Revisiting the Fairness of Pure Nash Equilibria of Round-Robin

Georgios Amanatidis[1,2], Haris Aziz[3], Georgios Birmpas[4], Federico Fusco[5], Stefano Leonardi[5], and Rebecca Reiffenhäuser[6]

[1]Department of Informatics, Athens University of Economics and Business, Greece.
[2]Archimedes / Athena RC, Greece.
[3]School of Computer Science and Engineering, UNSW Sydney, Australia.
[4]Department of Computer Science, University of Liverpool, UK.
[5]Department of Computer, Control, and Management Engineering "Antonio Ruberti", Sapienza University of Rome, Italy.
[6]Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands.

Abstract

We prove that for additive valuations, not every pure Nash equilibrium of the Round-Robin mechanism induces an allocation that is envy-free up to one good (EF1) with respect to the agents' true valuations. This corrects an earlier claim of Amanatidis et al. (WINE 2021; Math. Oper. Res. 2024). On the positive side, every pure Nash equilibrium of every fixed recursively balanced picking sequence mechanism induces an allocation that is $1/2$-EF1 and proportional up to one good (PROP1). We prove the approximate EF1 guarantee through a potential function comparing two executions, and the proportionality guarantee via a deviation that preserves the agent's most valuable goods. We also give a family of Round Robin equilibria showing that it is not possible to have a strictly better guarantee than $2/3$-EF1 for all equilibria.

***Keywords*—** Fair division; Round-Robin; Pure Nash equilibrium; Envy-freeness up to one good; Proportionality up to one good; Recursively balanced picking sequence

1 Introduction

In the fair division of indivisible goods, envy-freeness up to one good (EF1) is one of the central fairness notions, and the Round-Robin algorithm [7], in which agents take turns picking their favourite remaining good, is the simplest way to achieve it: when agents report their valuations truthfully, Round-Robin always returns an EF1 allocation [8].

Strategic agents, however, may misreport (see, e.g., [9, 5]). Amanatidis et al. [1, 2] claim that, for additive nonnegative valuations, every pure Nash equilibrium of the Round-Robin reporting mechanism still induces an allocation that is EF1 with respect to the true valuations [2, Theorem 4]. We give a counterexample under the reporting model stated in their paper.

Despite the failure of exact EF1, we prove that two guarantees survive: every pure Nash equilibrium is $1/2$-EF1 and proportional up to one good (PROP1). Both statements hold for the larger class of fixed recursively balanced picking sequences [4], in which the order of the agents may change between rounds. The $1/2$-EF1 guarantee is already known for Round-Robin as a special case of the result of Amanatidis et al. for subadditive cancelable valuations [3, Theorem 8]. We give a direct proof for additive valuations

and extend the guarantee to every fixed recursively balanced sequence. The existence of a pure Nash equilibrium for any recursively balanced picking sequence mechanism follows from the bluff profile of Aziz et al. [6]; see the discussion at the end of Section 4. Finally, we give a more elaborate family of Round-Robin equilibria showing that no factor strictly greater than $2/3$ can be guaranteed with respect to approximate EF1 for all equilibria.

1.1 Our contributions

We first show that exact EF1 need not hold at equilibrium, even with three agents and eight goods (Proposition 1). In our example, two agents value all goods equally, and the third uses only three distinct values. All reports are strict, so the example is independent of any tie-breaking rule used.

Our main positive result (Theorem 1) combines $1/2$-EF1 and PROP1 for every fixed recursively balanced sequence. The pairwise guarantee follows from a potential function that compares the equilibrium execution with a deviation targeting another agent's bundle. It requires only that the agent whose envy we examine best responds to the other reports. The proportionality guarantee uses a different deviation, in which the agent first takes all goods that she values more than every good she does not receive, and then picks truthfully. Again, this only requires that her report is a best response.

Finally, Theorem 2 gives a family of Round-Robin equilibria ruling out every universal EF1 approximation factor greater than $2/3$. All agents except the one whose envy we examine have identical unit valuations, and the latter uses only three distinct values. The main step is to show that the agent whose envy we examine best responds, and this reduces to a counting argument: for every two goods that she takes from the bundle she envies, she loses at least three of her own goods.

2 Setting

An instance consists of a set $N = \{1, \dots, n\}$ of agents, a set M of indivisible goods with $m = |M|$, and, for each agent $i \in N$, an additive valuation v_i with $v_i(g) \geq 0$ for every $g \in M$ and $v_i(S) = \sum_{g \in S} v_i(g)$ for every $S \subseteq M$. An *allocation* $(A_1, \dots, A_n)$ is a partition of M, where A_i is the bundle of agent i.

Definition 1 (α-EF1)**.** For $\alpha \in [0, 1]$, an allocation $(A_1, \dots, A_n)$ is *α-EF1* if, for every pair $i, j \in N$ with $A_j \neq \emptyset$, there is a good $g \in A_j$ such that

$$v_i(A_i) \geq \alpha \, v_i(A_j \setminus \{g\}).$$

When $\alpha = 1$ we just say that the allocation is EF1.

Definition 2 (PROP1)**.** An allocation $(A_1, \dots, A_n)$ is *proportional up to one good* if, for every agent i, either $v_i(A_i) \geq v_i(M)/n$ or there is a good $g \in M \setminus A_i$ such that $v_i(A_i \cup \{g\}) \geq v_i(M)/n$.

For nonnegative additive valuations, an EF1 allocation is also PROP1.

Sequential mechanisms. Each agent i reports a bid vector $b_i \in \mathbb{R}^M_{\geq 0}$, which need not coincide with v_i. A picking sequence, fixed independently of the reports, specifies the active agent at each turn, and the active agent receives the available good with her highest reported bid; ties are broken lexicographically. The reports are submitted once, before the allocation begins, and remain fixed throughout the execution. A picking sequence is *recursively balanced* if it consists of complete rounds, each of which is a permutation of N, followed possibly by an incomplete round in which no agent appears twice. Round-Robin repeats the order $(1, 2, \dots, n)$, truncating the last round when all goods have been allocated.

Under truthful reports, a recursively balanced sequence produces an EF1 allocation: for every r, the good that agent i picks at her rth turn is at least as valuable to her as the good that any other agent j picks at the $(r+1)$th turn of j; so, removing the first good of j's bundle eliminates the envy of agent i towards j. Hence, it also produces a PROP1 allocation. For Round-Robin, see also [8].

Since the outcome depends on a report only through the strict ranking of goods that it induces after tie-breaking, and every strict ranking is induced by some bid vector, we identify each report with a strict ranking $\widehat{\succ}_i$ of M, as in [2].

Definition 3 (Pure Nash equilibrium)**.** A report profile $(\widehat{\succ}_1, \ldots, \widehat{\succ}_n)$ is a *pure Nash equilibrium* if no agent i can obtain a bundle of strictly higher true value v_i by unilaterally changing her report.

3 Not every pure Nash equilibrium induces an EF1 allocation

Proposition 1. *There is a three-agent, eight-good additive instance with strictly positive valuations and a pure Nash equilibrium of the Round-Robin reporting mechanism whose outcome is not EF1.*

Proof. Let $M = \{a, b, c, x, y, z, p, q\}$, and let the picking sequence be $(1, 2, 3 \mid 1, 2, 3 \mid 1, 2)$. Agents 1 and 3 value every good at 1. Agent 2's values are

	a	b	c	x	y	z	p	q
v_2	8	8	8	5	5	5	2	2

and the reported rankings are

$$\begin{aligned}
\widehat{\succ}_1 &: \quad a \succ b \succ c \succ x \succ y \succ z \succ p \succ q, \\
\widehat{\succ}_2 &: \quad x \succ y \succ z \succ a \succ b \succ c \succ p \succ q, \\
\widehat{\succ}_3 &: \quad p \succ x \succ y \succ q \succ z \succ a \succ b \succ c.
\end{aligned}$$

These are strict rankings, and can be implemented by strictly decreasing bids. Thus no tie-breaking is needed. The goods are picked in the order a, x, p, b, y, q, c, z, giving

$$A_1 = \{a, b, c\}, \qquad A_2 = \{x, y, z\}, \qquad A_3 = \{p, q\}.$$

Agents 1 and 3 always receive three and two goods, respectively, so neither can improve her utility. Agent 2 receives value 15. To verify that she also has no profitable deviation, first observe that a is always taken immediately. Moreover, she can receive at most one of b, c, and only on her first or second turn. Indeed, agent 3 never takes b or c, since she ranks them last and at least three goods are available at each of her turns. So, if agent 2 takes one of b, c on her first turn, agent 1 takes the other on her second turn; otherwise, agent 1 takes b on her second turn and c on her third turn, unless agent 2 has already taken it.

Without b, c, agent 2 can obtain value at most 15. With one of b, c and at most one of x, y, z, her value is at most $8 + 5 + 2 = 15$. It remains to exclude a bundle containing one of b, c and two of x, y, z. For such a bundle, her first two picks must consist of one good from each of these two sets. At that point, agent 1 has taken a and the other good from $\{b, c\}$, and agent 3 has taken p. Only q and two goods from $\{x, y, z\}$ remain. At least one of these two goods is x or y, which agent 3 ranks above q, so agent 3 takes one of them on her second turn. Agent 1 ranks the other above q and takes it on her third turn. Agent 2 must therefore take q on her third turn, a contradiction. This proves that the profile is a pure Nash equilibrium.

Finally, $v_2(A_2) = 15$, whereas $v_2(A_1 \setminus \{g\}) = 16$ for every $g \in A_1$. Thus the outcome is not EF1. □

Remark 1. The gap in the argument of Amanatidis et al. [2] (which also appears in [1]) occurs in the proof of their Lemma 2, at inequality (2). Their recursive construction replaces v_1 by intermediate valuation functions v_1^{r+1}, but the best-response condition invoked at that point holds only for the original valuation function v_1. Preserving the total value of A_1 and the values of all goods outside A_1, as the construction does, does not preserve the values of deviating bundles that contain only some of the goods of A_1.

4 Equilibrium guarantees for recursively balanced sequences

Although EF1 can fail at equilibrium, every pure Nash equilibrium is still $1/2$-EF1 and PROP1. This remains true even when the order of the agents changes between rounds.

Theorem 1. *For nonnegative additive valuations, every pure Nash equilibrium of every fixed recursively balanced picking sequence mechanism induces an allocation that is $1/2$-EF1 and PROP1 with respect to the agents' true valuation functions.*

We prove the two guarantees separately. The first proposition gives a quantitative proportionality bound. The second gives the pairwise bundle value comparison needed for $1/2$-EF1.

Proposition 2. *Let $A = (A_1, \dots, A_n)$ be the allocation induced by a report profile for a fixed recursively balanced picking sequence in which agent i's report is a best response to the other reports. Let $h_i = \max_{g \in M \setminus A_i} v_i(g)$, where the maximum is defined as zero if $M \setminus A_i$ is empty. Then $v_i(A_i) + \frac{n-1}{n} h_i \geq \frac{v_i(M)}{n}$.*

A natural approach would be to compare A_i with the bundle that agent i would get by reporting truthfully, as in the standard proof that recursively balanced picking sequences produce EF1 allocations under truthful reports. However, a truthful report may cost agent i some of her most valuable goods. Instead, we let her first secure all goods that she values more than every good she does not receive in A_i, and only then pick truthfully. The next lemma shows that such a deviation is always possible.

Lemma 1. *Fix a picking sequence and the reports of the other agents. Let $g_1, \dots, g_c$ be any distinct goods that agent i receives under report $\widehat{\succ}_i$, listed in the order she acquires them, and let $\widehat{\succ}'_i$ be any report ranking $g_1, \dots, g_c$ first, in this order. Then under $\widehat{\succ}'_i$ agent i receives g_r at her rth turn, for every $r \leq c$.*

Proof. If $c = 0$, the statement is vacuous. Suppose $c \geq 1$. Let S_t and S'_t be the sets of goods allocated before turn t under $\widehat{\succ}_i$ and $\widehat{\succ}'_i$. We show by induction on t, up to agent i's cth turn, that she has received $g_1, \dots, g_{r-1}$ at her turns before t and that $S'_t \subseteq S_t \cup \{g_1, \dots, g_{r-1}\}$, where $r-1$ is the number of her turns before t. Initially $S_1 = S'_1 = \emptyset$, so the invariant holds. Assume it holds immediately before turn t.

If turn t is agent i's, then $g_r \notin S_t$, since under $\widehat{\succ}_i$ she acquires g_r at her rth turn or later, and $g_r \notin \{g_1, \dots, g_{r-1}\}$. So g_r is available, and as $g_1, \dots, g_{r-1}$ are gone, it is the highest-ranked available good of $\widehat{\succ}'_i$; she takes it. Adding g_r to S'_t and to the list of promoted goods preserves the inclusion at the next turn. If turn t belongs to $j \neq i$, let x be j's pick under $\widehat{\succ}_i$. Then $x \notin S_t$ and $x \notin \{g_1, \dots, g_{r-1}\}$, since these goods are allocated to agent i under $\widehat{\succ}_i$. The induction hypothesis therefore implies that x is available under $\widehat{\succ}'_i$. Hence j takes x or a good ranked above x. Any such higher-ranked good belongs to S_t, since otherwise j would have taken it under $\widehat{\succ}_i$. Thus the new pick belongs to S_{t+1}, preserving the invariant. □

We now prove the quantitative bound.

Proof of Proposition 2. Let $H = \{g \in M : v_i(g) > h_i\}$. By the definition of h_i, $H \subseteq A_i$, and every good outside H is worth at most h_i to agent i. Let $g_1, \dots, g_c$ be the goods of H in the order agent i acquires

them. Recursive balance gives two facts: every agent's rth turn lies in round r, and the numbers of turns of any two agents differ by at most one.

To exploit the best-response condition, we consider a unilateral deviation by agent i, while keeping all other reports fixed. We construct this deviation so that, denoting its resulting allocation by $(D_1, \dots, D_n)$, all but one good of every D_j can be paired with distinct goods of D_i that agent i values at least as much. The remaining good is worth at most h_i. This gives a lower bound on $v_i(D_i)$; because the original report is a best response, agent i values her actual bundle A_i at least as much as D_i.

Consider the report $\widehat{\succ}_i'$ ranking $g_1, \dots, g_c$ first, in this order, followed by the remaining goods in non-increasing order of v_i, breaking ties lexicographically, and denote the resulting allocation by $(D_1, \dots, D_n)$. By Lemma 1, agent i receives g_r at her rth turn for $r \le c$; at every later turn all of H is gone, so she takes a most valuable available good. Let d_r be her rth pick.

Fix $j \neq i$. We pair every good of D_j except its first with an earlier pick of agent i. If $D_j = \emptyset$, then $v_i(D_j) \le v_i(D_i) + h_i$ is immediate. Otherwise, list the goods of D_j in allocation order as $y_1, \dots, y_q$, where $q \le |D_i| + 1$. We claim $v_i(d_r) \ge v_i(y_{r+1})$ for $1 \le r \le q-1$. If $r \le c$, then $d_r = g_r$ is worth more than h_i, while $y_{r+1} \notin H$ is worth at most h_i. If $r > c$, then y_{r+1} is taken in round $r+1$, hence is still available at agent i's rth turn, where she takes a most valuable available good. Therefore $v_i(d_r) \ge v_i(y_{r+1})$ in this case as well.

Since $y_1 \notin H$, we have $v_i(y_1) \le h_i$, and thus

$$\begin{aligned} v_i(D_j) &= v_i(y_1) + \sum_{r=1}^{q-1} v_i(y_{r+1}) \\ &\le h_i + \sum_{r=1}^{q-1} v_i(d_r) \\ &\le h_i + v_i(D_i). \end{aligned}$$

Here, the first inequality uses $v_i(y_{r+1}) \le v_i(d_r)$ for every $r \le q-1$. For the second inequality, recall that $q - 1 \le |D_i|$, so $d_1, \dots, d_{q-1}$ form a subset of D_i; the inequality therefore follows from the nonnegativity of the valuations.

Since $(D_1, \dots, D_n)$ partitions M,

$$v_i(M) = v_i(D_i) + \sum_{j \neq i} v_i(D_j) \le n\, v_i(D_i) + (n-1)h_i.$$

Hence,

$$v_i(D_i) + \frac{n-1}{n} h_i \ge v_i(M)/n\,.$$

Finally, $v_i(A_i) \ge v_i(D_i)$ because the original report is a best response. Hence,

$$v_i(A_i) + \frac{n-1}{n} h_i \ge v_i(M)/n\,. \qquad \square$$

For the pairwise guarantee, we show $\frac{1}{2}$-EF1 by considering the deviation in which agent i ranks the goods of A_j first, from the most to the least valuable to her. The difficulty is that this deviation may change the picks of the other agents, and such changes may propagate until the end of the execution. To control them, we compare the two executions turn by turn, using a potential function that can increase only at the turns of agents i and j. Roughly speaking, each good of A_j that agent i takes in the deviation pays for at most two turns of agent j in the original execution, while the first turn of agent j is paid for by a most valuable good of A_j, i.e., by the good that is removed in the EF1 comparison. This is where the factor $1/2$ comes from.

Proposition 3. *Fix a recursively balanced picking sequence independently of the reports. If agent i's report is a best response to the other reports and the resulting allocation is $A = (A_1, \dots, A_n)$, then, for every $j \neq i$ with $A_j \neq \emptyset$,*

$$v_i(A_i) \geq \frac{1}{2}\Big(v_i(A_j) - \max_{g \in A_j} v_i(g)\Big).$$

Proof. Fix $j \neq i$ with $A_j \neq \emptyset$, and set $h = \max_{g \in A_j} v_i(g)$. Suppose that agent i deviates to a strict ranking $\succ_i'$ that places all goods of A_j first in nonincreasing order of v_i, breaking ties lexicographically, followed by the goods outside A_j in any strict order. We denote the deviating allocation by $D = (D_1, \dots, D_n)$.

Let k be the number of turns of i, and let d_r be her deviating pick on her rth turn, for $1 \leq r \leq k$. Define

$$x_0 = h, \qquad x_r = \begin{cases} v_i(d_r), & d_r \in A_j, \\ 0, & d_r \notin A_j. \end{cases}$$

Whenever a good of A_j is available, i takes a most valuable such good. If she takes a good outside A_j, no good of A_j remains. Since available sets only shrink, it follows that

$$x_0 \geq x_1 \geq \cdots \geq x_k \geq 0.$$

After her rth turn, every remaining good of A_j has adjusted value at most x_r. Additivity and nonnegativity also give

$$v_i(D_i) = \sum_{r=1}^{k} v_i(d_r) \geq \sum_{r=1}^{k} x_r. \tag{1}$$

These statements also apply when $k = 0$, with empty sums.

Both executions use the same picking sequence and allocate one good at each turn. We synchronize them and write R, S for the respective sets of remaining goods in the original and deviating executions, respectively. For a nonnegative *level* λ, define

$$\Phi(R, S, \lambda) = v_i(A_j \setminus R) + \sum_{g \in R \setminus S} \max\{\lambda, \mathbb{1}_{A_j}(g) \cdot v_i(g)\}. \tag{2}$$

The first term records the value allocated to j in the original execution. The second term assigns a *residual value* to every good that is still available in the original execution but not in the deviating one. The level λ ensures that each such good has residual value at least λ, even if it is not a good of A_j.

Initially set $\lambda = x_0 = h$. Immediately before i's rth turn, change the level to x_r, and retain this level until her next turn. Every level change is a decrease, so it does not increase Φ right at that moment. Initially $R = S = M$, hence $\Phi = 0$. Below, we bound the change of Φ at each turn, for the level that is in effect at that turn.

A turn of agent i: Let a_r be her original pick. Since $j \neq i$, we have $a_r \notin A_j$, so the first term of the potential does not change. The new residual value set is

$$((R \setminus S) \setminus \{a_r\}) \cup (\{d_r\} \cap (R \setminus \{a_r\})).$$

Thus a residual value term may be removed, and at most one is added, for d_r (only if it remains available in the original execution, i.e., if $d_r \in R \setminus \{a_r\}$). Because $\mathbb{1}_{A_j}(d_r) \cdot v_i(d_r) = x_r = \lambda$, the added residual value equals x_r. This turn therefore increases Φ by at most x_r.

A turn of an agent $\ell \neq i$: Let $u \in R$ and $z \in S$ be her original and deviating picks. Every good $g \in R \cap S$ satisfies $\mathbb{1}_{A_j}(g) \cdot v_i(g) \leq \lambda$. Before i's first turn this follows from $\lambda = h$; subsequently, it follows from the bound on goods remaining after her most recent deviating pick.

If $u \neq z$, the two goods cannot both be in $R \cap S$; otherwise the unchanged strict ranking $\widehat{\succ}_\ell$ would have to rank u above z in the original execution and z above u in the deviating execution. So, if $u \in R \cap S$, then either $u = z$ or $z \in S \setminus R$; in both cases the set $R \setminus S$ does not change. If $u \in R \setminus S$, then the residual value of u is removed, and $z \neq u$, as $u \notin S$. In this case, no residual value is added if $z \in S \setminus R$, whereas a residual value of $\max\{\mathbb{1}_{A_j}(z) \cdot v_i(z), \lambda\} = \lambda$ is added if $z \in R \cap S$.

If $\ell \neq j$, then $u \notin A_j$. The first term of Φ is unchanged. In the case where $u \in R \cap S$, the residual values are unchanged, whereas in the case where $u \in R \setminus S$, a residual value of λ is removed and at most a residual value of λ is added. Thus this turn does not increase Φ.

If $\ell = j$, then $u \in A_j$ and the first term increases by $v_i(u)$. When $u \in R \cap S$, this is at most λ, and residual values are unchanged. Otherwise (i.e., when $u \in R \setminus S$), by the discussion above, the total increase is at most

$$v_i(u) - \max\{v_i(u), \lambda\} + \lambda \leq \lambda.$$

Hence, every turn of j increases Φ by at most its current level.

Now we may bound the final value of Φ by summing over turns. Let q be the total number of turns of j; by recursive balance, we have $q \leq k + 1$. Before j's sth turn, suppose that i has taken p turns; again, by recursive balance, we have $p \geq s - 1$. The current level is x_p, and therefore

$$\lambda = x_p \leq x_{s-1}.$$

Note that this includes $s = 1$, as $x_0 = h$.

At termination $R = S = \emptyset$, so the final potential is $v_i(A_j)$. Summing the turn increases, including the nonpositive level changes, and then using $q - 1 \leq k$, we obtain

$$\begin{aligned} v_i(A_j) &= \Phi(\emptyset, \emptyset, x_k) \\ &\leq \Phi(M, M, x_0) + \sum_{r=1}^{k} x_r + \sum_{s=1}^{q} x_{s-1} \\ &= \sum_{r=1}^{k} x_r + h + \sum_{r=1}^{q-1} x_r \leq h + 2\sum_{r=1}^{k} x_r \\ &\leq h + 2v_i(D_i) \leq h + 2v_i(A_i). \end{aligned}$$

The second to last inequality is (1), whereas the last one uses the best-response condition for $\widehat{\succ}'_i$. Rearranging the terms gives the claimed inequality. □

We are now ready to prove Theorem 1 using the two propositions above.

Proof of Theorem 1. Let A be an allocation induced by a pure Nash equilibrium. Every agent is a best responder, so Proposition 3 applies to every ordered pair of distinct agents. For $A_j \neq \emptyset$, choose a good $g \in A_j$ maximizing $v_i(g)$. By additivity,

$$v_i(A_i) \geq \frac{1}{2}\left(v_i(A_j) - v_i(g)\right) = \frac{1}{2}v_i(A_j \setminus \{g\}).$$

Thus A is $1/2$-EF1.

For PROP1, fix i and use h_i from Proposition 2, i.e., $h_i = \max_{g \in M \setminus A_i} v_i(g)$. If $v_i(A_i) \geq v_i(M)/n$, proportionality already holds; otherwise, $M \setminus A_i$ must be nonempty. Consider a good $g \in M \setminus A_i$ with $v_i(g) = h_i$. Proposition 2 and nonnegativity give

$$v_i(A_i \cup \{g\}) = v_i(A_i) + h_i \geq v_i(A_i) + \frac{n-1}{n}h_i \geq \frac{v_i(M)}{n}.$$

Thus A is also PROP1. □

Corollary 1. *For two agents with nonnegative additive valuations, every pure Nash equilibrium of every fixed recursively balanced picking sequence mechanism induces an EF1 allocation.*

Proof. For $j \neq i$, Proposition 2 with $n = 2$ gives $2v_i(A_i) + h_i \geq v_i(M) = v_i(A_i) + v_i(A_j)$. Since $M \setminus A_i = A_j$, choosing a most valuable good of a nonempty A_j proves EF1. □

For Round-Robin, this was already known from [2, Theorem 3].

Approximate equilibria. The proofs above also give bounds when best responses are approximate, using the exact same analysis. For $\beta \in (0, 1]$, suppose that the report profile is a *β-approximate pure Nash equilibrium*, i.e., every agent i satisfies $v_i(A_i) \geq \beta v_i(D_i)$ for every unilateral deviation with resulting allocation D. The same deviations then give

$$v_i(A_i) \geq \frac{\beta}{2}\Big(v_i(A_j) - \max_{g \in A_j} v_i(g)\Big), \quad \text{for } j \neq i \text{ with } A_j \neq \emptyset,$$

and

$$v_i(A_i) + \beta\frac{n-1}{n}h_i \geq \beta\frac{v_i(M)}{n}.$$

Thus the allocation is $\beta/2$-EF1, and each agent can reach at least $\beta v_i(M)/n$ by adding at most one outside good. For Round-Robin, the pairwise guarantee is already known due to [3, Theorem 8].

Our counterexample separates exact EF1 from the two equilibrium guarantees of Theorem 1. EF1 may fail, whereas $1/2$-EF1 and PROP1 hold for every pure Nash equilibrium of every fixed recursively balanced mechanism. There also always exists an EF1 pure Nash equilibrium: the bluff profile reproduces the truthful allocation [6, Theorem 2].

Theorem 2 of Aziz et al. [6] is stated for positive values that induce strict rankings, but it extends to our setting. Refine true-value ties lexicographically and let $r_i(g)$ be a good's rank, with rank 1 denoting the most valuable good. For $\delta > 0$, replace $v_i(g)$ by $v_i(g) + \delta(m + 1 - r_i(g))$. The true strict rankings, truthful allocation, and bluff profile are unchanged. The profile is a pure Nash equilibrium for every $\delta > 0$; taking $\delta \downarrow 0$ in each best-response inequality proves that it is a pure Nash equilibrium for the original values as well. For Round-Robin, this was also shown in [2, Theorem 10].

5 An upper bound of $2/3$ for Round-Robin

The counterexample in Proposition 1 already gives an upper bound of $15/16$: agent 2 receives value 15, while every deletion from agent 1's bundle leaves value 16. This ratio is close to one. The example establishes the failure of exact EF1 with very few goods, but leaves open how much envy an equilibrium can sustain.

We now improve this upper bound to $2/3$. The construction has two principal agents, agent 1 and agent 2. Agent 2 is the agent whose envy we examine; she receives goods that she values at 2, whereas agent 1 receives goods that agent 2 values at 3. Agent 1 picks before agent 2 in every round and ranks her own goods first. So, if agent 2 takes some of these goods, then agent 1 runs out of them earlier and starts taking the goods of agent 2. At the same time, each of the remaining agents, whom we call *auxiliary*, removes a specific good of agent 2 by a prescribed deadline, unless agent 2 has already taken it. Together, these two effects force agent 2 to lose at least three of her own goods for every two goods of agent 1 that she takes. Therefore, no deviation is profitable for her.

The auxiliary agents implement the deadlines through fixed reports, while goods that only they value fill their remaining turns without disturbing the intended allocation. Giving unit valuations to all agents other than agent 2 makes their equilibrium conditions immediate. This leaves a single best-response condition to verify, namely that of agent 2.

Theorem 2. *For every integer $k \geq 4$, there is an additive instance with $k+1$ agents and $k(k+1)$ goods, and a pure Nash equilibrium of Round-Robin whose allocation is not α-EF1 for any*

$$\alpha \in \Big(\frac{2k}{3(k-1)}, 1\Big].$$

Consequently, it is not possible to guarantee α-EF1 for every pure Nash equilibrium of Round-Robin for any constant $\alpha > 2/3$.

Proof. We begin by describing the instance. There are $n = k+1$ agents: the principal agents 1 and 2, and the auxiliary agents $3, \ldots, k+1$, who pick in this order in every round. The set of goods is $M = H \cup G \cup Z_1 \cup \cdots \cup Z_{k-1}$, where

$$H = \{h_1, \ldots, h_k\}, \qquad G = \{g_1, \ldots, g_k\}, \qquad Z_t = \{z_{t,1}, \ldots, z_{t,k}\} \text{ for } 1 \leq t \leq k-1,$$

and all these sets are pairwise disjoint. The goods of Z_t are meant for agent $t+2$. We have $m = 2k + (k-1)k = kn$, so there are exactly k complete rounds and every agent receives exactly k goods, regardless of the reports. The valuation function of agent 2 is additive with

$$v_2(g) = \begin{cases} 2, & g \in H, \\ 3, & g \in G, \\ 0, & \text{otherwise}, \end{cases}$$

whereas every other agent values every good at 1.

For the reports of the agents, we first fix an arbitrary strict ranking of M. Every ranking displayed below is completed by appending all unlisted goods to its right, in this fixed order. The reports are

$$\begin{aligned} \widehat{\succ}_1 &: \quad g_1 \succ \cdots \succ g_k \succ h_1 \succ \cdots \succ h_k \succ \cdots, \\ \widehat{\succ}_2 &: \quad h_1 \succ \cdots \succ h_k \succ \cdots, \\ \widehat{\succ}_{t+2} &: \quad z_{t,1} \succ \cdots \succ z_{t,t-1} \succ h_t \succ z_{t,t} \succ \cdots \succ z_{t,k} \succ \cdots, \qquad 1 \leq t \leq k-1, \end{aligned}$$

where, for $t = 1$, no good of Z_1 precedes h_1 in agent 3's ranking. These are strict rankings and, as discussed in Section 2, they can be implemented by bid vectors.

We first argue about the resulting allocation. We claim that, in every round r, agent 1 takes g_r, agent 2 takes h_r, and agent $t+2$ takes $z_{t,r}$, for every $1 \leq t \leq k-1$. We use induction on r. Suppose that the claim holds for all rounds before round r; for $r = 1$ this assumption is vacuous. Then the goods allocated before round r are exactly g_s, h_s, and $z_{t,s}$, for $s < r$ and $1 \leq t \leq k-1$. In particular, g_r and h_r are available at the beginning of round r, whereas all goods that precede them in $\widehat{\succ}_1$ and $\widehat{\succ}_2$, respectively, have been allocated. Hence agent 1 takes g_r and agent 2 takes h_r. Now consider agent $t+2$. Since the agents before her in round r have taken g_r, h_r, and goods of $Z_1, \ldots, Z_{t-1}$, the goods of Z_t that are not available at her turn are exactly $z_{t,1}, \ldots, z_{t,r-1}$; in particular, $z_{t,r}$ is available. If $r < t$, the goods that she ranks above $z_{t,r}$ are $z_{t,1}, \ldots, z_{t,r-1}$. If $r \geq t$, these are $z_{t,1}, \ldots, z_{t,r-1}$ and h_t, where h_t has been taken by agent 2 in round t. In both cases, all goods that she ranks above $z_{t,r}$ are unavailable, so she takes $z_{t,r}$. This completes the induction. Therefore,

$$A_1 = G, \qquad A_2 = H, \qquad A_{t+2} = Z_t \text{ for } 1 \leq t \leq k-1,$$

and, in particular, $v_2(A_2) = 2k$.

Next, we show that the profile above is a pure Nash equilibrium. Every agent other than agent 2 receives k goods of value 1 to her. As every agent receives exactly k goods under every report profile, this is the maximum value she can obtain, and thus her report is a best response. It remains to show that $\widehat{\succ}_2$ is a best response for agent 2. To this end, consider any report of agent 2, keep the reports of all other agents fixed, and let $D = (D_1, \dots, D_n)$ be the resulting allocation. We are going to show that $v_2(D_2) \leq 2k$.

We begin by showing that, for every $1 \leq t \leq k-1$, the good h_t is allocated by the end of round t. If h_t is not available at the turn of agent $t+2$ in round t, then it has already been allocated. So, suppose that h_t is available at that turn. Then h_t was also available at each of her $t-1$ earlier turns, so at each of these turns she took a good that she ranks above h_t, i.e., one of $z_{t,1}, \dots, z_{t,t-1}$. These $t-1$ distinct picks exhaust $\{z_{t,1}, \dots, z_{t,t-1}\}$, and therefore h_t is the highest-ranked available good of agent $t+2$ in round t; hence, she takes it. Note that this argument applies to every report of agent 2, including reports under which she takes goods of Z_t in order to change the later picks of agent $t+2$.

If $|D_2 \cap G| = 0$, then $v_2(D_2) = 2|D_2 \cap H| \leq 2k$. So, assume that $|D_2 \cap G| \geq 1$. As agent 1 ranks the goods of G first, at each of her turns she takes a good of G if and only if such a good is available. Since the set of available goods only shrinks, the rounds in which a good of G is available at agent 1's turn are the first few rounds. As agent 1 has one turn per round and receives $|D_1 \cap G|$ goods of G, these are exactly the rounds $1, \dots, |D_1 \cap G|$. Moreover, as agent 1 picks before agent 2 in every round, if agent 2 takes a good of G in some round r, then this good was available at agent 1's turn in round r, and thus $r \leq |D_1 \cap G|$. Since agents 1 and 2 receive distinct goods of G, we have $|D_1 \cap G| + |D_2 \cap G| \leq k$; in particular, $|D_1 \cap G| \leq k-1$.

We now bound $|D_2 \cap H|$. In rounds $1, \dots, |D_1 \cap G|$, agent 2 has exactly $|D_1 \cap G|$ turns, and, by the above, she takes all $|D_2 \cap G|$ goods of $D_2 \cap G$ at these turns. So, she receives at most $|D_1 \cap G| - |D_2 \cap G|$ goods of H in these rounds. Next, consider rounds $|D_1 \cap G| + 1, \dots, k$. As $|D_1 \cap G| \leq k-1$, the goods $h_1, \dots, h_{|D_1 \cap G|}$ have been allocated by the end of round $|D_1 \cap G|$, and thus every good of H allocated in these rounds belongs to $\{h_{|D_1 \cap G|+1}, \dots, h_k\}$. Suppose that agent 2 takes some $h \in H$ in a round $r \geq |D_1 \cap G| + 1$. Then h was available at agent 1's turn in round r, whereas no good of G was, since $r > |D_1 \cap G|$. As agent 1 ranks the goods of H right after those of G, she took a good of H at that turn, which is different from h, since h was still available afterwards. That is, in every round $r \geq |D_1 \cap G| + 1$ in which agent 2 receives a good of H, agent 1 also receives a good of H. As every agent receives one good per round, in rounds $|D_1 \cap G| + 1, \dots, k$ agent 1 receives at least as many goods of H as agent 2 does. All these goods are distinct and belong to $\{h_{|D_1 \cap G|+1}, \dots, h_k\}$, so agent 2 receives at most $(k - |D_1 \cap G|)/2$ goods of H in these rounds. Combining the above, we get

$$\begin{aligned} v_2(D_2) &= 3|D_2 \cap G| + 2|D_2 \cap H| \\ &\leq 3|D_2 \cap G| + 2(|D_1 \cap G| - |D_2 \cap G|) + (k - |D_1 \cap G|) \\ &= k + |D_1 \cap G| + |D_2 \cap G| \leq 2k, \end{aligned}$$

where the last inequality follows from $|D_1 \cap G| + |D_2 \cap G| \leq k$.

Since $v_2(A_2) = 2k$, the report $\widehat{\succ}_2$ is a best response for agent 2. Together with the best responses of the other agents, this shows that the profile is a pure Nash equilibrium.

Finally, we show that the allocation $(A_1, \dots, A_n)$ fails to be α-EF1 for sufficiently large α. Every good of $A_1 = G$ is worth 3 to agent 2. Thus, for every $g \in A_1$, we have $v_2(A_1 \setminus \{g\}) = 3(k-1)$, while $v_2(A_2) = 2k$. That is, the allocation is not α-EF1 for any $\alpha > 2k/(3(k-1))$, and, since $k \geq 4$, we have $2k/(3(k-1)) < 1$. Note that

$$\frac{2k}{3(k-1)} = \frac{2}{3} + \frac{2}{3(k-1)}.$$

So, given any $\alpha \in (2/3, 1]$, we may choose an integer $k \geq 4$ so that $2k/(3(k-1)) < \alpha$, and the corresponding pure Nash equilibrium is not α-EF1. □

Note that the upper bound also holds for strictly positive valuations. For each k, we may increase the value of agent 2 for each good by $1/k$, leaving the reports and all other valuations unchanged. Since agent 2 receives exactly k goods under every report profile, the value of every bundle she may receive increases by exactly 1, and thus the profile remains an equilibrium. The EF1 comparison now gives

$$\frac{2k+1}{(3+1/k)(k-1)} \longrightarrow \frac{2}{3},$$

as $k \to \infty$. Hence, even strictly positive instances rule out every factor strictly greater than $2/3$.

6 Discussion and open problems

Our results leave a gap between the guaranteed factor of $1/2$ and the upper bound of $2/3$ established here with respect to EF1. To state the question precisely, let α_{RR} be the largest factor guaranteed at every pure Nash equilibrium of every Round-Robin instance, and let α_{RB} be the corresponding factor for all fixed recursively balanced sequences. Since Round-Robin belongs to the latter class,

$$\frac{1}{2} \leq \alpha_{\mathrm{RB}} \leq \alpha_{\mathrm{RR}} \leq \frac{2}{3}.$$

Our central open question is to determine these factors. In particular, does allowing the order to change between rounds affect the sharp guarantee? The lower-bound proof uses only that no agent is more than one turn ahead of another, whereas the upper-bound construction uses the same order in every round.

The dependence on the number of agents is another natural question. For two agents, the quantitative proportionality bound gives exact EF1 (Corollary 1). The small counterexample shows that this fails already for three agents, but the family approaching $2/3$ uses a growing number of auxiliary agents. Determining the best guaranteed factor for each fixed number of agents could distinguish which part of the gap comes from additional strategic interaction and which part comes from the number of rounds.

There is also a useful distinction between equilibrium existence and a guarantee for every equilibrium. The bluff profile provides an EF1 equilibrium, while the examples show that other equilibria can sustain envy after every deletion. It would be interesting to identify conditions on reports or valuations that force all equilibria to satisfy a stronger guarantee, without restricting the agents to truthful reports.

Acknowledgments

Rebecca Reiffenhäuser was partially supported by the NWO grant *Fair Allocation for Strategic Agents (FAStA)*, grant ID https://doi.org/10.61686/AZZOJ38639.

References

[1] Georgios Amanatidis, Georgios Birmpas, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Allocating indivisible goods to strategic agents: Pure Nash equilibria and fairness. In *Proceedings of the 17th International Conference on Web and Internet Economics (WINE)*, volume 13112 of *Lecture Notes in Computer Science*, pages 149–166. Springer, 2021.

[2] Georgios Amanatidis, Georgios Birmpas, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Allocating indivisible goods to strategic agents: Pure Nash equilibria and fairness. *Mathematics of Operations Research*, 49(4):2425–2445, 2024.

[3] Georgios Amanatidis, Georgios Birmpas, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Round-robin beyond additive agents: Existence and fairness of approximate equilibria. *Mathematics of Operations Research*, 50(4):2717–2737, 2025.

[4] H. Aziz, T. Walsh, and L. Xia. Possible and necessary allocations via sequential mechanisms. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 468–474, 2015.

[5] Haris Aziz, Sylvain Bouveret, Jérôme Lang, and Simon Mackenzie. Complexity of manipulating sequential allocation. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI 2017*, pages 328–334. AAAI Press, 2017.

[6] Haris Aziz, Paul Goldberg, and Toby Walsh. Equilibria in sequential allocation. In *Proceedings of the 5th International Conference on Algorithmic Decision Theory, ADT 2017*, volume 10576 of *Lecture Notes in Computer Science*, pages 270–283. Springer, 2017.

[7] S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 73–78. AAAI Press, 2011.

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

[9] D. A. Kohler and R. Chandrasekaran. A class of sequential games. *Operations Research*, 19(2):270–277, 1971.