

Fair Division with Two-Sided Preferences

Ayumi Igarashi¹, Yasushi Kawase¹, Warut Suksompong² and Hanna Sumita³

¹University of Tokyo

²National University of Singapore

³Tokyo Institute of Technology

{igarashi,kawase}@mist.i.u-tokyo.ac.jp, warut@comp.nus.edu.sg, sumita@c.titech.ac.jp

Abstract

We study a fair division setting in which a number of players are to be fairly distributed among a set of teams. In our model, not only do the teams have preferences over the players as in the canonical fair division setting, but the players also have preferences over the teams. We focus on guaranteeing envy-freeness up to one player (EF1) for the teams together with a stability condition for both sides. We show that an allocation satisfying EF1, swap stability, and individual stability always exists and can be computed in polynomial time, even when teams may have positive or negative values for players. Similarly, a balanced and swap stable allocation that satisfies a relaxation of EF1 can be computed efficiently. When teams have nonnegative values for players, we prove that an EF1 and Pareto optimal allocation exists and, if the valuations are binary, can be found in polynomial time. We also examine the compatibility between EF1 and justified envy-freeness.

1 Introduction

The new season of a youth sports league is starting in three months, and the league organizers need to allocate the available players to the participating teams. How can they accomplish this task in a satisfactory way, so that all parties involved can look forward to the upcoming season instead of grumble about the allocation?

A principal consideration in such allocation tasks is fairness, and the problem of fairly dividing resources (in this case, players) among interested recipients (in this case, teams) has long been studied under the name of *fair division* [Brams and Taylor, 1996; Moulin, 2003]. One of the most prominent fairness notions is *envy-freeness*, which means that no team should envy another team based on the sets of players that they receive.¹ Even though an envy-free allocation may not exist (e.g., if there is one highly-coveted superstar), an intuitive relaxation called *envy-freeness up to one player*

(EF1)—that is, any envy that one team has toward another team can be eliminated upon the removal of some player in the envied team—can always be fulfilled [Lipton *et al.*, 2004]. Another relevant criterion is *balancedness*, which requires the players to be distributed as equally among the teams as possible.² Balancedness can be especially desirable when allocating players to sports teams, as each team may need to have a fixed number of players due to the rules of the sport. Assuming that teams have additive and nonnegative values for players, an allocation that is both EF1 and balanced always exists and can be found efficiently via a simple round-robin algorithm (see, e.g., [Caragiannis *et al.*, 2019, p. 7]). In fact, this algorithm forms the basis of *draft* processes used in many sports leagues around the world.³

While EF1 provides a strong fairness guarantee with respect to the teams’ preferences, it overlooks the fact that the players may have preferences over the teams as well, for example, depending on their familiarity with the team managers or the proximity of their residence to the training grounds. Clearly, ignoring the preferences of the players may lead to a suboptimal allocation. As an extreme case, if every team is indifferent between all players, then swapping a pair of players keeps the teams as happy as before and may make both of the swapped players much happier. In addition to our sports league example, two-sided preferences also occur in the allocation of employees to different branches of a restaurant chain or volunteers to community service clubs. Moreover, the player preferences could signify the *suitability* of the teams for the players—for instance, the “players” could represent tasks (such as household chores or papers to be reviewed) and the “teams” have varying levels of ability to perform the tasks. Can we find an allocation that is fair to the teams and at the same time satisfies a stability condition with respect to the preferences of both sides?

1.1 Our Results

As is common in fair division, we assume that the teams have additive valuations over the players. Some of our results allow these values to be either positive or negative; this corresponds to the allocation of *indivisible goods and chores* [Aziz

¹Due to the setting that we study, throughout this paper we will use the terms *team* and *player* in place of the standard fair division terms *agent* and *item*, respectively.

²One could view balancedness as EF1 with respect to the number of allocated players.

³Please refer to [http://wikipedia.org/wiki/Draft_\(sports\)](http://wikipedia.org/wiki/Draft_(sports))

| Properties | Existence |
|---|----------------|
| EF[1,1] + balanced + swap stable | Yes (Thm. 3.2) |
| EF1 + balanced | No (Prop. 3.7) |
| balanced + individually stable | No (Prop. 3.8) |
| EF1 + swap stable + individually stable | Yes (Thm. 3.9) |
| EF1 + PO (nonnegative-value players) | Yes (Thm. 4.5) |
| EF1 + PO + team-PO (two teams) | Yes (Thm. 4.3) |
| EF1 + player-PO | No (Prop. 4.4) |
| EF1 + justified EF | No (Prop. 5.1) |

Table 1: Summary of our results on whether each combination of properties can always be satisfied simultaneously, with the corresponding theorem or proposition number.

et al., 2022]. For consistency of terminology, we will use the terms *nonnegative-value players* and *nonpositive-value players* instead of *goods* and *chores*, respectively. We formally define the notions that we consider and outline some relationships between them in Section 2.

In Section 3, we focus on *swap stability*, the requirement that no swap between two players makes at least one of the four involved parties better off and none of them worse off. First, we observe that even with nonnegative-value players, starting with an arbitrary EF1 allocation and letting players make beneficial swaps may result in an allocation that violates EF1. Despite this fact, for teams with arbitrary (positive or negative) values over players, we present a polynomial-time algorithm that produces a balanced and swap stable allocation satisfying $EF[1,1]$, a relaxation of EF1 where one player may be removed from each of the envying team and the envied team.⁴ Since EF[1,1] reduces to EF1 for nonnegative-value players as well as for nonpositive-value players, we obtain the same result for EF1 in each of these cases. We then note two ways in which the arbitrary-value result cannot be improved: EF[1,1] cannot be strengthened to EF1, and we cannot simultaneously attain *individual stability*—the condition that no deviation of a player to another team makes the player better off and neither of the involved teams worse off. Nevertheless, we show that if we give up balancedness, both of these improvements become possible: an allocation that satisfies EF1, swap stability, and individual stability exists and can be found efficiently.

Next, in Section 4, we consider the notion of *Pareto optimality (PO)*—no allocation can make a party (i.e., either player or team) better off without making another party worse off—which is stronger than both swap stability and individual stability. We prove that deciding whether an allocation is PO or not is coNP-complete even for two teams with identical valuations, nonnegative-value players, and a balanced allocation. On the other hand, for two teams with arbitrary valuations, we show that an extension of the *generalized adjusted winner procedure* of Aziz *et al.* [2022] computes an

⁴When both positive and negative values are allowed, EF1 permits one player to be removed from either the envying team or the envied team (but not both) [Aziz *et al.*, 2022].

EF1 and PO allocation in polynomial time. For any number of teams and nonnegative-value players, we observe that an EF1 and PO allocation always exists. Moreover, we demonstrate that such an allocation can be found efficiently in two special cases: (i) the teams have binary valuations over the players, and (ii) there are three teams with identical valuations, and each player has a favorite team and is indifferent between the other two teams. We also provide a pseudopolynomial-time algorithm when the number of teams is constant.

Finally, in Section 5, we examine *justified envy*, a stability notion from the two-sided matching literature: player p_i is said to have justified envy toward another player p_j assigned to team k if k prefers p_i to p_j and p_i prefers k to her assigned team. Perhaps surprisingly, we show that an EF1 and justified envy-free allocation may not exist, even for two teams and nonnegative-value players who all prefer the same team. We then prove that deciding whether such an allocation exists is NP-complete even for nonnegative-value players who have strict preferences over the teams. On the other hand, if one adds the condition that there are two teams, we show that the problem becomes polynomial-time solvable.

Our (non-)existence results are summarized in Table 1.

1.2 Related Work

Even though fair division has given rise to a sizable literature, the vast majority of the literature assumes one-sided preferences—in our terminology, the teams have preferences over the players, but not vice versa. A small number of recent papers have combined fairness concepts with two-sided preferences. Freeman *et al.* [2021] considered many-to-many matching and proposed the notion of *double envy-freeness up to one match (DEF1)*, which requires EF1 to hold for both sides simultaneously. Note that in our many-to-one setting, DEF1 is meaningless on the player side because it is always trivially satisfied. Gollapudi *et al.* [2020] studied many-to-many matching in a dynamic setting; their positive results primarily hold for symmetric binary valuations, which are much more restrictive than the valuations that we allow. Patro *et al.* [2020] investigated fairness in two-sided platforms between producers and customers, but assumed that producers are indifferent between customers. We stress that none of these papers studied a model that suitably captures our motivating examples such as the allocation of sports players to teams.

While most fair division papers assume that the items (in our terminology, players) are goods and some assume that they are chores, a recent line of work has relaxed these assumptions by allowing items to be either goods or chores, with this evaluation possibly varying across agents (in our terminology, teams) [Aleksandrov and Walsh, 2020; Bérczi *et al.*, 2020; Kulkarni *et al.*, 2021; Aziz *et al.*, 2022].

Lastly, we remark that justified envy is commonly studied in two-sided matching. Indeed, it forms the basis of the stability notion in one-to-one matching [Gale and Shapley, 1962], and has also been used in many-to-one matching [Wu and Roth, 2018; Yokoi, 2020].

2 Preliminaries

For each positive integer z , let $[z] := \{1, \dots, z\}$. Let $T = [n]$ be the set of teams and $P = \{p_1, \dots, p_m\}$ the set of players; we sometimes refer to either a team or a player as a *party*. Each player $p \in P$ has a weak transitive preference \succsim_p over the teams; denote by \succ_p and \sim_p the strict and equivalence part of \succsim_p , respectively.⁵ The *rank* of a team i for a player p is defined as 1 plus the number of teams j such that $j \succ_p i$. Each team $i \in T$ has a valuation function (or utility function) $v_i: 2^P \rightarrow \mathbb{R}$ over subsets of players. We assume that the valuations are additive, i.e., $v_i(P') = \sum_{p \in P'} v_i(\{p\})$ for all $i \in T$ and $P' \subseteq P$. For convenience, we write $v_i(p)$ instead of $v_i(\{p\})$. An *instance* consists of the teams and players, as well as the valuations and preferences of both sides. Sometimes we will consider the setting of *nonnegative-value players* (resp., *nonpositive-value players*), which means that $v_i(p) \geq 0$ (resp., $v_i(p) \leq 0$) for all $i \in T$ and $p \in P$.

An *allocation* $A = (A_1, A_2, \dots, A_n)$ is an ordered partition of P into n parts, where the part A_i is assigned to team i . We will investigate several fairness and stability notions for allocations. A basic fairness consideration on the team side is (almost) *envy-freeness*.

Definition 2.1. An allocation A is said to satisfy

- *EF1* if for all distinct $i, j \in T$, it holds that $v_i(A_i \setminus X) \geq v_i(A_j \setminus Y)$ for some $X \subseteq A_i$ and $Y \subseteq A_j$ with $|X \cup Y| \leq 1$;
- *EF[1,1]* if for all distinct $i, j \in T$, it holds that $v_i(A_i \setminus X) \geq v_i(A_j \setminus Y)$ for some $X \subseteq A_i$ and $Y \subseteq A_j$ with $|X|, |Y| \leq 1$.

EF1 was first studied for nonnegative-value players by Lip-ton et al. [2004] and subsequently extended to arbitrary-value players by Aziz et al. [2022], while EF[1,1] was recently introduced by Shoshan et al. [2022]. It follows immediately from the definition that EF1 implies EF[1,1]. Moreover, if all players yield nonnegative value, there is no reason to remove a player from A_i , so EF1 and EF[1,1] coincide in this case; an analogous statement holds for nonpositive-value players with A_i replaced by A_j .

Our next criterion is *balancedness*, which requires the players to be distributed as equally among the teams as possible.

Definition 2.2. An allocation A is said to be *balanced* if $||A_i| - |A_j|| \leq 1$ for all $i, j \in T$.

Observe that if there exists a constant $c \neq 0$ such that $v_i(p) = c$ for all $i \in T$ and $p \in P$, then both EF1 and EF[1,1] coincide with balancedness.

We now define stability concepts, several of which take into account the preferences of both sides. We say that a party is *better off* (resp., *worse off*) if it receives a better (resp., worse) outcome with respect to its valuation function (for a team) or preference (for a player).

Definition 2.3. Given an allocation A , a swap between players $p \in A_i$ and $q \in A_j$ (for some $i, j \in T$) is a *beneficial swap* if it makes at least one of the four involved parties better off and none of them worse off. A deviation of a player p

to another team is a *beneficial deviation* if it makes p better off and neither of the teams involved worse off.

An allocation A is said to be *swap stable* if it does not admit a beneficial swap. It is said to be *individually stable* if it does not admit a beneficial deviation.⁶

Definition 2.4. An allocation A is said to be *Pareto dominated* by another allocation A' if no party is worse off in A' than in A and at least one party is better off; in this case, A' is a *Pareto improvement* of A . An allocation A is *Pareto optimal* (*PO*) if it is not Pareto dominated by any other allocation.

We define *team-Pareto dominated*, *team-Pareto optimal* (*team-PO*), *player-Pareto dominated*, and *player-Pareto optimal* (*player-PO*) similarly, with “party” replaced by “team” and “player”, respectively.

Although PO clearly implies both swap stability and individual stability, it implies neither team-PO nor player-PO.

Proposition 2.5. *PO does not necessarily imply team-PO.*

Proposition 2.6. *PO does not necessarily imply player-PO.*

The proofs of Propositions 2.5 and 2.6, along with all other omitted proofs, can be found in the full version of our paper [Igarashi et al., 2022].

Finally, we define the concept of *justified envy*.

Definition 2.7. Given an allocation A , a player $p \in A_i$ is said to have *justified envy* toward a player $q \in A_j$ if $j \succ_p i$ and $v_j(p) > v_j(q)$. An allocation A is *justified envy-free* (*justified EF*) if no player has justified envy toward another player.⁷

3 Swap Stability

In this section, we focus on swap stability. A natural idea for obtaining an EF1 and swap stable allocation is to start with an arbitrary EF1 allocation and let players swap as long as a beneficial swap exists. Note that determining whether beneficial swaps exist (and, if so, finding such a swap) can be done in polynomial time since we can simply check all pairs of players. However, as can be seen in the following example, this approach does not always result in an EF1 allocation, even for nonnegative-value players.

Example 3.1. Consider the following instance with $n = 3$ and $m = 6$:

- $v_i(p_j) = 0$ for $i \in [2]$ and $j \in [6]$;
- $v_3(p_1) = v_3(p_2) = 1$ and $v_3(p_3) = v_3(p_4) = v_3(p_5) = v_3(p_6) = 0$;
- each player has a unique favorite team and is indifferent between the other two teams: p_1 and p_2 prefer team 1, p_4 and p_5 prefer team 2, and p_3 and p_6 prefer team 3.

⁶This is analogous to the notion of *contractual individual stability* in hedonic games [Aziz and Savani, 2016]. If we only require that the deviation does not make the player’s new team worse off (as in *individual stability* in hedonic games), then supposing that all players yield nonnegative value, the only stable allocations are ones in which every player is assigned to one of her most preferred teams.

⁷One could also consider a weaker version of justified envy where one of the two sides may be indifferent—this leads to a stronger version of justified EF. Our non-existence result (Proposition 5.1) carries over to this version as well.

⁵All notions considered in this paper take into account only the players’ ordinal preferences, so we do not assume cardinal utilities.

Algorithm 1: For computing an EF[1,1], swap stable, and balanced allocation

- 1 Construct a complete bipartite graph $G = (Q, P; E)$ with weight function $w: E \rightarrow \mathbb{R}$ where $Q = [m]$ and $w(q, p) = v_{f(q)}(p)$ for each $(q, p) \in Q \times P$; // The set Q mimics the available positions within the teams
 - 2 Compute a perfect matching $\mu \subseteq Q \times P$ such that the weight of the edge adjacent to vertex $1 \in Q$ is as large as possible, and subject to this condition, the weight of the edge adjacent to vertex $2 \in Q$ is as large as possible, and so on until vertex $m \in Q$; // Simulate round-robin with only the players' values being assigned to teams
 - 3 Let $E^* = \{(q, p) \in Q \times P \mid w(q, p) = w(q, \mu_q)\}$; // Keep edges that are as good for the teams as those in μ
 - 4 Compute a perfect matching μ^* in $G^* = (Q, P; E^*)$ such that the sum over all players $p \in P$ of the rank of team $f(\mu_p^*)$ for player p is minimized;
 - 5 Return the allocation A such that p is allocated to team $f(q)$ for each $(q, p) \in \mu^*$;
-

The allocation $A = (\{p_1, p_4\}, \{p_2, p_5\}, \{p_3, p_6\})$ is EF1. The swap between p_2 and p_4 is the unique beneficial swap; let A' be the allocation after this swap. The allocation A' is swap stable, but it is not EF1 because team 3 envies team 1 by more than one player.

In spite of this example, we show that an EF[1,1] and swap stable allocation that is moreover balanced always exists and can be found efficiently.

Theorem 3.2. *For any instance, a balanced allocation that satisfies EF[1,1] and swap stability exists and can be computed in polynomial time.*

Since EF[1,1] reduces to EF1 for nonnegative-value players as well as for nonpositive-value players, Theorem 3.2 implies the following corollary.

Corollary 3.3. *For any nonnegative-value player instance, a balanced allocation that satisfies EF1 and swap stability exists and can be computed in polynomial time. The same holds for any nonpositive-value player instance.*

Our algorithm for Theorem 3.2 proceeds in a round-robin manner. However, instead of assigning a player to a team in each turn as is usually done, we only assign a player's *value* to the team; this ensures that more possibilities are available in later turns. Then, among the allocations that satisfy the determined values for teams, we compute an allocation that minimizes the sum of the players' ranks for the teams. Formally, the algorithm is described as Algorithm 1. For each positive integer q , we denote by $f(q)$ the unique integer in $[n]$ such that $f(q) \equiv q \pmod{n}$. Also, for a matching μ with $(q, p) \in \mu$, we define the notation μ_q and μ_p so that $\mu_q = p$ and $\mu_p = q$. Note that each $q \in Q$ corresponds to a copy of team $f(q)$.

It is clear that the allocation produced by Algorithm 1 is balanced. To establish Theorem 3.2, we prove the remaining properties of the algorithm, including its polynomial running time, in the following three lemmas.

Lemma 3.4. *The output allocation A of Algorithm 1 is EF[1,1].*

Proof. By definition of EF[1,1], we need to show that, for all distinct $i, j \in T$, we have $v_i(A_i \setminus X) \geq v_i(A_j \setminus Y)$ for some $X \subseteq A_i$ and $Y \subseteq A_j$ with $|X|, |Y| \leq 1$. The statement holds trivially if $m \leq n$ since each team receives at most one player, so assume that $m > n$. Fix arbitrary distinct $i, j \in T$.

We consider three cases. (In what follows, μ and μ^* are as defined in Algorithm 1.)

First, suppose that $|A_i| = |A_j|$. Let $k := |A_i| \geq 1$. Then,

$$\begin{aligned} v_i(A_i \setminus \{\mu_{n(k-1)+i}^*\}) &= \sum_{\ell=1}^{k-1} v_i(\mu_{n(\ell-1)+i}^*) = \sum_{\ell=1}^{k-1} v_i(\mu_{n(\ell-1)+i}) \\ &\geq \sum_{\ell=1}^{k-1} v_i(\mu_{n\ell+j}) = \sum_{\ell=2}^k v_i(\mu_{n(\ell-1)+j}) \\ &= \sum_{\ell=2}^k v_i(\mu_{n(\ell-1)+j}^*) = v_i(A_j \setminus \{\mu_j^*\}), \end{aligned}$$

where $v_i(\mu_{n(\ell-1)+i}) \geq v_i(\mu_{n\ell+j})$ holds because otherwise the weight of the edge in μ adjacent to vertex $n(\ell-1) + i \in Q$ can be increased without decreasing the weights of the edges adjacent to vertices $1, 2, \dots, n(\ell-1) + i - 1 \in Q$, contradicting the definition of μ .

The remaining cases $|A_i| > |A_j|$ and $|A_i| < |A_j|$ can be handled similarly; we defer the details to the full version of our paper [Igarashi *et al.*, 2022]. \square

Lemma 3.5. *The output allocation A of Algorithm 1 is swap stable.*

Proof. Let us consider a swap between players μ_q^* and μ_r^* , where $q, r \in Q$ with $q < r$. Suppose that this swap is *possibly* a beneficial swap, i.e., $v_{f(q)}(\mu_q^*) \leq v_{f(q)}(\mu_r^*)$, $v_{f(r)}(\mu_r^*) \leq v_{f(r)}(\mu_q^*)$, $f(q) \succsim_{\mu_q^*} f(r)$, and $f(r) \succsim_{\mu_r^*} f(q)$. We will show that this swap cannot make any of the involved parties better off. Denote by μ^{**} the matching that results from this swap.

If $v_{f(q)}(\mu_q^*) < v_{f(q)}(\mu_r^*)$, the matching μ can be improved by using μ^{**} instead, a contradiction. So $v_{f(q)}(\mu_q^*) = v_{f(q)}(\mu_r^*)$. Similarly, if $v_{f(r)}(\mu_r^*) < v_{f(r)}(\mu_q^*)$, then because $v_{f(q)}(\mu_q^*) = v_{f(q)}(\mu_r^*)$, the matching μ can again be improved by using μ^{**} instead, a contradiction. So $v_{f(r)}(\mu_r^*) = v_{f(r)}(\mu_q^*)$. Hence, the matching μ^{**} after the swap remains a feasible perfect matching in G^* . As μ^* minimizes the sum of the players' rank for teams among the perfect matchings in G^* , we get $f(q) \sim_{\mu_q^*} f(r)$ and $f(r) \sim_{\mu_r^*} f(q)$. Therefore, the swap is not a beneficial swap, and the allocation A is swap stable. \square

Lemma 3.6. *Algorithm 1 can be implemented to run in polynomial time.*

Next, we observe two ways in which Theorem 3.2 cannot be improved: the condition EF[1,1] cannot be strengthened to EF1, and it is not possible to add individual stability to the

Algorithm 2: For computing an EF1, swap stable, and individually stable allocation

- 1 Partition P into $P^+ := \{p \in P \mid \max_{i \in T} v_i(p) \geq 0\}$ and $P^- := \{p \in P \mid \max_{i \in T} v_i(p) < 0\}$;
 - 2 Let \hat{P}^+ consist of P^+ together with $(n-1)|P^+| + n$ dummy players, where each dummy player yields value 0 to every team and is indifferent between all teams;
 - 3 Let \hat{P}^- consist of P^- together with $(n-1)|P^-| + n$ dummy players, where each dummy player yields value 0 to every team and is indifferent between all teams;
 - 4 Let A^+ be the allocation obtained by executing Algorithm 1 on \hat{P}^+ with the teams in the forward order $1, 2, \dots, n$;
 - 5 Let A^- be the allocation obtained by executing Algorithm 1 on \hat{P}^- with the teams in the backward order $n, n-1, \dots, 1$;
 - 6 Return the allocation A which is the union of A^+ and A^- with the dummy players removed;
-

list of guarantees. In fact, the first observation was also made by Shoshan et al. [2022], although their work only deals with one-sided preferences.

Proposition 3.7 ([Shoshan et al., 2022]). *Even for two teams with identical valuations, there does not necessarily exist a balanced EF1 allocation.*

Proposition 3.8. *Even for two teams and nonnegative-value players, there does not necessarily exist a balanced and individually stable allocation.*

Nevertheless, if we give up balancedness, we can attain EF1, swap stability, and individual stability simultaneously. To this end, we combine our Algorithm 1 with the *double round-robin algorithm* introduced by Aziz et al. [2022]. In the first phase, the players who yield nonnegative value to at least one team are allocated by Algorithm 1 in the forward order of the teams, while in the second phase, the remaining players are allocated by Algorithm 1 in the backward order of the teams. Intuitively, EF1 is guaranteed because, for each pair of teams i and j with $i < j$, i does not envy j in the first phase whereas j does not envy i in the second phase. Moreover, we add a sufficient number of dummy players, who yield value 0 to every team and are indifferent between all teams, in order to guarantee individual stability. This leads to each team receiving at least one dummy player, and a beneficial deviation in the resulting situation can be captured by a beneficial swap between the deviating player and a dummy player. The algorithm is formally described as Algorithm 2.

Theorem 3.9. *For any instance, Algorithm 2 returns an EF1, swap stable, and individually stable allocation in polynomial time.*

4 Pareto Optimality

In this section, we turn our attention to Pareto optimality, which is a stronger requirement than both swap stability and individual stability. Firstly, while it is easy to check whether an allocation is swap stable or individually stable by checking for all (polynomial number of) possible beneficial swaps or deviations, the same is not true for PO.

Theorem 4.1. *Deciding whether an allocation is PO or not is coNP-complete, even for two teams with identical valuations, nonnegative-value players, and a balanced allocation.*

Note that even though the same decision problem is also coNP-complete for two teams with one-sided preferences

[Aziz et al., 2019, Thm. 1], it becomes trivial for any number of teams with identical valuations and one-sided preferences, because every allocation is PO in that case.

In light of Theorem 4.1, we cannot hope to reach a PO allocation in polynomial time by starting with an arbitrary allocation and iteratively finding Pareto improvements. However, a PO allocation can be efficiently computed by simply assigning each player to a team with the highest value for her, breaking ties in favor of a team that the player prefers most. Can we attain PO along with fairness for the teams? The next example shows that round-robin-based algorithms such as Algorithms 1 and 2 do not work, even for two teams with identical valuations and nonnegative-value players.

Example 4.2. Consider the following instance with $n = 2$ and $m = 8$:

- For $i \in [2]$, $v_i(p_1) = v_i(p_2) = 4$, $v_i(p_3) = v_i(p_4) = 3$, $v_i(p_5) = v_i(p_6) = 2$, and $v_i(p_7) = v_i(p_8) = 1$;
- $1 \succ_{p_j} 2$ for $j \in \{1, 2, 7, 8\}$ and $2 \succ_{p_j} 1$ for $j \in \{3, 4, 5, 6\}$.

Given this instance, Algorithms 1 and 2 return an allocation A that assigns to each team exactly one player from each of the sets $\{p_1, p_2\}$, $\{p_3, p_4\}$, $\{p_5, p_6\}$, and $\{p_7, p_8\}$. However, A is Pareto dominated by the allocation $A' = (\{p_1, p_2, p_7, p_8\}, \{p_3, p_4, p_5, p_6\})$.

Nevertheless, for two teams and arbitrary-value players, we can find an EF1 and PO allocation by extending the *generalized adjusted winner procedure* of Aziz et al. [2022]. The algorithm operates in a similar way as Aziz et al.'s algorithm but we need to employ a tie-breaking rule among players with the same ratio between the teams' values.

Theorem 4.3. *Given any instance with two teams, there exists an algorithm that outputs an allocation that is EF1, PO, and team-PO in time $O(m^2)$.*

Although EF1, PO, and team-PO can be guaranteed simultaneously in the case of two teams, EF1 and player-PO are already incompatible in this case.

Proposition 4.4. *Even for two teams with identical valuations and nonnegative-value players, there does not necessarily exist an EF1 and player-PO allocation.*

Note also that since PO is a strengthening of individual stability, Proposition 3.8 implies that we cannot guarantee PO and balancedness simultaneously.

We now move on to the general setting where the number of teams can be arbitrary. Unfortunately, even for nonpositive-value players and one-sided preferences, it is unknown whether EF1 and PO can always be satisfied together [Ebadian *et al.*, 2022; Garg *et al.*, 2022]. We therefore restrict our attention to nonnegative-value players in the remainder of this section. By building upon a well-known result of Caragiannis *et al.* [2019], we can establish the existence of an EF1 and PO allocation. For any allocation A , its *Nash welfare* is defined as the product $\prod_{i \in T} v_i(A_i)$. An allocation is said to be a *maximum Nash welfare (MNW) allocation* if it maximizes the Nash welfare among all allocations.⁸

Theorem 4.5. *For any instance with nonnegative-value players, there exists an EF1 and PO allocation.*

Proof. Let \mathcal{W} be the set of all MNW allocations, and let A be an allocation that is PO within \mathcal{W} —such an allocation must exist because otherwise there would be an infinite sequence of Pareto improvements in \mathcal{W} . It is known that every MNW allocation is EF1 [Caragiannis *et al.*, 2019], so A is EF1. We claim that A is PO within the set of all allocations. Suppose to the contrary that there is a Pareto improvement A' of A . Since $v_i(A'_i) \geq v_i(A_i)$ for all $i \in T$, A' must also be an MNW allocation. However, this contradicts the assumption that A is PO within \mathcal{W} . \square

Given Theorem 4.5, a natural question is whether there exists a polynomial-time algorithm that computes an allocation guaranteed by the theorem. However, this question is open even for one-sided preferences.⁹ We demonstrate next that, in two special cases, such an algorithm exists. The first case is when the teams have *binary valuations*, meaning that each team has value either 0 or 1 for each player. In this case, it turns out that Algorithm 2 computes an EF1 and PO allocation in polynomial time.¹⁰

Theorem 4.6. *For any instance with binary valuations, Algorithm 2 computes an EF1 and PO allocation in polynomial time.*

Proof. Since EF1 and polynomial-time computability were already shown in the proof of Theorem 3.9, it is sufficient to establish PO.

Let A be the outcome of Algorithm 2, and suppose to the contrary that there is a Pareto improvement A' of A . For each player p , we denote by $A(p)$ and $A'(p)$ the team that p is allocated to in A and A' , respectively. Note that $A'(p) \succ_p A(p)$ for all $p \in P$ and $v_i(A'_i) \geq v_i(A_i)$ for all $i \in T$. We claim that $v_{A(p)}(p) \geq v_{A'(p)}(p)$ for each player p . Indeed, if this is not the case, then $v_{A(p)}(p) = 0$ and $v_{A'(p)}(p) = 1$ and moreover $A'(p) \succ_p A(p)$. However, a similar proof as that for individual stability in Theorem 3.9 (given in the full version of our paper) shows that such a deviation by p from

⁸If the maximum possible Nash welfare is 0, an MNW allocation should yield nonzero utility to the largest possible number of teams and, subject to that, maximize the product of utilities of these teams.

⁹A pseudopolynomial-time algorithm for this problem was given by Barman *et al.* [2018].

¹⁰With binary valuations, the set \widehat{P}^- in Algorithm 2 is empty, so the algorithm can be simplified.

$A(p)$ to $A'(p)$, which hurts neither p nor $A(p)$ and strictly helps $A'(p)$, cannot exist, thereby proving the claim.

Now, since A' is a Pareto improvement of A , we have

$$\sum_{p \in P} v_{A(p)}(p) = \sum_{i \in T} v_i(A_i) \leq \sum_{i \in T} v_i(A'_i) = \sum_{p \in P} v_{A'(p)}(p).$$

Since $v_{A(p)}(p) \geq v_{A'(p)}(p)$ for all $p \in P$, we must have $v_{A(p)}(p) = v_{A'(p)}(p)$ for all $p \in P$ and $v_i(A_i) = v_i(A'_i)$ for all $i \in T$. Thus, we can construct a better matching than μ^* in Algorithm 1 on \widehat{P}^+ (Line 4 of Algorithm 2) by a round-robin sequence in which each team i picks players in A'_i as early as possible, because the Pareto improvement makes no player worse off and at least one player strictly better off. However, this contradicts the definition of μ^* . \square

Next, we focus on the case of three teams with identical valuations and specific player preferences.

Theorem 4.7. *Suppose that there are $n = 3$ teams with identical valuations, all players yield nonnegative value, and each player prefers one team and is indifferent between the other two teams. Then, there exists an algorithm that computes an EF1 and PO allocation in polynomial time.*

Finally, we provide a pseudopolynomial-time algorithm for the case where the number of teams is constant.

Theorem 4.8. *For any instance with a constant number of teams, each of which has a nonnegative integer value for each player, an EF1 and PO allocation can be computed in pseudopolynomial time.*

5 Justified Envy-Freeness

In this section, we consider justified EF. Note that if $m = n$ and all players yield nonnegative value to every team, the existence of an EF1 and justified EF allocation follows from the celebrated result in two-sided matching of Gale and Shapley [1962] (with arbitrary tie-breaking). We show that, perhaps surprisingly, this guarantee cannot be extended to the case $m > n$.

Proposition 5.1. *Even for two teams and nonnegative-value players who prefer the same team, there may not exist an EF1 and justified EF allocation.*

Proof. Consider an instance with $n = 2$ and $m = 4$ such that team 1 has value 3, 3, 2, 2 for the four players, respectively, team 2 has value 1, 1, 0, 0, respectively, and all players strictly prefer team 1 to team 2. Team 2 needs at least one of p_1 and p_2 in order for EF1 to be fulfilled—assume without loss of generality that it receives p_2 . Given this, team 1 needs at least one of p_3 and p_4 in order for EF1 to be fulfilled—assume without loss of generality that it receives p_3 . However, this results in p_2 having justified envy toward p_3 . \square

Can we efficiently determine whether an EF1 and justified EF allocation exists in a given instance? The following theorem gives a negative answer, provided that $P \neq NP$.

Theorem 5.2. *Even for nonnegative-value players with strict preferences, deciding whether there exists an EF1 and justified EF allocation is NP-complete.*

| | player a | player b | player c_j ($c_j \in C$) | player p_w ($w \in V$) | |
|-------------|--------------------------|-------------------------|------------------------------|-------------------------------|--|
| Preferences | $s \succ d \succ \dots$ | $s \succ d \succ \dots$ | arbitrary | $s \succ \phi(w) \succ \dots$ | |
| Valuations | team s | 1 | 1 | $1 - \varepsilon$ | 1 |
| | team t_e ($e \in E$) | 0 | 0 | 0 | 1 (if w is an endpoint of e) or 0 (otherwise) |
| | team d | 1 | 1 | ε | ε |

Table 2: Player preferences and team valuations in the proof of Theorem 5.2.

Proof. The problem clearly belongs to NP. We prove the hardness by reducing from INDEPENDENT SET, the problem of deciding whether a graph G admits an independent set of size k . Consider an instance (G, k) of INDEPENDENT SET, where $G = (V, E)$. Without loss of generality,¹¹ assume that $k \geq 2$ and $|E| \geq |V| \geq 4$. We construct an instance of our problem as follows.

We set $\tau := 2 + |E|$, which will be the number of teams in our instance, and $\eta := |V| + (k + 1)(\tau - 1) + 3$, which will be the number of players in our instance. Clearly, $\eta > \tau$ since $(k + 1)(\tau - 1) > \tau$ for $\tau \geq 2$. We set $\varepsilon \in (0, 1)$ such that $\varepsilon < \min\{\frac{1}{\eta - 2}, 1 - \frac{k}{k + 1}\}$; hence, we have $1 > (\eta - 2)\varepsilon$ and $(k + 1)(1 - \varepsilon) > k$.

Create one special team s . For each edge $e \in E$, create an edge team t_e . For each vertex $w \in V$, create a vertex player p_w . The special team s assigns value 1 to each vertex player p_w for $w \in V$. Each edge team t_e assigns value 1 to a vertex player p_w if w is an endpoint of edge e and value 0 otherwise. Create an injective map $\phi: V \rightarrow E$; this is possible since $|V| \leq |E|$. Each vertex player p_w prefers s the most and $\phi(w)$ the second most, and has an arbitrary preference over the other teams (including one that will be defined later).

For the special team s , we create the following gadget I that forces s to get k vertex players in our desired allocation. The gadget I consists of two teams s and d , two players a and b , and $(k + 1)(\tau - 1) + 1$ players c_j for $j \in [(k + 1)(\tau - 1) + 1]$. Let $C = \{c_1, c_2, \dots, c_{(k + 1)(\tau - 1) + 1}\}$. Players a and b prefer s the most and d the second most, and have an arbitrary preference over the other teams. Each $c_j \in C$ has an arbitrary (strict) preference over the teams. Teams s and d are the only teams that have positive value for the players in the gadget I . Team s assigns value 1 to each of a and b and value $(1 - \varepsilon)$ to each player $c_j \in C$. Team d assigns value 1 to each of a and b , and value ε to the remaining players (including those outside the gadget). See Table 2 for a summary of the players’ preferences and teams’ valuations in our constructed instance. This completes the description of our construction, which clearly runs in polynomial time.

To finish the proof, we show that there exists an allocation A satisfying EF1 and justified EF if and only if there exists an independent set $S \subseteq V$ of size k in G . The details for establishing this claim are provided in the full version of our paper [Igarashi *et al.*, 2022]. \square

¹¹If $k = 1$ or $|V| \leq 3$, the problem is trivial. If $|V| > |E|$, we repeatedly remove an isolated vertex and decrease k by 1 each time until $|V| \leq 2|E|$, then add an isolated clique of size $|V|$ and increase k by 1.

Despite Theorem 5.2, the problem becomes efficiently solvable if there are two teams. Note that this special case covers the example in the proof of Proposition 5.1.

Theorem 5.3. *For two teams and nonnegative-value players with strict preferences, there is a polynomial-time algorithm that decides whether an EF1 and justified EF allocation exists (and, if so, computes such an allocation).*

Finally, we prove that if the two teams have identical valuations, then an EF1 and justified EF allocation always exists.

Theorem 5.4. *For two teams with identical valuations and nonnegative-value players, there exists an EF1 and justified EF allocation. Moreover, such an allocation can be computed in polynomial time.*

6 Conclusion and Future Work

In this work, we have investigated the setting of fair division with two-sided preferences, which serves to model the allocation of players to sports teams, employees to branches of a restaurant chain, or volunteers to community service clubs. We showed that EF1, swap stability, and individual stability are compatible in this setting, and an allocation satisfying these properties can be computed in polynomial time even when the teams may have positive or negative values for the players. If all players yield nonnegative value to the teams, an EF1 and PO allocation always exists, and such an allocation can be found efficiently provided that the values are binary. Furthermore, we demonstrated that an EF1 and justified EF allocation does not always exist and determining whether such an allocation exists is NP-complete.

For future work, it would be worth examining the interplay between other common (one-sided) fairness notions and two-sided stability conditions such as swap stability and justified EF. For instance, while we have focused on the important fairness notion of EF1, one might try to achieve an approximation of *maximin share fairness (MMS)* [Budish, 2011; Kurokawa *et al.*, 2018], either in place of or in conjunction with EF1. Note that EF1 implies a $1/n$ -approximation of MMS for nonnegative-value players [Amanatidis *et al.*, 2018], so our relevant results also hold for this approximation. It could also be interesting to extend our results to accommodate teams with varying *entitlements* [Farhadi *et al.*, 2019; Chakraborty *et al.*, 2021]; this would allow us to capture restaurant branches or community service clubs of different sizes. Finally, one could attempt to bring other concepts from the rich matching literature into consideration as well.

Acknowledgments

This work was partially supported by JSPS KAKENHI Grant Numbers JP17K12646, JP20K19739, JP21K17708, and JP21H03397, by JST PRESTO Grant Numbers JPMJPR2122 and JPMJPR20C1, by Value Exchange Engineering, a joint research project between Mercari, Inc. and the RIISE, by the Singapore Ministry of Education under grant number MOE-T2EP20221-0001, and by an NUS Start-up Grant. We would like to thank the anonymous reviewers for their valuable comments.

References

- [Aleksandrov and Walsh, 2020] Martin Aleksandrov and Toby Walsh. Two algorithms for additive and fair division of mixed manna. In *Proceedings of the 43rd German Conference on Artificial Intelligence (KI)*, pages 3–17, 2020.
- [Amanatidis *et al.*, 2018] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. Comparing approximate relaxations of envy-freeness. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 42–48, 2018.
- [Aziz and Savani, 2016] Haris Aziz and Rahul Savani. Hedonic games. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 15, pages 356–376. Cambridge University Press, 2016.
- [Aziz *et al.*, 2019] Haris Aziz, Péter Biró, Jérôme Lang, Julien Lesca, and Jérôme Monnot. Efficient reallocation under additive and responsive preferences. *Theoretical Computer Science*, 790:1–15, 2019.
- [Aziz *et al.*, 2022] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi-Agent Systems*, 36(1):3:1–3:21, 2022.
- [Barman *et al.*, 2018] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 557–574, 2018.
- [Bérczi *et al.*, 2020] Kristóf Bérczi, Erika R. Bérczi-Kovács, Endre Boros, Fekadu Tolessa Gedefa, Naoyuki Kamiyama, Telikepalli Kavitha, Yusuke Kobayashi, and Kazuhisa Makino. Envy-free relaxations for goods, chores, and mixed items. *CoRR*, abs/2006.04428, 2020.
- [Brams and Taylor, 1996] Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- [Budish, 2011] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2019] 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.
- [Chakraborty *et al.*, 2021] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. Weighted envy-freeness in indivisible item allocation. *ACM Transactions on Economics and Computation*, 9(3):18:1–18:39, 2021.
- [Ebadian *et al.*, 2022] Soroush Ebadian, Dominik Peters, and Nisarg Shah. How to fairly allocate easy and difficult chores. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 372–380, 2022.
- [Farhadi *et al.*, 2019] Alireza Farhadi, Mohammad Ghodsi, MohammadTaghi Hajiaghayi, Sebastien Lahaie, David Pennock, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research*, 64:1–20, 2019.
- [Freeman *et al.*, 2021] Rupert Freeman, Evi Micha, and Nisarg Shah. Two-sided matching meets fair division. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 203–209, 2021.
- [Gale and Shapley, 1962] David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.
- [Garg *et al.*, 2022] Jugal Garg, Aniket Murhekar, and John Qin. Fair and efficient allocations of chores under bivalued preferences. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5043–5050, 2022.
- [Gollapudi *et al.*, 2020] Sreenivas Gollapudi, Kostas Kollias, and Benjamin Plaut. Almost envy-free repeated matching in two-sided markets. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*, pages 3–16, 2020.
- [Igarashi *et al.*, 2022] Ayumi Igarashi, Yasushi Kawase, Warut Suksompong, and Hanna Sumita. Fair division with two-sided preferences. *CoRR*, abs/2206.05879, 2022.
- [Kulkarni *et al.*, 2021] Rucha Kulkarni, Ruta Mehta, and Setareh Taki. Indivisible mixed manna: on the computability of MMS+PO allocations. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, pages 683–684, 2021.
- [Kurokawa *et al.*, 2018] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM*, 64(2):8:1–8:27, 2018.
- [Lipton *et al.*, 2004] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Economics and Computation (EC)*, pages 125–131, 2004.
- [Moulin, 2003] Hervé Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.
- [Patro *et al.*, 2020] Gourab K. Patro, Arpita Biswas, Niloy Ganguly, Krishna P. Gummadi, and Abhijnan Chakraborty. FairRec: Two-sided fairness for personalized recommendations in two-sided platforms. In

Proceedings of the 29th Web Conference (WWW), pages 1194–1204, 2020.

[Shoshan *et al.*, 2022] Hila Shoshan, Erel Segal-Halevi, and Noam Hazon. Efficient nearly-fair division with capacity constraints. *CoRR*, abs/2205.07779, 2022.

[Wu and Roth, 2018] Qingyun Wu and Alvin E. Roth. The lattice of envy-free matchings. *Games and Economic Behavior*, 109:201–211, 2018.

[Yokoi, 2020] Yu Yokoi. Envy-free matchings with lower quotas. *Algorithmica*, 82(2):188–211, 2020.