the preference profile $\succ^b = (\succ_1^b, \succ_2^b, \dots, \succ_n^b)$, where everyone ranks the goods in the order they were allocated in Round-Robin($\succ^*$). The following fact follows directly from the definition of the bluff profile and the description of Round-Robin.

FACT 3.1. *If $\succ^*$ is a strict version of the truthful preference profile and $\succ^b$ is the corresponding bluff profile, then Round-Robin($\succ^b$) and Round-Robin($\succ^*$) both return the same allocation.*

An interesting observation about this fact is that, combined with Proposition 2.5 and Theorem 3.2, it implies that there is at least one PNE of Mechanism 1 which is EF1! Of course, it is now known that all exact PNE of Round-Robin are EF1 for agents with *additive* valuation functions and,

as we will see later on, even approximate PNE have (approximate) EF1 guarantees for much more general instances, including the case of *subadditive cancelable* valuation functions.

THEOREM 3.2. *When all agents have cancelable valuation functions, the bluff profile is an exact PNE of Mechanism 1.*

We first need to prove the following lemma that generalizes a straightforward property of additive functions for cancelable functions.

LEMMA 3.3. *Suppose that $v(\cdot)$ is a cancelable valuation function. Consider sets $X = \{x_1, x_2, \dots, x_k\}$ and $Y = \{y_1, y_2, \dots, y_k\}$. If for every $j \in [k]$, we have that $v(x_j) \geq v(y_j)$, then $v(X) \geq v(Y)$.*

PROOF. We begin by arguing that it is without loss of generality to first assume that the elements of X are ordered by non-increasing value with respect to v and then also assume that $y_j \notin \{x_1, x_2, \dots, x_{j-1}\}$, for any $j \in [k]$. The former is indeed a matter of reindexing, if necessary, the elements of X and consistently reindexing the corresponding elements of Y . For the latter, suppose that there exist j such that $y_j = x_t$ for $t \leq j-1$ and consider the smallest t for which this happens. We have $v(x_t) \geq v(x_{t+1}) \geq \dots \geq v(x_j)$ by the assumption on the ordering of the elements of X , $v(x_j) \geq v(y_j)$ by hypothesis, and $v(y_j) = v(x_t)$. Thus, $v(x_t) = v(x_{t+1}) = \dots = v(x_j)$. Now we may rename the elements of Y to $\{y'_1, \dots, y'_k\}$ by inserting y_j to the t -th position, i.e., $y'_t = y_j$, $y'_s = y_{s-1}$, for $t+1 \leq s \leq j$, and $y'_s = y_s$, for $s < t$ or $s > j$. Since only $y_t, y_{t+1}, \dots, y_j$ changed indices but $v(x_t) = v(x_{t+1}) = \dots = v(x_j)$, we again have that $v(x_j) \geq v(y'_j)$ for every $j \in [k]$. Moreover, now the smallest ℓ for which there exist $j > \ell$ such that $y_j = x_\ell$ is strictly larger than t . By repeating this renaming of the elements of Y we end up with a renaming $\{y_1^*, \dots, y_k^*\}$ such that for every $j \in [k]$, $v(x_j) \geq v(y_j^*)$ and $y_j^* \notin \{x_1, x_2, \dots, x_{j-1}\}$.

So, assuming that the elements of X are ordered in non-increasing value with respect to v and that $y_j \notin \{x_1, x_2, \dots, x_{j-1}\}$, for any $j \in [k]$, suppose towards a contradiction that $v(X) < v(Y)$. That is, $v(\{x_1, x_2, \dots, x_k\}) < v(\{y_1, y_2, \dots, y_k\})$. Now notice that if $v(\{x_1, x_2, \dots, x_{k-1}\}) \geq v(\{y_1, y_2, \dots, y_{k-1}\})$, this would imply that $v(\{x_1, \dots, x_{k-1}, y_k\}) \geq v(\{y_1, \dots, y_{k-1}, y_k\})$, by the definition of cancelable valuations and the fact that $y_k \notin \{x_1, \dots, x_{k-1}\} \cup \{y_1, \dots, y_{k-1}\}$. This leads to $v(\{x_1, \dots, x_{k-1}, x_k\}) \geq v(\{x_1, \dots, x_{k-1}, y_k\}) \geq v(\{y_1, \dots, y_{k-1}, y_k\})$, where the first inequality follows from $v(x_k) \geq v(y_k)$ and Fact 2.2, something that contradicts our initial assumption. Therefore, $v(\{x_1, \dots, x_{k-1}\}) < v(\{y_1, \dots, y_{k-1}\})$. By repeating the same argument $k-2$ more times, we end up with $v(x_1) < v(y_1)$, a contradiction. $\square$

PROOF OF THEOREM 3.2. Now we show that the bluff profile for cancelable valuations is an exact PNE. Consider the goods named $h_1, \dots, h_m$ as in the bluff profile, i.e., by the order in which they are picked when each agent reports their preference order to be the one induced by all singleton good values. Consider agent i . Her assigned set of goods under the bluff profile is $A_i^b = \{h_i, h_{n+i}, \dots, h_{(k-1)n+i}\}$, where $k = m/n$. Assume now that she deviates from $\succ^b$ to $\succ_i$, resulting in some allocated set $A_i = \{y_1, y_2, \dots, y_k\}$, where we assume y_r to be allocated in round r . We need to show $v_i(A_i^b) \geq v_i(A_i)$.

To this end, we compare the goods allocated to agent i in both reports, one by one. If $v_i(y_r) \leq v_i(h_{(r-1)n+i})$ for every $r \in [k]$, then we are done by applying Lemma 3.3 with A_i^b and A_i . If some of these inequalities fail, let r denote the latest round such that $v_i(y_r) > v_i(h_{(r-1)n+i})$. Therefore, in the execution of Mechanism 1 with the bluff profile as input, y_r was no longer available in round r . However, y_r becomes available in round r once agent i deviates. This can only stem from the fact that at some point before round r , a good h_t with $t > (r-1)n+i$ was picked (since the overall number of goods picked per round always stays the same). Clearly, the only agent who could have done so (since she is the only one deviating from the common bluff order) is agent i .

Therefore, it holds that $h_t = y_j$ for some $j < r$. Now, we replace the ordered set $Y = (y_1, y_2, \dots, y_k)$ by $Y' = (y_1, \dots, y_{j-1}, y_r, y_{j+1}, \dots, y_{r-1}, y_j, y_{r+1}, \dots, y_k)$, i.e., we simply exchange y_r and y_j . It will be convenient to rename $y_1, \dots, y_k$ so that $Y' = (y'_1, y'_2, \dots, y'_k)$.

We claim that if agent i reports a preference ranking $\succ'_i$ that starts with all goods in Y' , in that specific order, followed by everything else, in any order, she still gets A_i but the goods are allocated in the order suggested by Y' . Indeed, first notice that the first $j - 1$ rounds of Round-Robin will be the same as in the run with the original deviation $\succ_i$. Further, $y'_j = y_r$ is allocated earlier under $\succ'_i$ than under $\succ_i$, and thus it surely is available at the time. After that, rounds $j - 1$ to $r - 1$ will be the same as in the run with the deviation $\succ_i$. Now $y'_r = y_j$ is allocated later than before, namely in round r , but it is not among the first $(r - 1)n + i$ goods in the bluff order, as noted above, which means it is not allocated to any other agent in any round before the r -th under $\succ'_i$. Finally, rounds $r + 1$ to k will be the same as in the run with $\succ_i$.

Although agent i still is assigned the same set A_i by deviating to $\succ'_i$, we now have $v_i(y'_r) = v_i(y_j) \leq v_i(h_{(r-1)n+i})$, where the inequality holds because both goods are available in round r of the bluff run, and agent one prefers $h_{(r-1)n+i}$. Also, all later goods in Y' remain unchanged, i.e., $y'_s = y_s$ for $s > r$. Therefore, the latest occurrence of some $y'_t > h_{(r-1)n+i}$ now happens at an earlier point in the sequence, if at all. Repeating this process until no such occurrence is left yields an ordering $Y^* = (y^*_1, y^*_2, \dots, y^*_k)$ of A_i such that for all $r \in [k]$, $v_i(y^*_r) \leq v_i(h_{(r-1)n+i})$. Now using Lemma 3.3 completes the proof. $\square$

3.2 Submodular valuations

We move on to the much more general class of submodular valuations. In order to define the bluff profile here, we again would like to start from the truthful profile. However, recall that Round-Robin restricts each agent's report to specifying an ordering on the good set M and these preference rankings are not expressive enough to fully capture submodular valuation functions. In fact, it is not obvious what 'truthful' means here without further assumptions on what information is known by the agents. Still, we define a *truthfully greedy* allocation and use this as our starting point.

Imagine that, instead of having a full preference profile from the beginning, we only ask the active agent i (i.e., the agent to which we are about to allocate a new good) for the good with the largest marginal value with respect to her current set of goods A_i and give this to her. Let $h_1, h_2, \dots, h_m$ be a renaming of the goods according to the order in which they would be allocated in this hypothetical truthfully greedy scenario and $\succ^b$ be the corresponding total order. Like in the cancelable case, the bluff profile is the preference profile $\succ^b = (\succ^b_1, \succ^b_2, \dots, \succ^b_n)$.

Formally, the renaming of the goods is performed as described in Algorithm 2 below. It should be noted that this definition of the bluff profile is consistent with the definition for cancelable functions, assuming that all ties are resolved lexicographically.

Algorithm 2 Greedy renaming of goods for defining the bluff profile

Input: N, M , value oracles for $v_1(\cdot), \dots, v_n(\cdot)$

- 1: $X_i = \emptyset$ for $i \in [n]$
 - 2: **for** $j = 1, \dots, m$ **do**
 - 3: $i = (j - 1) \pmod{n} + 1$
 - 4: $h_j = \arg \max_{g \in M \setminus \bigcup_{\ell < j} X_\ell} v_i(g \mid X_i)$ // Ties are broken lexicographically.
 - 5: $X_i = X_i \cup \{h_j\}$
 - 6: **return** $(h_1, h_2, \dots, h_m)$
-

Also notice that the allocation $\text{Round-Robin}(\succ^b)$ produced under the bluff profile is exactly $(X_1, X_2, \dots, X_n)$, as described in Algorithm 2, i.e., $X_i = A_i^b = \{h_i, h_{n+i}, \dots, h_{(k-1)n+i}\}$, where recall that $k = m/n$.

The main result of this section is Theorem 3.7 stating that the bluff profile is a $\frac{1}{2}$ -approximate PNE when agents have submodular valuation functions. While this sounds weaker than Theorem 3.2, it should be noted that for submodular agents Mechanism 1 does not have PNE in general, even for relatively simple instances, as stated in Proposition 3.4. In fact, even the existence of approximate equilibria can be seen as rather surprising, given the generality of the underlying valuation functions.

PROPOSITION 3.4. *There exists an instance where all agents have submodular valuation functions such that Mechanism 1 has no $(\frac{3}{4} + \varepsilon)$ -approximate PNE.*

Moving towards the proof of Theorem 3.7 for the submodular case, we note that although it is very different from that of Theorem 3.2, we will still need an analog of the main property therein, i.e., the existence of a good-wise comparison between the goods an agent gets under the bluff profile and the ones she gets by deviating. As expected, the corresponding property here (see Lemma 3.5) is more nuanced and does not immediately imply Theorem 3.7 as we are now missing the analog of Lemma 3.3.

Throughout this section, we are going to argue about an arbitrary agent i . To simplify the notation, let us rename $X_i = A_i^b = \{h_i, h_{n+i}, \dots, h_{(k-1)n+i}\}$ to simply $X = \{x_1, x_2, \dots, x_k\}$, where we have kept the order of indices the same, i.e., $x_j = h_{(j-1)n+i}$. This way, the goods in X are ordered according to how they were allocated to agent i in the run of Mechanism 1 with the bluff profile as input.

We also need to define the ordering of the goods agent i gets when she deviates from the bluff bid $\succ^b$ to another preference ranking $\succ_i$. Let $A_i = Y = \{y_1, y_2, \dots, y_k\}$ be this set of goods. Instead of renaming the elements of Y in a generic fashion like in the proof of Theorem 3.2, doing so becomes significantly more complicated, and we need to do it in a more systematic way, see Algorithm 3.

Algorithm 3 Greedy renaming of goods for the deviating agent i

Input: $X = \{x_1, x_2, \dots, x_k\}$, Y , and a value oracle for $v_i(\cdot)$

- 1: $Z = Y$
 - 2: **for** $j = |Y|, \dots, 1$ **do**
 - 3: $y'_j = \arg \min_{g \in Z} v_i(g \mid \{x_1, \dots, x_{j-1}\})$ // Ties are broken lexicographically.
 - 4: $Z = Z \setminus \{y'_j\}$
 - 5: **return** $(y'_1, y'_2, \dots, y'_{|Y|})$
-

In what follows, we assume that the indexing $y_1, y_2, \dots, y_k$ is already the result of Algorithm 3. This renaming is crucial and it will be used repeatedly. In particular, we need this particular ordering in order to prove that $v_i(x_j \mid \{x_1, \dots, x_{j-1}\}) \geq v_i(y_j \mid \{x_1, \dots, x_{j-1}\})$, for all $j \in [k]$, in Lemma 3.5 below. Towards that, we need to fix some notation for the sake of readability. For $j \in [k]$, we use X_-^j and X_+^j to denote the sets $\{x_1, x_2, \dots, x_j\}$ and $\{x_j, x_{j+1}, \dots, x_k\}$, respectively. The sets Y_-^j and Y_+^j , for $j \in [k]$, are defined analogously. We also use $X_-^0 = Y_-^0 = \emptyset$. The main high-level idea of the proof is that if $v_i(y_\ell \mid X_-^{\ell-1}) > v_i(x_\ell \mid X_-^{\ell-1})$ for some ℓ , then it must be the case that during the execution of $\text{Round-Robin}(\succ^b)$ every good in $Y_-^\ell = \{y_1, \dots, y_\ell\}$ is allocated before the turn of agent i in round ℓ . Then, using a simple counting argument, we show that agent i cannot receive all the goods in Y_-^ℓ when deviating, leading to a contradiction.

LEMMA 3.5. *Let $X = \{x_1, x_2, \dots, x_k\}$ be agent i 's bundle in Round-Robin($\succ^b$), where goods are indexed in the order they were allocated, and $Y = \{y_1, y_2, \dots, y_k\}$ be i 's bundle in Round-Robin($\succ_i, \succ_{-i}^b$), where goods are indexed by Algorithm 3. Then, for every $j \in [k]$, we have $v_i(x_j | X^{j-1}) \geq v_i(y_j | Y^{j-1})$.*

PROOF. The way goods in X are indexed, we have that x_j is the good allocated to agent i in round j of Round-Robin($\succ^b$). Suppose, towards a contradiction, that there is some $\ell \in [k]$, for which we have $v_i(y_\ell | Y^{\ell-1}) > v_i(x_\ell | X^{\ell-1})$. First notice that $\ell \neq 1$, as x_1 is, by the definition of the bluff profile, a singleton of maximum value for agent i excluding the goods allocated to agents 1 through $i - 1$ in round 1, regardless of agent i 's bid. Thus, $\ell \geq 2$.

Let $B \subseteq M$ and $D \subseteq M$ be the sets of goods allocated (to any agent) up to right before a good is allocated to agent i in round ℓ in Round-Robin($\succ^b$) and Round-Robin($\succ_i, \succ_{-i}^b$), respectively. Clearly, $|B| = |D| = (\ell - 1)n + i - 1$. In fact, we claim that in this case the two sets are equal.

CLAIM 3.6. *It holds that $\{y_1, \dots, y_\ell\} \subseteq B = D$.*

PROOF OF THE CLAIM. For the first part of the claim, i.e., that $\{y_1, \dots, y_\ell\} \subseteq B$, we begin by observing that $v_i(y_j | Y^{j-1}) \geq v_i(y_\ell | Y^{\ell-1}) > v_i(x_\ell | X^{\ell-1})$, for every $j \in [\ell - 1]$, where the first inequality follows from way Algorithm 3 ordered the elements of Y . Now consider the execution of Round-Robin($\succ^b$). Since x_ℓ was the good allocated to agent i in round ℓ , x_ℓ had maximum marginal value for agent i with respect to $X^{\ell-1}$ among the available goods. Thus, none of the goods $y_1, \dots, y_\ell$ were available at the time. That is, $y_1, \dots, y_\ell$ were all already allocated to some of the agents (possibly including agent i herself). We conclude that $\{y_1, \dots, y_\ell\} \subseteq B$.

Moving to the second part of the claim, suppose for a contradiction that $D \neq B$ and consider the execution of Round-Robin($\succ_i, \succ_{-i}^b$). Recall that the goods in B are still the $(\ell - 1)n + i - 1$ most preferable goods for every agent in $N \setminus \{i\}$ according to the profile $(\succ_i, \succ_{-i}^b)$. Therefore, all agents in $N \setminus \{i\}$ will get goods from B allocated to them up to the point when a good is allocated to agent i in round ℓ , regardless of what $\succ_i$ is. If agent i also got only goods from B allocated to her in the first $\ell - 1$ rounds of Round-Robin($\succ_i, \succ_{-i}^b$), then D would be equal to B . Thus, at least one good which is not in B (and thus, not in $\{y_1, \dots, y_\ell\}$) must have been allocated to agent i in the first $\ell - 1$ rounds. As a result, at the end of round $\ell - 1$, there are at least two goods from $\{y_1, \dots, y_\ell\}$ that have not yet been allocated to i .

However, we claim that up to right before a good is allocated to agent i in round $\ell + 1$, all goods in B (and thus in $\{y_1, \dots, y_\ell\}$ as well) will have been allocated, leaving i with at most $\ell - 1$ goods from $\{y_1, \dots, y_\ell\}$ in her final bundle and leading to a contradiction. Indeed, this follows from a simple counting argument. Right before a good is allocated to agent i in round $\ell + 1$, the goods allocated to agents in $N \setminus \{i\}$ are exactly $\ell(n - 1) + i - 1 \geq (\ell - 1)n + i - 1 = |B|$. As noted above, agents in $N \setminus \{i\}$ will get goods from B allocated to them as long as they are available. Thus, no goods from B , or from $\{y_1, \dots, y_\ell\}$ in particular, remain unallocated right before a good is allocated to agent i in round $\ell + 1$. Therefore, agent i may get at most $\ell - 1$ goods from $\{y_1, \dots, y_\ell\}$ (at most $\ell - 2$ in the first $\ell - 1$ rounds and one in round ℓ), contradicting the definition of the set Y . We conclude that $D = B$. ■

Given the claim, it is now easy to complete the proof. Clearly, in the first $\ell - 1$ rounds of Round-Robin($\succ_i, \succ_{-i}^b$) at most $\ell - 1$ goods from $\{y_1, \dots, y_\ell\}$ have been allocated to agent i . However, when it is i 's turn in round ℓ , only goods in $M \setminus D$ are available, by the definition of D . By Claim 3.6, we have $\{y_1, \dots, y_\ell\} \subseteq D$, and thus there is at least one good $\{y_1, \dots, y_\ell\}$ that is allocated to another agent, which contradicts the definition of Y . □

We are now ready to state and prove the main result of this section.

THEOREM 3.7. *When all agents have submodular valuation functions, the bluff profile is a $\frac{1}{2}$ -approximate PNE of Mechanism 1. Moreover, this is tight, i.e., for any $\varepsilon > 0$, there are instances where the bluff profile is not a $(\frac{1}{2} + \varepsilon)$ -approximate PNE.*

PROOF. We are going to use the notation used so far in the section and consider the possible deviation of an arbitrary agent i . Like in the statement of Lemma 3.5, $X = \{x_1, \dots, x_k\}$ is agent i 's bundle in Round-Robin($\succ^b$), with goods indexed in the order they were allocated, and $Y = \{y_1, y_2, \dots, y_k\}$ is i 's bundle in Round-Robin($\succ_i, \succ_{-i}^b$), with goods indexed by Algorithm 3. Also, recall that $X_-^j = \{x_1, \dots, x_j\}$ and $X_+^j = \{x_j, \dots, x_k\}$ (and similarly for Y_-^j and Y_+^j). We also use the convention that $Y_+^{k+1} = \emptyset$. For any $j \in [k]$, we have

$$\begin{aligned} v_i(X_-^j) - v_i(X_-^{j-1}) &= v_i(x_j \mid X_-^{j-1}) \\ &\geq v_i(y_j \mid X_-^{j-1}) \\ &\geq v_i(y_j \mid X_-^{j-1} \cup Y_+^{j+1}) \\ &= v_i(X_-^{j-1} \cup Y_+^{j+1} \cup \{y_j\}) - v_i(X_-^{j-1} \cup Y_+^{j+1}) \\ &= v_i(X_-^{j-1} \cup Y_+^j) - v_i(X_-^{j-1} \cup Y_+^{j+1}) \\ &\geq v_i(X_-^{j-1} \cup Y_+^j) - v_i(X_-^j \cup Y_+^{j+1}). \end{aligned}$$

The first inequality holds because Lemma 3.5 applies on X and Y , whereas the second inequality holds because of submodularity. Finally, the last inequality holds since $X_-^{j-1} \subseteq X_-^j$ and $v_i(\cdot)$ is non-decreasing, for every $i \in N$. Using these inequalities along with a standard expression of the value of a set as a sum of marginals, we have

$$\begin{aligned} v_i(X) &= v_i(X_-^k) - v_i(X_-^0) \\ &= \sum_{j=1}^k (v_i(X_-^j) - v_i(X_-^{j-1})) \\ &\geq \sum_{j=1}^k (v_i(X_-^{j-1} \cup Y_+^j) - v_i(X_-^j \cup Y_+^{j+1})) \\ &= v_i(X_-^0 \cup Y_+^1) - v_i(X_-^k \cup Y_+^{k+1}) \\ &= v_i(Y) - v_i(X). \end{aligned}$$

Thus, we have $v_i(X) \geq \frac{1}{2} \cdot v_i(Y)$, and we conclude that $\succ^b$ is a $\frac{1}{2}$ -approximate PNE of Mechanism 1.

To show that the result is tight, consider an example with two agents and five goods. The valuation function of agent 1 is additive and defined as follows on the singletons:


$$v_1(g_1) = 2 \quad v_1(g_2) = 1 \quad v_1(g_3) = 1 - \varepsilon_1 \quad v_1(g_4) = 1 - \varepsilon_2 \quad v_1(g_5) = 1 - \varepsilon_3,$$

where $1 \gg \varepsilon_3 > \varepsilon_2 > \varepsilon_1 > 0$.

The valuation function of agent 2 is OXS² and defined by the maximum matchings in the bipartite graph below, e.g., $v_2(\{g_1, g_2\}) = 2 + 1 = 3$ and $v_2(\{g_1, g_4, g_5\}) = 2 + 1 - \varepsilon_2 = 3 - \varepsilon_2$.

It is not hard to see that the bluff profile for this instance consists of the following declared ordering by both agents: $g_1 > g_2 > g_3 > g_4 > g_5$. The allocation produced by Mechanism 1 for the bluff profile is then $A = (A_1, A_2)$, where $A_1 = \{g_1, g_3, g_5\}$, and $A_2 = \{g_2, g_4\}$. Observe that $v_1(A_1) = 4 - \varepsilon_1 - \varepsilon_3$ and $v_2(A_2) = 1$. It is easy to see that there is no profitable deviation for agent 1, while the maximum value that agent 2 can attain by deviating is $2 - \varepsilon_1 - \varepsilon_2$. Agent 2 achieves

²Roughly speaking, OXS functions generalize unit-demand functions. The set of OXS functions is a strict superset of additive functions and a strict subset of submodular functions. See [Lehmann et al. 2006; Leme 2017].



this by reporting the preference ranking: $g_3 > g_4 > g_1 > g_2 > g_5$ and getting goods $\{g_3, g_4\}$. This implies that for any $\varepsilon > 0$ one can choose appropriately small $\varepsilon_1, \varepsilon_2, \varepsilon_3$ so that the bluff profile is not a $(\frac{1}{2} + \varepsilon)$ -approximate PNE. $\square$

In Section 4, we show that every approximate PNE of Mechanism 1 results in an approximately EF1 allocation. Here, as a warm-up, we start this endeavor with an easy result which holds specifically for the bluff profile (and can be extended to approximate PNE where all agents submit the same preference ranking) but shows a better fairness guarantee than our general Theorem 4.4.

THEOREM 3.8. *When all agents have submodular valuation functions $v_1, \dots, v_n$, the allocation returned by Round-Robin($\succ^b$) is $\frac{1}{2}$ -EF1 with respect to $v_1, \dots, v_n$. Moreover, this is tight, i.e., for any $\varepsilon > 0$, there are instances where this allocation is not $(\frac{1}{2} + \varepsilon)$ -EF1.*

PROOF. Suppose, towards a contradiction, that the allocation $(A_1^b, A_2^b, \dots, A_n^b)$ returned by Round-Robin($\succ^b$) is not $\frac{1}{2}$ -EF1. That is, there exist agents i and j such that $v_i(A_i^b) < 0.5 \cdot v_i(A_j^b \setminus \{g\})$, for all $g \in A_j^b$. We are going to show that this allows us to construct a deviation for agent i where she gets value more than $2v_i(A_i^b)$, contradicting the fact that $\succ^b$ is a $\frac{1}{2}$ -approximate PNE. Recall that using the renaming $h_1, h_2, \dots$ produced by Algorithm 2, we have $A_i^b = \{h_i, h_{n+i}, \dots, h_{(k-1)n+i}\}$ and $A_j^b = \{h_j, h_{n+j}, \dots, h_{(k-1)n+j}\}$.

Let δ be the indicator variable of the event $j < i$, i.e., δ is 1 if $j < i$ and 0 otherwise. We will show that it is possible for agent i to get the set $\{h_{\delta n+j}, h_{(1+\delta)n+j}, h_{(2+\delta)n+j}, \dots, h_{(k-1)n+j}\}$, which is either the entire A_j^b (when $i < j$) or $A_j^b \setminus \{h_j\}$ (when $j < i$). In particular, let $\succ_i$ be a preference ranking that starts with all goods in A_j^b in the same order as they were allocated to agent j in Round-Robin($\succ^b$), followed by everything else, in any order.

Consider the execution of Round-Robin($\succ_i, \succ_{-i}^b$). The crucial, yet simple, observation (that makes an inductive argument work) is that the first $i - 1$ goods $h_1, \dots, h_{i-1}$ are allocated as before, then good $h_{\delta n+j}$ (rather than h_i) is allocated to agent i , and after that the $n - 1$ top goods for all agents in $N \setminus \{i\}$ according to $\succ_{-i}^b$ are $h_i, h_{i+1}, \dots, h_{\delta n+j-1}, h_{\delta n+j+1}, \dots, h_{n+i-1}$, and these are allocated in the next $n - 1$ steps of the algorithm. As a result, right before a second good is allocated to agent i , the available goods are $h_{n+i}, h_{n+i+1}, \dots, h_m$ exactly as in the execution of Round-Robin($\succ^b$).

More generally, right before an r -th good is allocated to i , her bundle is $\{h_{\delta n+j}, h_{(1+\delta)n+j}, h_{(2+\delta)n+j}, \dots, h_{(r-2+\delta)n+j}\}$, and the available goods are $h_{(r-1)n+i}, h_{(r-1)n+i+1}, \dots, h_m$ (as they were in the execution of Round-Robin($\succ^b$)). Then good $h_{(r-1+\delta)n+j}$ (rather than $h_{(r-1)n+i}$) is allocated to agent i , and after that the $n - 1$ top goods for all agents according to $\succ_{-i}^b$ are

$$h_{(r-1)n+i}, h_{(r-1)n+i+1}, \dots, h_{(r-1+\delta)n+j-1}, h_{(r-1+\delta)n+j+1}, \dots, h_{rn+i-1},$$

and they are allocated in the next $n - 1$ steps of the algorithm. At the end, agent i gets the entire A_j^b or $A_j^b \setminus \{h_j\}$ plus some arbitrary good, depending on whether $i < j$ or $j < i$. In either case, by monotonicity, agent i 's value for her bundle is at least $v_i(A_j^b \setminus \{h_j\}) > 2v_i(A_i^b)$, where the last inequality follows from our assumption that $(A_1^b, A_2^b, \dots, A_n^b)$ is not $\frac{1}{2}$ -EF1. Therefore, by deviating from $\succ^b$ to $\succ_i$, agent i increases her value by a factor strictly greater than 2, contradicting Theorem 3.7.

To show that this factor is tight, we again turn to the example given within the proof of Theorem 3.7. Recall the allocation produced by Mechanism 1 for the bluff profile is $A = (A_1, A_2)$, with $A_1 = \{g_1, g_3, g_5\}$ and $A_2 = \{g_2, g_4\}$. Observe that agent 1 is envy-free towards agent 2 as $v_1(A_1) = 4 - \varepsilon_1 - \varepsilon_3 > 2 - \varepsilon_2 = v_1(A_2)$. On the other hand, $v_2(A_2) = 1$, whereas $v_2(A_1) = 4 - \varepsilon_1 - \varepsilon_3$ and $v_2(A_1 \setminus \{g_1\}) = 2 - \varepsilon_1 - \varepsilon_3$. The latter implies that for any $\varepsilon > 0$ one can choose appropriately small $\varepsilon_1, \varepsilon_2, \varepsilon_3$ so that the bluff profile does not result in a $(\frac{1}{2} + \varepsilon)$ -EF1 allocation with respect to the true valuation functions of the agents. $\square$

4 FAIRNESS PROPERTIES OF PNE

In Section 2.3, Proposition 2.5, we state the fairness guarantees of Round-Robin—viewed as an algorithm—when all agents have cancelable valuation functions. So far, we have not discussed this matter for the submodular case. It is not hard to see, however, that Theorem 3.8 and the definition of the bluff profile via Algorithm 2 imply that when we have (value oracles for) the valuation functions, then we can use Round-Robin to algorithmically produce $\frac{1}{2}$ -EF1 allocations. Using similar arguments, we show next that for any preference profile $\succ = (\succ_1, \dots, \succ_n)$ and any $i \in N$, there is always a response $\succ'_i$ of agent i to $\succ_{-i}$, such that the allocation returned by Round-Robin($\succ'_i, \succ_{-i}$) is $\frac{1}{2}$ -EF1 from agent i 's perspective.

Towards this, we first need a variant of Algorithm 2 that considers everyone in $N \setminus \{i\}$ fixed to their report in $\succ_{-i}$ and greedily determines a “good” response for agent i . An intuitive interpretation of what Algorithm 4 below is doing, can be given if one sees Mechanism 1 as a sequential game. Then, given that everyone else stays consistent with $\succ_{-i}$, agent i picks a good of maximum marginal value every time her turn is up.

Algorithm 4 Greedy response of agent i

Input: $N, M, \succ_{-i}$, value oracle for v_i

```

1:  $S = M; X = \emptyset$ 
2: for  $j = 1, \dots, m$  do
3:    $\ell = (j - 1) \pmod{n} + 1$ 
4:   if  $\ell = i$  then
5:      $x_{\lceil j/n \rceil} = \arg \max_{g \in S} v_i(g \mid X)$  // Ties are broken lexicographically.
6:      $X = X \cup \{x_{\lceil j/n \rceil}\}$ 
7:      $S = S \setminus \{x_{\lceil j/n \rceil}\}$ 
8:   else
9:      $g = \text{top}(\succ_\ell, S)$ 
10:     $S = S \setminus \{g\}$ 
11: return  $x_1 \succ'_i x_2 \succ'_i \dots \succ'_i x_k \succ'_i \dots$  // Arbitrarily complete  $\succ'_i$  with goods in  $M \setminus X$ .
```

Proving the next lemma closely follows the proof of Theorem 3.7 but without the need of an analog of Lemma 3.5, as we get this for free from the way the greedy preference profile $\succ'_i$ is constructed.

LEMMA 4.1. Assume that agent i has a submodular valuation function v_i . If $\succ'_i$ is the ranking returned by Algorithm 4 when given $N, M, \succ_{-i}, v_i$, then the allocation $(A'_1, A'_2, \dots, A'_n)$ returned by Round-Robin($\succ'_i, \succ_{-i}$) is such that for every $j \in N$, with $A'_j \neq \emptyset$, there exists a good $g \in A'_j$, so that $v_i(A'_i) \geq \frac{1}{2} \cdot v_i(A'_j \setminus \{g\})$.

4.1 The Case of Two Agents

As a warm-up, we begin with the easier case of $n = 2$. Not only the proofs of our main results for submodular and additive functions are much simpler here, but the fairness guarantees are stronger as well.

THEOREM 4.2. Let $\alpha \in (0, 1]$. Assume we have a fair division instance with two agents, whose valuation functions v_1, v_2 are submodular. Then any allocation that corresponds to a α -approximate PNE of the Round-Robin mechanism is $\frac{\alpha}{2-\alpha}$ -EF1 with respect to v_1, v_2 .

For additive valuation functions we can get a slightly stronger fairness guarantee, which we show that is also tight for any α , with an even easier proof. Note that this reproduces the result of Amanatidis et al. [2021] for exact PNE in the case of two agents.

THEOREM 4.3. Let $\alpha \in (0, 1]$. Assume we have a fair division instance with two agents, whose valuation functions v_1, v_2 are additive. Then any allocation that corresponds to a α -approximate PNE of the Round-Robin mechanism is $\frac{\alpha}{2-\alpha}$ -EF1 with respect to v_1, v_2 . This is tight, i.e., for any $\varepsilon > 0$, there are instances where a α -approximate PNE does not correspond to a $(\frac{\alpha}{2-\alpha} + \varepsilon)$ -EF1 allocation.

PROOF. Let $\succ = (\succ_1, \succ_2)$, A_1, A_2 be as in the proof of Theorem 4.2, but now consider the deviation of agent i to $\succ'_i$ which is a strict version of her true preference ranking $\succ_i^*$. Again, let (A'_1, A'_2) be the allocation returned by Round-Robin($\succ'_i, \succ_{-i}$).

Let g be good of maximum value in A'_j according to v_i . Since $\succ'_i$ is a true preference ranking of agent i , according to Proposition 2.5 (A'_1, A'_2) is EF1 from the point of view of agent i . That is, we have $v_i(A'_i) \geq v_i(A'_j \setminus \{g\})$ and, thus, $v_i(A'_i) \geq \frac{1}{2} \cdot v_i(M \setminus \{g\})$. Therefore,

$$\begin{aligned} v_i(A_j \setminus \{g\}) &= v_i(M \setminus \{g\}) - v_i(A_i) \\ &\leq 2 \cdot v_i(A'_i) - v_i(A_i) \\ &\leq \frac{2}{\alpha} \cdot v_i(A_i) - v_i(A_i) \\ &= \frac{2-\alpha}{\alpha} \cdot v_i(A_i), \end{aligned}$$

where the second inequality follows from the fact that $\succ$ is a α -approximate PNE and thus $v_i(A_i) \geq \alpha \cdot v_i(A'_i)$. We conclude that (A_1, A_2) is $\frac{\alpha}{2-\alpha}$ -EF1 with respect to v_1, v_2 .

To see that this guarantee is tight, consider an instance with two agents, and a set of five goods $\{g_1, g_2, \dots, g_5\}$. In addition, let the valuation functions of the agents to be additive and defined by:

$$v_1(g_j) = \begin{cases} 6, & \text{if } j = 1 \\ 3 + \delta, & \text{if } j = 2 \\ 3, & \text{if } j = 3 \\ 0.5 + \delta, & \text{if } j = 4 \\ 0.5, & \text{if } j = 5 \end{cases} \quad v_2(g_j) = \begin{cases} 6\beta, & \text{if } j = 1 \\ 3\beta + \delta, & \text{if } j = 2 \\ 3\beta, & \text{if } j = 3 \\ 0.5 + \delta, & \text{if } j = 4 \\ 0.5, & \text{if } j = 5 \end{cases}$$

where $0.5 \gg \delta$, and $\beta > \frac{1}{6} + \delta$. Now suppose that the agents bid as follows: Agent 1 bids truthfully (i.e., an ordering $\succ_1$ that is consistent with her true valuation function), while agent 2 bids $g_5 \succ_2$

$g_4 \succ_2 g_1 \succ_2 g_2 \succ_2 g_3$. It is easy to confirm that the produced allocation is $A = (A_1, A_2) = (\{g_1, g_2, g_3\}, \{g_4, g_5\})$. Regarding agent 1, she takes her three most desirable goods in this allocation so there is no profitable deviation for her. For the same reason, she is envy-free towards agent 2.

Moving to agent 2, by observing her valuation function, we immediately derive that she is $\frac{1+\delta}{6\beta+\delta}$ -EF1 towards agent 1. The only thing that remains, is to check how much agent 2 can improve her utility through deviating. Initially notice that agent 2 cannot get good g_1 regardless of her bid as this good is taken by agent 1 in round 1. At the same time, it is easy to verify that she cannot get both goods g_2 and g_3 due to the declared ordering of agent 1. Thus, the best bundle of goods that she can acquire is $\{g_2, g_4\}$ by deviating to the bid: $g_2 \succ'_2 g_4 \succ'_2 g_1 \succ'_2 g_3 \succ'_2 g_5$ and attain a value of $3\beta + 0.5 + 2\delta$.

By setting $\alpha = \frac{1+\delta}{3\beta+0.5+2\delta}$ we trivially have that $(\succ_1, \succ_2)$ is a α -approximate PNE. On the other hand, for a given $\varepsilon > 0$, we have $\frac{\alpha}{2-\alpha} + \varepsilon = \frac{1+\delta}{6\beta+3\delta} + \varepsilon$ which is strictly larger than $\frac{1+\delta}{6\beta+\delta}$ for sufficiently small δ . That is, there is a choice of δ so that the α -approximate PNE $(\succ_1, \succ_2)$ is not $\frac{\alpha}{2-\alpha} + \varepsilon$ -EF1. $\square$

4.2 The Case of n Agents

Looking back at the proofs of Theorems 4.2 and 4.3, the obvious fact that everything not in A_i or A'_i was allocated to agent j played a key role in proving our sharp bounds. Moving to the general case of n agents, there is no reason to expect that we have some control on how the goods are redistributed between agents in $N \setminus \{i\}$ when agent i deviates from an (approximate) equilibrium. Surprisingly, we show that this redistribution does not favor any agent too much from i 's perspective when the valuation functions are submodular or subadditive cancelable (Lemmata 4.6 and 4.7). Consequently, the main results of this section have similar flavor not only with respect to their statements, but with respect to their proofs as well.

THEOREM 4.4. *Let $\alpha \in (0, 1]$. For instances with submodular valuation functions $\{v_i\}_{i \in N}$, any α -approximate PNE of the Round-Robin mechanism is $\frac{\alpha}{3}$ -EF1 with respect to $\{v_i\}_{i \in N}$.*

THEOREM 4.5. *Let $\alpha \in (0, 1]$. For instances with subadditive cancelable valuation functions $\{v_i\}_{i \in N}$, any α -approximate PNE of the Round-Robin mechanism is $\frac{\alpha}{2}$ -EF1 with respect to $\{v_i\}_{i \in N}$.*

As the proofs of both theorems have the same general structure and share Lemmata 4.6 and 4.7, we begin with some common wording and notation, consistent with our proofs for two agents. Given any instance, we use $\succ = (\succ_1, \dots, \succ_n)$ for an arbitrary α -approximate PNE of Mechanism 1. We then consider the deviation of some agent i to a preference ranking $\succ'_i$; in the submodular case $\succ'_i$ is the output of Algorithm 4 when given $\succ_{-i}$ and v_i , whereas in the cancelable case $\succ'_i$ is a strict version of i 's true preference ranking $\succ_i^*$. We use $(A_1, \dots, A_n)$ and $(A'_1, \dots, A'_n)$ to denote the allocations returned by Round-Robin($\succ$) and Round-Robin($\succ'_i, \succ_{-i}$), respectively.

In order to show that $(A_1, \dots, A_n)$ is $\frac{\alpha}{\kappa}$ -EF1 from agent i 's perspective (where κ is 3 for submodular and 2 for cancelable functions), we use the stronger EF1 guarantees that $(A'_1, \dots, A'_n)$ has from her perspective. To this end, we use h_r^ℓ to denote the good that was allocated to an agent $\ell \in N$ in round r of Round-Robin($\succ'_i, \succ_{-i}$). In particular, $A'_i = \{h_1^i, h_2^i, \dots, h_k^i\}$; recall that $k = m/n$. Further, given that we have fixed agent i , we use S_r and S'_r , for $0 \leq r \leq k-1$, to denote the set of goods that had been allocated up to right before a good was allocated to i in round $r+1$ of Round-Robin($\succ$) and Round-Robin($\succ'_i, \succ_{-i}$), respectively. That is, for $0 \leq r \leq k-1$, S_r and S'_r contain the goods allocated in steps 1 through $rn+i-1$ of Round-Robin($\succ$) and Round-Robin($\succ'_i, \succ_{-i}$), respectively.

For the next technical lemma we assume that the valuation functions are either submodular or cancelable and, in each case, we use the corresponding $\succ'_i$ as described above.

LEMMA 4.6. *For any $r \in [k]$, right before an r -th good is allocated to agent i in Round-Robin($\succ$), there are at most $r-1$ goods from S'_{r-1} that are still unallocated, i.e., $|S'_{r-1} \setminus S_{r-1}| \leq r-1$.*

PROOF. We will prove the statement using induction on r . For $r = 1$, it is straightforward that $S_0 = S'_0$, as the preference rankings of agents 1 through $i - 1$ are the same in the two runs of the mechanism and, thus, the first goods allocated to them are exactly the same.

Now suppose that the statement is true for every round up to round r ; we will show that it is true for round $r + 1$ as well. Initially, observe that if the number of unallocated goods from S'_{r-1} is $r - 1$ right before a good is allocated to agent i in round r , it will trivially be at most $r - 1$ right before a good is allocated to agent i in round $r + 1$ (as the number of unallocated goods from any set cannot increase as the allocation progresses). That is, $|S'_{r-1} \setminus S_r| \leq r - 1$.

Notice that the goods that might cause $S'_r \setminus S_r$ to increase are the elements of

$$S'_r \setminus S'_{r-1} = \{h_r^i, h_r^{i+1}, \dots, h_r^n, h_{r+1}^1, h_{r+1}^2, \dots, h_{r+1}^{i-1}\},$$

and suppose that there are λ goods therein which are still unallocated right before a good is allocated to agent i in round $r + 1$ of Round-Robin($\succ$). Clearly, if $\lambda \leq 1$, we are done. So, assume that $\lambda \geq 2$. This means that there are $\lambda - 1 \geq 1$ unallocated goods in $(S'_r \setminus S'_{r-1}) \setminus \{h_r^i\}$. Let g be one of these goods and let j be the agent to whom g was given, i.e., $g = h_{\bar{r}}^j$, where $\bar{r} = r$, if $j > i$, and $\bar{r} = r + 1$, if $j < i$. In either case, notice that according to $\succ_j$ the good g is better than any good in $M \setminus S'_r$ or else it would not have been allocated to j at round $\bar{r}$ of Round-Robin($\succ'_i, \succ_{-i}$) when everything in $M \setminus S'_r$ is still available. We claim that g does not increase the number of elements in $S'_r \setminus S_r$. Indeed, given that g was available during step $(\bar{r} - 1)n + j$ of Round-Robin($\succ$) and that j 's declared preference ranking is still $\succ_j$, the only possibility is that during that step one of the unallocated goods from $S'_{r-1} \cup \{h_r^i, h_r^{i+1}, \dots, h_{\bar{r}}^{j-1}\}$ was allocated to j instead.

Therefore, the only good out of the λ candidate goods of $S'_r \setminus S'_{r-1}$ which might count towards the number of elements in $S'_r \setminus S_r$ is h_r^i . We conclude that $S'_r \setminus S_r \leq (r - 1) + 1 = r$. $\square$

Lemma 4.6 is global, illustrating that the sets S_r and S'_r cannot differ in more than a $1/n$ -th of their elements. The next lemma shows that no agent can accumulate too many goods from S'_r , for any $0 \leq r \leq k - 1$. Again, we assume that the valuation functions are either submodular or cancelable and, in each case, the appropriate $\succ'_i$ is used as discussed after the statements of Theorems 4.2 and 4.3. Note that S'_0 in the lemma's statement contains exactly these goods which we will exclude when showing the EF1 guarantee for our two theorems.

LEMMA 4.7. *For any $r \in [k]$ and any $j \in N$, agent j gets at most $2(r - 1)$ goods from $S'_{r-1} \setminus S'_0$ in the allocation $(A_1, \dots, A_n)$ returned by Round-Robin($\succ$), i.e., $|A_j \cap (S'_{r-1} \setminus S'_0)| \leq 2(r - 1)$.*

PROOF. Fix an $r \in [k]$ and a $j \in N$. Consider the end of step $(r - 1)n + i - 1$ of Round-Robin($\succ$), i.e., right before an r -th good is allocated to agent i . Ignoring all the goods allocated before i got her first good, agent j has received exactly $r - 1$ goods up to this point. As a result, the number of goods allocated to j from $S'_{r-1} \setminus S'_0$ at this point is at most $r - 1$.

At the same time, the number of goods from $S'_{r-1} \setminus S'_0$ that might end up in A_j in any future steps of Round-Robin($\succ$) are at most as many as the goods from S'_{r-1} that are still unallocated at the end of step $(r - 1)n + i - 1$. The latter, by Lemma 4.6, are also at most $r - 1$.

From these two observations, we have that the final bundle A_j of agent j may contain at most $2(r - 1)$ goods from $S'_{r-1} \setminus S'_0$. $\square$

With Lemma 4.7 at hand, we are now ready to prove Theorems 4.4 and 4.5; the proof of the latter is deferred to the full paper [Amanatidis et al. 2023].

PROOF OF THEOREM 4.4. We, of course, adopt the notation that has been used throughout this section, focusing on an arbitrary agent $i \in N$ and assuming that her deviation $\succ'_i$ has been the

output of Algorithm 4 with input $\succ_{-i}$ and v_i . In particular, $(A_1, \dots, A_n)$ and $(A'_1, \dots, A'_n)$ are the allocations returned by Round-Robin($\succ$) and Round-Robin($\succ'_i, \succ_{-i}$), respectively.

Consider another agent $j \in N \setminus \{i\}$. Let $A'_i = \{x_1, x_2, \dots, x_k\}$ and $A_j = \{y_1, y_2, \dots, y_k\}$, where in both sets goods are indexed in the order in which they were allocated in the run of Round-Robin($\succ'_i, \succ_{-i}$). For A'_i , this means that x_r was allocated in round r for all $r \in [k]$. For A_j , this indexing guarantees that for every $0 \leq \ell < r \leq k-1$, the goods in $A_j \cap (S'_\ell \setminus S'_{\ell-1})$ all have smaller indices than the goods in $A_j \cap (S'_r \setminus S'_{r-1})$ (where we use the convention that $S'_{-1} = \emptyset$). We further partition $A_j \setminus \{y_1\}$ to $Y_1 = \{y_1^1, \dots, y_{\tau_1}^1\}$ and $Y_2 = \{y_1^2, \dots, y_{\tau_2}^2\}$ which contain the goods of $A_j \setminus \{y_1\}$ with odd and even indices, respectively, and are both renamed according to Algorithm 3 with inputs A'_i, Y_1, v_i , and A'_i, Y_2, v_i , respectively. Clearly, $\tau_1 = \lfloor \frac{k-1}{2} \rfloor$ and $\tau_2 = \lceil \frac{k-1}{2} \rceil$.

By Lemma 4.7, we have that A_j contains at most $2(r-1)$ goods from $S'_{r-1} \setminus S'_0$, for any $r \in [k]$. The original ordering $y_1, y_2, \dots$ of the goods in A_j and the way $A_j \setminus \{y_1\}$ was partitioned into Y_1 and Y_2 imply that $||Y_1 \cap (S'_{r-1} \setminus S'_0)| - |Y_2 \cap (S'_{r-1} \setminus S'_0)|| \leq 1$ and, thus, each of Y_1 and Y_2 contains at most $r-1$ goods from $S'_{r-1} \setminus S'_0$.

We also claim that, for $\ell \in \{1, 2\}$ and $r \in [\tau_\ell]$, we have

$$v_i(x_r \mid \{x_1, \dots, x_{r-1}\}) \geq v_i(y_r^\ell \mid \{x_1, \dots, x_{r-1}\}). \quad (1)$$

Suppose not. That is, there are $\ell \in \{1, 2\}$ and $r \in [\tau_\ell]$ so that (1) is violated. Note that, by the way Algorithm 3 ordered the elements of Y_1 and Y_2 , this implies

$$v_i(x_r \mid \{x_1, \dots, x_{r-1}\}) < v_i(y_r^\ell \mid \{x_1, \dots, x_{r-1}\}) \leq v_i(y_t^\ell \mid \{x_1, \dots, x_{r-1}\}),$$

for all $t \in [r]$. Since x_r was the good allocated to agent i at step $(r-1)n + i$ of Round-Robin($\succ'_i, \succ_{-i}$), x_r had maximum marginal value for i with respect to $\{x_1, \dots, x_{r-1}\}$ among the available goods. Thus, none of the goods $y_1^\ell, \dots, y_r^\ell$ were available at the time, i.e., $y_1^\ell, \dots, y_r^\ell \in S'_{r-1}$. Given that the only good of A_j that could possibly be in $S'_0 = S_0$ was y_1 which is not in $Y_1 \cup Y_2$. Therefore, $y_1^\ell, \dots, y_r^\ell \in S'_{r-1} \setminus S'_0$, which contradicts the fact that $|Y_\ell \cap (S'_{r-1} \setminus S'_0)| \leq r-1$. We conclude that (1) holds for all $\ell \in \{1, 2\}$ and $r \in [\tau_\ell]$.

We are now ready to apply Theorem 2.1 to bound the value of $A_j \setminus \{y_1\}$. We have

$$\begin{aligned} v_i(A_j \setminus \{y_1\}) &\leq v_i(A'_i) + \sum_{g \in (A_j \setminus \{y_1\}) \setminus A'_i} v(g \mid A'_i) \\ &= v_i(A'_i) + \sum_{g \in Y_1 \setminus A'_i} v(g \mid A'_i) + \sum_{g \in Y_2 \setminus A'_i} v(g \mid A'_i) \\ &= v_i(A'_i) + \sum_{\ell=1}^{\tau_1} v(y_\ell^1 \mid A'_i) + \sum_{\ell=1}^{\tau_2} v(y_\ell^2 \mid A'_i) \\ &\leq v_i(A'_i) + \sum_{\ell=1}^{\tau_1} v(y_\ell^1 \mid \{x_1, \dots, x_{\ell-1}\}) + \sum_{\ell=1}^{\tau_2} v(y_\ell^2 \mid \{x_1, \dots, x_{\ell-1}\}) \\ &\leq v_i(A'_i) + \sum_{\ell=1}^{\tau_1} v(x_\ell \mid \{x_1, \dots, x_{\ell-1}\}) + \sum_{\ell=1}^{\tau_2} v(x_\ell \mid \{x_1, \dots, x_{\ell-1}\}) \\ &\leq v_i(A'_i) + 2 \cdot \sum_{\ell=1}^k v(x_\ell \mid \{x_1, x_2, \dots, x_{\ell-1}\}) \\ &= v_i(A'_i) + 2 \cdot v_i(A'_i) \\ &\leq \frac{3}{\alpha} \cdot v_i(A_i), \end{aligned}$$

where the first inequality follows directly from Theorem 2.1, the second one follows from submodularity, the third inequality holds because of (1), the fourth one follows from the monotonicity of v_i , and the last inequality follows from the fact that $\succ$ is a α -approximate PNE and thus $v_i(A_i) \geq \alpha \cdot v_i(A'_i)$. We conclude that $(A_1, A_2, \dots, A_n)$ is $\frac{\alpha}{3}$ -EF1 with respect to the underlying valuation functions. $\square$

The $\alpha/(2 - \alpha)$ upper bound of Theorem 4.3 for the additive case applies to both submodular and subadditive cancelable valuation functions, leaving a very small gap for the latter. For the submodular case, we improve this upper bound to $\alpha/2$.

PROPOSITION 4.8. *Let $\alpha, \varepsilon \in (0, 1]$. For instances with submodular valuation functions $\{v_i\}_{i \in N}$, a α -approximate PNE of the Round-Robin mechanism may not be $(\frac{\alpha}{2} + \varepsilon)$ -EF1 with respect to $\{v_i\}_{i \in N}$.*

5 DISCUSSION AND FUTURE DIRECTIONS

In this work we studied the existence and fairness guarantees of the approximate pure Nash equilibria of the Round-Robin mechanism for agents with cancelable and submodular valuation functions. In both cases, we generalized the surprising connection between the stable states of the mechanism and its fairness properties, a connection that was only known for exact equilibria and additive valuation functions. For the function classes considered, we provide tight or almost tight bounds, thus giving a complete picture of the strengths and the limitations of the Round-Robin mechanism for these scenarios. There are several interesting related directions, some of which we discuss below.

An obvious first direction is to explore function classes beyond the ones studied here, with XOS or subadditive functions being prominent candidates. Since our results heavily rely on the properties of cancelable and submodular functions, it is likely that different approaches are needed for this endeavour. As we mention in the introduction, a second interesting direction, related to this one, is the study of the stability and fairness properties of variants of the Round-Robin mechanism that allow the agents to be more expressive. Analyzing mechanisms that take as an input value oracles seems to be highly non-trivial, and although some of our results might transfer in this setting, we suspect that, in general, strong impossibility results hold regarding the fairness guarantees of approximate PNE.

Finally, although here we focused on Round-Robin and EF1, most fair division algorithms have not been considered in the strategic setting. One promising such algorithm, which is both fundamental in a number of variants of the problem and simple enough, is the Envy-Cycle-Elimination algorithm of Lipton et al. [2004] which is known to compute EF1 allocations for general non-decreasing valuation functions. An appealing alternative here is studying the existence of equilibria of approximation algorithms for MMS allocations. An important advantage in this case is that once the existence of an approximate PNE is shown, the corresponding MMS guarantee comes for free (see also the related discussion of Amanatidis et al. [2021]).

ACKNOWLEDGMENTS

This work was supported by ERC Advanced Grant 788893 AMDROMA, PNRR MUR project PE0000013-FAIR, MIUR PRIN project ALGADIMAR “Algorithms, Games, and Digital Markets”, and NWO Veni project No. VI.Veni.192.153.

REFERENCES

Hannaneh Akrami, Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, and Ruta Mehta. 2022. EFX Allocations: Simplifications and Improvements. *CoRR* abs/2205.07638 (2022). <https://doi.org/10.48550/arXiv.2205.07638>

- Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. 2022. Fair Division of Indivisible Goods: A Survey. *CoRR* abs/2208.08782 (2022). <https://doi.org/10.48550/arXiv.2208.08782>
- Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. 2017a. Truthful Allocation Mechanisms Without Payments: Characterization and Implications on Fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC' 17*. ACM, 545–562.
- Georgios Amanatidis, Georgios Birmpas, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. 2021. Allocating Indivisible Goods to Strategic Agents: Pure Nash Equilibria and Fairness. In *Proceedings of the 17th International Conference on Web and Internet Economics, WINE 2021 (LNCS, Vol. 13112)*. 149–166.
- Georgios Amanatidis, Georgios Birmpas, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. 2023. Round-Robin Beyond Additive Agents: Existence and Fairness of Approximate Equilibria. *CoRR* abs/2301.13652 (2023). <https://doi.org/10.48550/arXiv.2301.13652>
- Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2016. On Truthful Mechanisms for Maximin Share Allocations. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence, IJCAI '16*. IJCAI/AAAI Press, 31–37.
- Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. 2017b. Approximation Algorithms for Computing Maximin Share Allocations. *ACM Trans. Algorithms* 13, 4 (2017), 52:1–52:28.
- Georgios Amanatidis, Evangelos Markakis, and Apostolos Ntokos. 2020. Multiple birds with one stone: Beating 1/2 for EFX and GMMS via envy cycle elimination. *Theor. Comput. Sci.* 841 (2020), 94–109.
- Haris Aziz, Sylvain Bouveret, Jérôme Lang, and Simon Mackenzie. 2017a. Complexity of Manipulating Sequential Allocation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence, AAAI '17*. AAAI Press, 328–334.
- Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. 2022. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi Agent Systems* 36, 1 (2022), 3.
- Haris Aziz, Paul Goldberg, and Toby Walsh. 2017b. Equilibria in Sequential Allocation. In *Proceedings of the 5th International Conference on Algorithmic Decision Theory, ADT '17 (LNCS, Vol. 10576)*. Springer, 270–283.
- Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2021. Fair and Truthful Mechanisms for Dichotomous Valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence, AAAI 2021*. AAAI Press, 5119–5126.
- Siddharth Barman and Sanath Kumar Krishnamurthy. 2020. Approximation Algorithms for Maximin Fair Division. *ACM Trans. Economics and Comput.* 8, 1 (2020), 5:1–5:28.
- Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. 2022. Almost Full EFX Exists for Four Agents. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence, AAAI 2022*. AAAI Press, 4826–4833.
- Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. 2016. Fair Allocation of Indivisible Goods. In *Handbook of Computational Social Choice*. Cambridge University Press, 284–310.
- Sylvain Bouveret and Jérôme Lang. 2014. Manipulating picking sequences. In *Proceedings of the 21st European Conference on Artificial Intelligence - ECAI 2014*, Vol. 263. IOS Press, 141–146.
- Sylvain Bouveret and Michel Lemaître. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems* 30, 2 (2016), 259–290.
- Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- Ioannis Caragiannis, Jugal Garg, Nidhi Rathi, Eklavya Sharma, and Giovanna Varricchio. 2022. Existence and Computation of Epistemic EFX Allocations. *CoRR* abs/2206.01710 (2022). <https://doi.org/10.48550/arXiv.2206.01710>
- Ioannis Caragiannis, Christos Kaklamani, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2009. On Low-Envy Truthful Allocations. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory, ADT 2009 (LNCS, Vol. 5783)*. Springer, 111–119.
- Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Economics and Comput.* 7, 3 (2019), 12:1–12:32.
- 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 2021*. AAAI Press, 5269–5276.
- Duncan K. Foley. 1967. Resource Allocation and the Public Sector. *Yale Economics Essays* 7 (1967), 45–98.
- Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2022. Fair allocation of indivisible goods: Beyond additive valuations. *Artif. Intell.* 303 (2022), 103633.
- 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 International Conference on Web and Internet Economics, WINE 2020 (LNCS, Vol. 12495)*. Springer, 370–383.
- David Kurokawa. 2017. *Fair Division in Game Theoretic Settings*. Ph.D. Dissertation. Carnegie Mellon University.
- Benny Lehmann, Daniel Lehmann, and Noam Nisan. 2006. Combinatorial auctions with decreasing marginal utilities. *Games Econ. Behav.* 55, 2 (2006), 270–296.

- Renato Paes Leme. 2017. Gross substitutability: An algorithmic survey. *Games Econ. Behav.* 106 (2017), 294–316.
- 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, EC '04*. ACM, 125–131.
- Pasin Manurangsi and Warut Suksompong. 2021. Closing Gaps in Asymptotic Fair Division. *SIAM Journal on Discrete Mathematics* 35, 2 (2021), 668–706.
- Evangelos Markakis. 2017. Approximation Algorithms and Hardness Results for Fair Division with Indivisible Goods. In *Trends in Computational Social Choice*. AI Access, Chapter 12.
- Evangelos Markakis and Christos-Alexandros Psomas. 2011. On Worst-Case Allocations in the Presence of Indivisible Goods. In *Proceedings of the 7th International Conference on Web and Internet Economics, WINE 2011 (LNCS, Vol. 7090)*. Springer, 278–289.
- Luisa Montanari, Ulrike Schmidt-Kraepelin, Warut Suksompong, and Nicholas Teh. 2022. Weighted Envy-Freeness for Submodular Valuations. *CoRR* abs/2209.06437 (2022). <https://doi.org/10.48550/arXiv.2209.06437>
- George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. 1978. An analysis of approximations for maximizing submodular set functions - I. *Math. Program.* 14, 1 (1978), 265–294.
- Benjamin Plaut and Tim Roughgarden. 2020. Almost Envy-Freeness with General Valuations. *SIAM J. Discret. Math.* 34, 2 (2020), 1039–1068.
- Ariel D. Procaccia. 2016. Cake Cutting Algorithms. In *Handbook of Computational Social Choice*. Cambridge University Press, 311–330.
- Alexandros Psomas and Paritosh Verma. 2022. Fair and Efficient Allocations Without Obvious Manipulations. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022*.
- Hugo Steinhaus. 1949. Sur la division pragmatique. *Econometrica* 17 (Supplement) (1949), 315–319.
- Hal R. Varian. 1974. Equity, Envy and Efficiency. *Journal of Economic Theory* 9 (1974), 63–91.