

Position Fair Mechanisms Allocating Indivisible Goods

Ryoga Mahara¹, Ryuhei Mizutani², Taihei Oki^{3,4}, Tomohiko Yokoyama¹

¹The University of Tokyo, Japan

²Keio Univeristy, Japan

³Hokkaido Univeristy, Japan

⁴RIKEN, Japan

{mahara, tomohiko_yokoyama}@mist.u-tokyo.ne.jp, mizutani@math.keio.ac.jp, oki@icredd.hokudai.ac.jp

Abstract

Fair division mechanisms for indivisible goods require agent orderings to deterministically select one allocation when running the algorithm in practice. We introduce position envy-freeness up to one good (PEF1) as a fairness criterion for mechanisms: a mechanism is said to satisfy PEF1 if for any pair of agent orderings, no agent prefers their bundle determined under one ordering to that under another ordering by more than the utility of a single good. First, we propose a scale-invariant, polynomial-time mechanism that satisfies PEF1 and yields an envy-freeness up to one good (EF1) allocation. For the case of two agents, we establish that any mechanism producing a maximum Nash welfare allocation eliminates envy based on positions by removing one good, provided that utilities are positive. Additionally, we present a polynomial-time mechanism based on the adjusted winner procedure, which satisfies PEF1 and produces an EF1 and Pareto optimal allocation for two agents. In contrast, we demonstrate that well-known mechanisms such as round-robin and envy-cycle elimination do not generally satisfy PEF1.

1 Introduction

Fair division of indivisible goods among agents is a fundamental problem in economics and computer science with significant practical importance. Applications range from course allocation in universities (Budish et al. 2016) and inheritance division (Goldman and Procaccia 2015) to various other real-world settings (Igarashi and Yokoyama 2023; Han and Suksompong 2024). For surveys, see (Walsh 2020; Aziz et al. 2022b; Amanatidis et al. 2023). In fair division, agents are typically assumed to have *additive* utilities over bundles of goods, and the challenge is to find an allocation that meets certain fairness criteria. One well-studied criterion is *envy-freeness* (Foley 1967), which requires an allocation to satisfy that no agent prefers another agent’s bundle to their own.

In this paper, we consider fairness properties not of allocations but of *mechanisms*, which have received little attention in existing work. Particularly, we deal with fairness regarding an input order of agents. To illustrate this, consider a situation where two agents have identical utilities for one good, requiring mechanisms to establish clear rules for

determining which agent receives it. Namely, mechanisms must incorporate an ordering among the agents.

We formalize *mechanisms* as follows (see Section 2.2 and Figure 1 for formal definitions). The input of a mechanism is a tuple of utilities arranged by an *agent ordering*. The mechanism prioritizes agents based on their *positions* in the agent ordering, and outputs an allocation. Under this framework, utilities that agents obtain by the mechanism can vary substantially depending on their positions in the agent ordering.

Previous work by Manabe and Okamoto (2012) introduced a fairness concept for mechanisms in the *divisible* goods setting, based on agent orderings. This concept, called *meta-envy-freeness*, requires a mechanism to ensure that each agent obtains the same utility regardless of the agent ordering. A meta-envy-free mechanism always exists for divisible goods since divisible goods can be split equally among agents. Specifically, when n agents desire a single good, each can simply receive $1/n$ of the good, obtaining identical utilities.

In contrast, for indivisible goods, the situation differs significantly. A meta-envy-free mechanism cannot exist even with two agents and a single good, as only the agent in the higher priority position can receive the good.

A similar difficulty arises for envy-freeness in the indivisible goods setting, where envy-free allocations do not always exist. This motivated the development of relaxations such as *envy-freeness up to one good (EF1)* (Budish 2011). An allocation is said to be EF1 if any agent’s envy toward another agent can be eliminated by removing a single good from the envied agent’s bundle. Importantly, an EF1 allocation always exists for agents with additive utilities (Lipton et al. 2004; Caragiannis et al. 2019).

Inspired by EF1 and the concept of meta-envy-freeness, we introduce *position envy-freeness up to one good (PEF1)*¹ as a fairness criterion for mechanisms with respect to agent orderings. A mechanism is said to satisfy PEF1 if for any agent and two agent orderings, their envy towards a bundle under one ordering over that under another ordering can be eliminated by removing a single good from the envied bundle.

Since PEF1 is a property of mechanisms concerning dif-

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹We use the term “position envy-free” instead of “meta-envy-free” to clarify the meaning of “meta.”

ferent agent orderings, a PEF1 mechanism may not produce a fair allocation. This raises a fundamental question: can we design a mechanism that satisfies PEF1 and is guaranteed to produce an EF1 allocation?

Our Results First, we answer the above question affirmatively by presenting a PEF1 mechanism that always produces an EF1 allocation (Theorem 3). The mechanism employs a maximum-weight matching to determine a bundle for agents in each round. Notably, the mechanism runs in polynomial time.

Second, for two agents, we establish that any mechanism that maximizes the Nash welfare (i.e., the geometric mean of agents’ utilities) satisfies PEF1. Since a maximum Nash welfare (MNW) allocation is both EF1 and Pareto optimal (PO, i.e., no allocation makes some agent better off without making another agent worse off) (Caragiannis et al. 2019), this mechanism outputs an allocation that satisfies both EF1 and PO. While computing a MNW allocation is NP-hard (Nguyen et al. 2014) and even APX-hard (Lee 2017), we present a modified version of the adjusted winner mechanism (Brams and Taylor 1996; Aziz et al. 2015, 2022a) that is PEF1, runs in polynomial time, and produces an EF1 and PO allocation for the case of two agents.

Finally, we analyze the round-robin mechanism, which produces an EF1 allocation (Caragiannis et al. 2019). We show that the mechanism satisfies PEF1 for the case of at most three agents and it lacks PEF1 when the number of agents is at least four. In addition, we show that the envy-cycle mechanism, which returns an EF1 allocation when agents have monotone utility functions (Lipton et al. 2004), also fails to satisfy PEF1 in general.

We remark that all mechanisms proposed in this paper are scale-invariant; namely, scaling an agent’s utility by any positive constant does not affect the output.

Further Related Work A mechanism is said to be *anonymous* if for any agent ordering, the outcome *allocation* of the mechanism remains unchanged (Gibbard 1973; Satterthwaite 1975). In this paper, we distinguish between anonymity and position (meta-)envy-freeness. Position envy-freeness ensures that the *utilities* that agents receive do not change, while the allocation produced by the mechanism may vary depending on the agent ordering.

The *equal-treatment-of-equals* (ETE) is also a fairness concept for mechanisms (Moulin 2004). A mechanism satisfies ETE if agents with identical preferences receive the same bundle of goods.

Similar to position envy-freeness, both anonymity and ETE are impossible to achieve for indivisible goods. For divisible goods, the compatibility of anonymity and ETE with fairness and efficiency properties has been extensively studied (Shapley and Scarf 1974; Zhou 1990; Bogomolnaia and Moulin 2001; Roth, Sönmez, and Ünver 2005; Bei, Huzhang, and Suksompong 2020).

Our work also relates to a fairness concept in allocation rules. An *allocation rule* is a map from utility profiles to sets of allocations satisfying specified criteria (Sönmez 1999). An allocation rule is *essentially-single-valued* if for any utility profile and any two allocations in its output set, each

agent receives equal utility from these allocations (Sönmez 1999). If a mechanism always selects its output from allocations determined by an essentially-single-valued allocation rule, then the mechanism is meta-envy-free. It is known that the MNW allocation rule, defined as a mapping to all MNW allocations, is essentially-single-valued for agents with continuous utility functions over divisible goods (Dubins and Spanier 1961; Segal-Halevi and Sziklai 2019).

2 Preliminaries

2.1 Fair Division Model

Let M be the set of m goods and N be the set of n agents. A subset of M is termed a *bundle*. Each agent $a \in N$ has a non-negative utility function $u_a : 2^M \rightarrow \mathbb{R}_{\geq 0}$. For simplicity, we denote $u_a(g)$ as $u_a(\{g\})$ for each $g \in M$. We assume that u_a is *additive*, that is, we have $u_a(S) = \sum_{g \in S} u_a(g)$ for any $S \subseteq M$. A family $u = \{u_a\}_{a \in N}$ of the utility functions of all agents is called a *profile*. Let $U_{\geq 0}$ denote the set of all profiles. An *allocation* $A = \{A_a\}_{a \in N}$ is a partition of M into n bundles, where A_a denotes the bundle of agent $a \in N$.

We now introduce a fairness concept of an allocation. An allocation A is said to be *envy-free* if no agent envies any other agent, i.e., $u_a(A_a) \geq u_a(A_{a'})$ for all $a, a' \in N$. An allocation A is called *envy-free up to one good* (EF1) if for all $a, a' \in N$ with $A_{a'} \neq \emptyset$, there exists a good $g \in A_{a'}$ such that $u_a(A_a) \geq u_a(A_{a'} \setminus \{g\})$.

Next, we define an efficiency concept. An allocation A is said to *Pareto dominate* another allocation A' if $u_a(A_a) \geq u_a(A'_a)$ for all $a \in N$ and $u_{a'}(A_{a'}) > u_{a'}(A'_a)$ for some $a' \in N$. An allocation A is called *Pareto optimal* (PO) if there is no allocation that Pareto dominates A .

2.2 Mechanism and Position Fairness

We now introduce several notions related to agent orderings. An *agent ordering* is defined as a bijection from N to $[n] = \{1, 2, \dots, n\}$. We call $\pi(a)$ the position of agent a under π . Let Π denote the set of all agent orderings. An *ordered profile* $u = (u_1, u_2, \dots, u_n)$ is an ordered n -tuple of utility functions. Given an agent ordering π , let u_π be an ordered profile generated from an original profile $u = \{u_a\}_{a \in N}$ by mapping utility functions according to π , i.e., $(u_\pi)_i = u_{\pi^{-1}(i)}$ for each $i \in [n]$. An *ordered allocation* $A = (A_1, A_2, \dots, A_n)$ is a partition of M into n bundles indexed from 1 to n .

A *mechanism* \mathcal{M} is defined as a map from ordered profiles to ordered allocations. More precisely, an input of a mechanism is an ordered profile u_π generated from a profile u and an agent ordering π , and the mechanism cannot access u and π ; in other words, the mechanism is ignorant of the correspondence between agents and their positions in π . Then, $\mathcal{M}(u_\pi)$ means the ordered allocation returned by \mathcal{M} when the input is u_π . See Figure 1 for an illustration.

We now define a fairness concept of a mechanism concerning agent orderings, called *position envy-freeness*. This concept means that no agent envies the bundle under one agent ordering π compared to that under another agent ordering π' .

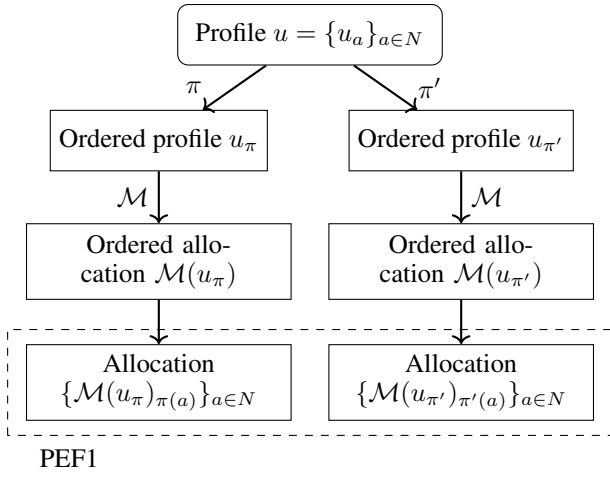


Figure 1: An illustration of a mechanism \mathcal{M} and fairness concepts.

Definition 1. A mechanism \mathcal{M} satisfies *position envy-freeness* with respect to $U_{\geq 0}$ if for any profile $u \in U_{\geq 0}$, for any agent orderings $\pi, \pi' \in \Pi$, and for any agent $a \in N$,

$$u_a(\mathcal{M}(u_\pi)_{\pi(a)}) \geq u_a(\mathcal{M}(u_{\pi'})_{\pi'(a)}).$$

As mentioned in Section 1, this concept is already known as *meta-envy-freeness* (Manabe and Okamoto 2012) in fair division for divisible goods.

Unfortunately, a position envy-free mechanism may not exist for indivisible goods. Consider a setting with two agents and a single good, where both agents have identical utility functions. The input of a mechanism is an ordered tuple of the same two utility functions. Without knowing the correspondence between agents and their positions, the mechanism must allocate the good consistently to a fixed position (e.g., position 1). This leads to a violation of position envy-freeness since they receive the good when assigned to position 1 but not when assigned to position 2.

Following the spirit of EF1, we introduce a relaxation of the position envy-freeness, called *position envy-freeness up to one good* (PEF1). This relaxation allows for some degree of position-based disparity, bounded by the utility of at most one good. See also Figure 1.

Definition 2. A mechanism \mathcal{M} satisfies *position envy-freeness up to one good* (PEF1) with respect to $U_{\geq 0}$ if for any profile $u \in U_{\geq 0}$, agent orderings $\pi, \pi' \in \Pi$, and agent $a \in N$ with $\mathcal{M}(u_{\pi'})_{\pi'(a)} \neq \emptyset$, there exists a good $g \in \mathcal{M}(u_{\pi'})_{\pi'(a)}$ such that

$$u_a(\mathcal{M}(u_\pi)_{\pi(a)}) \geq u_a(\mathcal{M}(u_{\pi'})_{\pi'(a)} \setminus \{g\}).$$

When we write PEF1 without further specification, we mean PEF1 with respect to $U_{\geq 0}$.

Additionally, we define the scale-invariance of mechanisms. For a profile u and a tuple of positive real numbers $\alpha = (\alpha_a)_{a \in N} \in \mathbb{R}_{>0}^n$ indexed by agents in N , let αu denote the profile defined as $(\alpha u)_a(S) = \alpha_a \cdot u_a(S)$ for every $a \in N$ and $S \subseteq M$. A mechanism \mathcal{M} is called *scale-invariant* if $\mathcal{M}(u_\pi) = \mathcal{M}((\alpha u)_\pi)$ for every profile u , tuple $\alpha \in \mathbb{R}_{>0}^n$, and agent ordering $\pi \in \Pi$.

Algorithm 1: Round-Robin Mechanism

Input: Ordered profile (u_1, u_2, \dots, u_n)
Output: Ordered allocation (A_1, A_2, \dots, A_n)

- 1: Fix indices of goods.
- 2: $A_i \leftarrow \emptyset$ for all $i \in [n]$
- 3: $i \leftarrow 1$
- 4: **while** $A_1 \cup A_2 \cup \dots \cup A_n \subsetneq M$ **do**
- 5: Take $g \in \operatorname{argmax}_{g' \in M \setminus (A_1 \cup A_2 \cup \dots \cup A_n)} u_i(g')$
- 6: $A_i \leftarrow A_i \cup \{g\}$
- 7: $i \leftarrow (i \bmod n) + 1$
- 8: **return** (A_1, A_2, \dots, A_n)

To explain these concepts, we present the round-robin mechanism (Caragiannis et al. 2019) described in Algorithm 1. The mechanism first gives a total order of the goods for tie breaking. Then it operates by having agents sequentially select their most preferred remaining good, following the order specified by $\pi^{-1}(1)$ to $\pi^{-1}(n)$. Ties are broken by choosing the good with the smallest order. This process continues until all goods are allocated. The round-robin mechanism possesses two important properties: it is scale-invariant and outputs EF1 allocation for any profile and agent ordering (Caragiannis et al. 2019).

While the round-robin mechanism gives an EF1 allocation, it does not satisfy PEF1. To illustrate this, consider an instance with four agents, five goods, and a profile u as shown in Table 1 with positive values $x > y > z > 0$. Consider two agent orderings $\pi, \pi' \in \Pi$ such that $\pi(a_i) = i$ and $\pi'(a_i) = 5 - i$ for each $i \in \{1, 2, 3, 4\}$. Let \mathcal{M} be the round-robin mechanism and let $A = \mathcal{M}(u_\pi)$ and $B = \mathcal{M}(u_{\pi'})$. Under π , agent a_1 receives $A_{\pi(a_1)} = \{g_1, g_5\}$, while under π' , agent a_1 receives $B_{\pi'(a_1)} = \{g_4\}$. For any $g \in A_{\pi(a_1)}$, we have $u_{a_1}(A_{\pi(a_1)} \setminus \{g\}) > z = u_{a_1}(B_{\pi'(a_1)})$, which violates PEF1.

	g_1	g_2	g_3	g_4	g_5	
Agent a_1	x	0	0	z	y	
Agent a_2	0	x	0	0	y	
Agent a_3	x	0	y	0	0	
Agent a_4	0	x	z	y	0	

Table 1: An example utility profile where the round-robin mechanism violates PEF1.

More generally, for the round-robin mechanism \mathcal{M} , if $\lceil \frac{m}{n} \rceil \geq \lfloor \log_2 n \rfloor$, there exist a profile u and agent orderings π, π' such that even after removing any $\lfloor \log_2 n \rfloor - 1$ goods from $\mathcal{M}(u_\pi)$, the agent still prefers $\mathcal{M}(u_{\pi'})$ (Theorem 10). Further discussion of the round-robin mechanism can be found in Section 5.

3 Existence of a Scale-Invariant PEF1 Mechanism Producing an EF1 Allocation

In this section, we present our main result.

Algorithm 2: A Scale-Invariant and PEF1 Mechanism Producing an EF1 Allocation

Input: Ordered profile (u_1, u_2, \dots, u_n)
Output: Ordered allocation (A_1, A_2, \dots, A_n)

- 1: Fix indices of goods as $M = \{g_1, g_2, \dots, g_m\}$.
- 2: Add dummy goods until m is divisible by n .
- 3: $A_i \leftarrow \emptyset$ for all $i \in [n]$ and $I \leftarrow M$
- 4: **for** $r = 1$ to $\lceil \frac{m}{n} \rceil$ **do**
- 5: Compute a maximum-weight matching μ_r with respect to w defined by the equation (1) in $G_r = ([n] \cup I, E_r)$.
- 6: Let $\mu_r(i)$ denote the good in I matched with $i \in [n]$ under μ_r .
- 7: **for** $i = 1$ to n **do**
- 8: $A_i \leftarrow A_i \cup \{\mu_r(i)\}$
- 9: $I \leftarrow I \setminus \{\mu_r(i)\}$
- 10: **return** (A_1, A_2, \dots, A_n)

Theorem 3. *There exists a scale-invariant, PEF1 mechanism that always produces an EF1 allocation in polynomial time.*

We will prove Theorem 3 by presenting Algorithm 2, which constructs a maximum-weight matching iteratively. This mechanism is similar to one proposed by Brustle et al. (2020) in the context of fair division with subsidy. Our mechanism is distinguished by its scale invariance and the tie-breaking method.

In the mechanism, we first give an arbitrary total order of the goods M . Let g_1, \dots, g_m be the goods aligned according to this order. We then ensure that m is divisible by n by adding dummy goods valued at zero by all agents if necessary. Initially, set $A_i = \emptyset$ for every position $i \in [n]$, and let I be a set of all unallocated goods.

The core of the mechanism consists of $\frac{m}{n}$ rounds (the for loop of Lines 4-9). Let $G = ([n] \cup M, E)$ denote a complete bipartite graph with two disjoint vertex sets $[n]$ and M , where $E = \{\{i, g\} \mid i \in [n], g \in M\}$. In each round r , we consider the remaining subgraph $G_r = ([n] \cup I, E_r)$, where $E_r = \{\{i, g\} \mid i \in [n], g \in I\}$.

The mechanism utilizes a weight function $w : [n] \times M \rightarrow \mathbb{R}_{\geq 0}$ on E defined by

$$w(i, g) = 2^{m+1}nw_1(i, g) + w_2(i, g) \quad (1)$$

for each edge $\{i, g\} \in E$. Here, two weight functions w_1 and w_2 are defined as follows.

The first weight function w_1 is based on utilities. For each good g and the agent in each position i , we define *rank* $R(i, g)$ as follows: $R(i, g) = k$ if g has the k th highest utility among all goods in M for the agent in position i . If multiple goods have the same utility, they are assigned the same rank, and the next rank is assigned as if no ties occurred. For example, if four goods g_1, g_2, g_3, g_4 have utilities 10, 10, 8, and 7, respectively, then $R(i, g_1) = R(i, g_2) = 1$, $R(i, g_3) = 2$, and $R(i, g_4) = 3$. For each $i \in [n]$ and $g \in M$, we define

$$w_1(i, g) = m - R(i, g).$$

The second weight function w_2 is based on the total order of the goods: for each $i \in [n]$ and $\ell \in [m]$, we define

$$w_2(i, g_\ell) = 2^{m-\ell}.$$

Additionally, recall the fundamental concept from matching theory. A subset of edges is a *matching* if no two edges in the subset share a common vertex. For a matching μ , let $w(\mu)$ denote the total weight of the matching with respect to w . Let $\mu(i)$ denote the good matched with i under μ .

In each r th round, the mechanism computes a maximum-weight matching μ_r with respect to w in G_r (Line 5). While there may be multiple maximum-weight matchings, any such matching can be selected. Then, according to μ , goods are allocated to each position. As we will show in Lemma 4, maximizing w ensures that we first maximize w_1 , then w_2 among matchings maximizing w_1 .

To illustrate the behavior of Algorithm 2, we apply it to the same instance from Table 1. Recall the profile with four agents and five goods with utilities satisfying $x > y > z > 0$. For this profile, the edge weights of the bipartite graph are given as in Table 2.

	g_1	g_2	g_3	g_4	g_5
Agent a_1	1040	264	260	514	769
Agent a_2	528	1032	516	514	769
Agent a_3	1040	520	772	514	513
Agent a_4	272	1032	516	770	257

Table 2: Edge weights for the bipartite graph in Algorithm 2.

For any agent ordering π , in the bipartite graph, there exists a unique maximum-weight matching $\{(\pi(a_1), g_1), (\pi(a_2), g_2), (\pi(a_3), g_3), (\pi(a_4), g_4)\}$. This means that regardless of which agent ordering is chosen, each agent receives the same good in round 1. In the second round, only good g_5 remains unallocated. Here, the maximum-weight matchings between the remaining good and agents are $\{(\pi(a_1), g_5)\}$ and $\{(\pi(a_2), g_5)\}$. The algorithm selects either maximum-weight matching. For instance, we can choose to prioritize the agent with the smaller position value under π . Although agent a_1 experiences position envy between different orderings, this envy is bounded by at most one good, satisfying PEF1.

3.1 Proof of Theorem 3

The following lemma shows that a single weight function can achieve lexicographical maximization.

Lemma 4. *A maximum-weight matching of G_r with respect to w maximizes w_1 first, and among all such matchings, maximizes w_2 .*

Proof. Let μ and ν be any two matchings, and let $w_i(\mu)$ and $w_i(\nu)$ denote the total weights of μ and ν with respect to w_i for $i \in \{1, 2\}$. Since each matching contains at most n edges, and $w_2(e) \leq 2^m$ for all $e \in E$, we have $w_2(\mu) - w_2(\nu) \geq -2^m n$. If $w_1(\mu) > w_1(\nu)$, then we get $w_1(\mu) - w_1(\nu) \geq 1$ and

$$w(\mu) - w(\nu)$$

$$\begin{aligned}
&= 2^{m+1}n(w_1(\mu) - w_1(\nu)) + (w_2(\mu) - w_2(\nu)) \\
&\geq 2^{m+1}n - 2^m n > 0.
\end{aligned}$$

This implies that a maximum-weight matching with respect to w must maximize w_1 . If $w_1(\mu) = w_1(\nu)$ and $w_2(\mu) > w_2(\nu)$, then $w(\mu) > w(\nu)$. Thus, among matchings maximizing w_1 , a maximum-weight matching with respect to w must maximize w_2 . \square

For each agent ordering $\pi \in \Pi$ and each $r \in \{1, 2, \dots, \lceil \frac{m}{n} \rceil\}$, let $\mu_{\pi,r}$ denote a maximum-weight matching in G_r with respect to w computed in r th round when the input is u_π . Let $\Gamma_{\pi,r}$ be the set of goods matched by $\mu_{\pi,r}$. For agent $a \in N$, we denote by $g_{\pi,r,a}$ the good matched with position $\pi(a)$ under $\mu_{\pi,r}$. By the definition of w_2 , we can show that the set $\Gamma_{\pi,r}$ is determined independently of the agent ordering π .

Lemma 5. *For every round $r \in \{1, 2, \dots, \lceil \frac{m}{n} \rceil\}$ and any pair of agent orderings $\pi, \pi' \in \Pi$, we have $\Gamma_{\pi,r} = \Gamma_{\pi',r}$.*

Proof. We prove the lemma by induction on r . Fix any ordering π and, for $r = 1$, suppose there exist two maximum-weight matchings $\mu_{\pi,1}$ and $\mu'_{\pi,1}$ with respect to w . By the definition of w , these matchings must have the same w_2 weight sum. Since w_2 uses powers of 2 based on the goods' indices, this equality implies that μ and μ' match the same set of goods. Thus, $\Gamma_{\pi,1}$ is unique. Moreover, since w_2 does not depend on positions, $\Gamma_{\pi,1}$ is independent of π . The same argument applies inductively for each subsequent round $r > 1$, completing the proof. \square

Finally, we prove Theorem 3.

Proof of Theorem 3. Let \mathcal{M} denote Algorithm 2. We first prove that \mathcal{M} satisfies PEF1. To this end, we compare any pair of two agent orderings $\pi, \pi' \in \Pi$. By Lemma 5, $\Gamma_{\pi,r} = \Gamma_{\pi',r}$ for all $r = 1, 2, \dots, \lceil \frac{m}{n} \rceil$. This implies that $g_{\pi,r+1,a} \notin \bigcup_{r'=1}^r \Gamma_{\pi,r'}$ for any $r = 1, 2, \dots, \lceil \frac{m}{n} \rceil - 1$ and any agent $a \in N$. By Lemma 4, matching $\mu_{\pi',r}$ is a maximum-weight matching with respect to w_1 . Thus, for any agent $a \in N$, we obtain $w_1(\pi(a), g_{\pi,r+1,a}) \leq w_1(\pi'(a), g_{\pi',r,a})$ since otherwise good $g_{\pi,r+1,a}$ is included in the maximum-weight matching $\mu_{\pi',r}$. By the definition of w_1 , this implies $u_a(g_{\pi,r+1,a}) \leq u_a(g_{\pi',r,a})$. This leads that

$$\begin{aligned}
u_a(\mathcal{M}(u_{\pi'})_{\pi'(a)}) &= \sum_{r=1}^{\lceil \frac{m}{n} \rceil} u_a(g_{\pi',r,a}) \geq \sum_{r=1}^{\lceil \frac{m}{n} \rceil - 1} u_a(g_{\pi',r,a}) \\
&\geq \sum_{r=2}^{\lceil \frac{m}{n} \rceil} u_a(g_{\pi,r,a}) = u_a(\mathcal{M}(u_\pi)_{\pi(a)} \setminus \{g_{\pi,1,a}\}),
\end{aligned}$$

which implies that the mechanism is PEF1.

Next, we show that the mechanism always produces an EF1 allocation. Fix any agent ordering π . Since we choose a maximum-weight matching with respect to w_1 , we have $u_a(g_{\pi,r+1,a'}) \leq u_a(g_{\pi,r,a})$ for any two agents $a, a' \in N$ and $r = 1, 2, \dots, \lceil \frac{m}{n} \rceil - 1$. Then, for all $a, a' \in N$, we have

$$u_a(\mathcal{M}(u_\pi)_{\pi(a)}) = \sum_{r=1}^{\lceil \frac{m}{n} \rceil} u_a(g_{\pi,r,a}) \geq \sum_{r=1}^{\lceil \frac{m}{n} \rceil - 1} u_a(g_{\pi,r,a})$$

$$\geq \sum_{r=2}^{\lceil \frac{m}{n} \rceil} u_a(g_{\pi,r,a'}) = u_a(\mathcal{M}(u_\pi)_{\pi(a')} \setminus \{g_{\pi,1,a'}\}).$$

We finally consider the time complexity and scale-invariance of the mechanism. In each round, we can find a maximum-weight matching with respect to w in polynomial time (Lovász and Plummer 2009). Thus, the mechanism runs in polynomial time. Furthermore, the weight function w is unchanged if the profile is multiplied by a tuple of positive reals. Therefore, the mechanism is scale-invariant. \square

4 The Case of Two Agents

In this section, we focus on the case of two agents.

4.1 Maximize Nash Welfare

The *Nash welfare* of an allocation $A = \{A_a\}_{a \in N}$ is defined as $\text{NW}(A) = (\prod_{a \in N} u_a(A_a))^{1/n}$. An allocation A is said to be *maximum Nash welfare (MNW)* if it maximizes $\text{NW}(A)$ among all allocations. Let $U_{>0}$ be the class of profiles where $u_a : 2^M \rightarrow \mathbb{R}_{>0}$ for all agents $a \in N$.

We will show that PEF1 can be achieved by a mechanism that maximizes the Nash welfare for two agents. To this end, we first prove the following theorem.

Theorem 6. *When $n = 2$, for any $u \in U_{>0}$, any two MNW allocations A and B , and any agent $a \in N$, if $B_a \neq \emptyset$, then there exists $g \in B_a$ such that $u_a(A_a) \geq u_a(B_a \setminus \{g\})$.*

Proof. Let $N = \{a_1, a_2\}$ denote the set of two agents. Let $A = \{A_{a_1}, A_{a_2}\}$ and $B = \{B_{a_1}, B_{a_2}\}$ be two distinct MNW allocations.

We show that for agent a_1 , if $B_{a_1} \neq \emptyset$, there exists some good $g \in B_{a_1}$ such that $u_{a_1}(A_{a_1}) \geq u_{a_1}(B_{a_1} \setminus \{g\})$. The same argument can be applied to a_2 . Without loss of generality, we can assume that $B_{a_1} \setminus A_{a_1} \neq \emptyset$. Suppose, towards a contradiction, that $u_{a_1}(A_{a_1}) < u_{a_1}(B_{a_1} \setminus \{g\})$ for every $g \in B_{a_1} \setminus A_{a_1}$. Take a good $h \in B_{a_1} \setminus A_{a_1}$ (note that $h \in A_{a_2}$). Consider allocations A' and B' where $A'_{a_1} = A_{a_1} \cup \{h\}$, $A'_{a_2} = A_{a_2} \setminus \{h\}$, $B'_{a_1} = B_{a_1} \setminus \{h\}$, and $B'_{a_2} = B_{a_2} \cup \{h\}$.

We will show that $\text{NW}(A')\text{NW}(B') > \text{NW}(A)\text{NW}(B)$, contradicting the optimality of A and B . Observe that

$$\begin{aligned}
&\frac{\text{NW}(A')^2}{\text{NW}(A)^2} \cdot \frac{\text{NW}(B')^2}{\text{NW}(B)^2} \\
&= \frac{u_{a_1}(A'_{a_1})}{u_{a_1}(A_{a_1})} \cdot \frac{u_{a_2}(A'_{a_2})}{u_{a_2}(A_{a_2})} \cdot \frac{u_{a_1}(B'_{a_1})}{u_{a_1}(B_{a_1})} \cdot \frac{u_{a_2}(B'_{a_2})}{u_{a_2}(B_{a_2})} \\
&= \left(1 + \frac{u_{a_1}(h)(u_{a_1}(B_{a_1}) - u_{a_1}(A_{a_1}) - u_{a_1}(h))}{u_{a_1}(A_{a_1})u_{a_1}(B_{a_1})} \right) \\
&\quad \cdot \left(1 + \frac{u_{a_2}(h)(u_{a_2}(A_{a_2}) - u_{a_2}(B_{a_2}) - u_{a_2}(h))}{u_{a_2}(A_{a_2})u_{a_2}(B_{a_2})} \right).
\end{aligned}$$

By our assumption, $u_{a_1}(A_{a_1}) < u_{a_1}(B_{a_1} \setminus \{h\})$. This implies that the first term in the product on the right-hand side is strictly greater than 1.

Since A is MNW, $\text{NW}(A) \geq \text{NW}(B')$, implying $u_{a_1}(A_{a_1})u_{a_2}(A_{a_2}) \geq u_{a_1}(B_{a_1} \setminus \{h\})u_{a_2}(B_{a_2} \cup \{h\})$. Combined with $u_{a_1}(A_{a_1}) < u_{a_1}(B_{a_1} \setminus \{h\})$, we obtain $u_{a_2}(A_{a_2}) \geq u_{a_2}(B_{a_2} \cup \{h\})$. Thus, the second

Algorithm 3: Adjusted Winner Mechanism

Input: Ordered profile (u_1, u_2)
Output: Ordered allocation (A_1, A_2)

- 1: Fix total order of the goods.
 - 2: Normalize utilities.
 - 3: Let $A_1 = \{g \in M \mid u_1(g) \geq 0 \wedge u_2(g) = 0\}$ and $A_2 = \{g \in M \mid u_1(g) = 0 \wedge u_2(g) > 0\}$.
 - 4: Let $M^+ = \{g \in M \mid u_1(g) > 0 \wedge u_2(g) > 0\}$.
 - 5: **if** $M^+ \neq \emptyset$ **then**
 - 6: Arrange the goods in M^+ in non-increasing order based on their utility ratios: $\frac{u_1(g_1)}{u_2(g_1)} \geq \dots \geq \frac{u_1(g_\ell)}{u_2(g_\ell)}$.
 - 7: Find $P_1 = \{g_1, \dots, g_{k-1}\}$, $P_2 = \{g_{k+1}, \dots, g_\ell\}$ and g_k with λ_1 and λ_2 (where $\lambda_1 + \lambda_2 = 1$ and $\lambda_1, \lambda_2 \geq 0$) such that $\frac{1}{u_1(M^+)}(u_1(P_1) + \lambda_1 u_1(g_k)) = \frac{1}{u_2(M^+)}(u_2(P_2) + \lambda_2 u_2(g_k))$.
 - 8: **if** $\lambda_1 \geq \lambda_2$ **then**
 - 9: $A_1 \leftarrow A_1 \cup P_1 \cup \{g_k\}$ and $A_2 \leftarrow A_2 \cup P_2$
 - 10: **else**
 - 11: $A_1 \leftarrow A_1 \cup P_1$ and $A_2 \leftarrow A_2 \cup P_2 \cup \{g_k\}$
 - 12: **return** (A_1, A_2)
-

term is at least 1. This yields $\text{NW}(A')^2 \text{NW}(B')^2 > \text{NW}(A)^2 \text{NW}(B)^2$, a contradiction. \square

Consider a mechanism that returns an MNW allocation, breaking ties according to agents' positions. This mechanism is scale-invariant by the definition of Nash welfare. Furthermore, the resulting allocation satisfies EF1 and PO (Caragiannis et al. 2019). By Theorem 6, we obtain the following theorem.

Theorem 7. *When $n = 2$, any mechanism that returns an MNW allocation is scale-invariant and satisfies PEF1 with respect to $U_{>0}$. Moreover, such allocations are EF1 and PO.*

4.2 Adjusted Winner Mechanism

For two agents, we prove the existence of a scale-invariant and PEF1 mechanism that always produces EF1 and PO in polynomial time by considering the adjusted winner mechanism (Brams and Taylor 1996; Aziz et al. 2015, 2022a).

Theorem 8. *When $n = 2$, there exists a scale-invariant, PEF1 mechanism that always returns an EF1 and PO allocation in polynomial time.*

See Algorithm 3. The mechanism first fixes a total order of goods and partitions the goods into three sets: A_1 containing goods valued only by position 1, A_2 containing goods valued only by position 2, and M^+ containing goods positively valued by both agents. We then arrange goods in M^+ in non-increasing order of utility ratios $\frac{u_1(g)}{u_2(g)}$, breaking ties by the indices of goods. We denote the sequence by g_1, g_2, \dots, g_ℓ where $\ell = |M^+|$.

We consider dividing fractionally these ordered goods using a boundary line: goods to the left of the boundary are allocated to agent in position 1, and goods to the right are allocated to agent in position 2. Formally, as we move a boundary from left to right, there exists a unique

$k \in [\ell]$ and parameters λ_1, λ_2 where allocating bundle $P_1 = \{g_1, g_2, \dots, g_{k-1}\}$ to agent in position 1, bundle $P_2 = \{g_{k+1}, g_{k+2}, \dots, g_\ell\}$ to agent in position 2, and splitting good g_k in proportions λ_1, λ_2 gives equal utility to both agents. Specifically, $\frac{1}{u_1(S)}(u_1(P_1) + \lambda_1 u_1(g_k)) = \frac{1}{u_2(S)}(u_2(P_2) + \lambda_2 u_2(g_k))$. Based on the comparison of λ_1 and λ_2 , we allocate the boundary good g_k entirely to one of the agents (see Lines 8 and 10 in Algorithm 3).

From Section 3 of (Aziz et al. 2015), we have the following lemma.

Lemma 9 (Aziz et al. (2015)). *For every pair of positions $i, j \in \{1, 2\}$, $u_i(P_i) + \lambda_i u_i(g_k) \geq u_i(P_j) + \lambda_j u_i(g_k)$. Moreover, no partition of M^+ between the two agents can make one agent better off without making the other agent worse off compared to either $(P_1 \cup \{g_k\}, P_2)$ or $(P_1, P_2 \cup \{g_k\})$.*

Using this lemma, we prove Theorem 8.

Proof of Theorem 8. Let $N = \{a_1, a_2\}$ be the set of two agents. For each $a \in N$, let $M_a = \{g \in M \mid u_a(g) > 0 \wedge u_{a'}(g) = 0\}$ be the set of goods valued only by agent a where a' denotes the other agent, and let $M_0 = \{g \in M \mid u_{a_1}(g) = u_{a_2}(g) = 0\}$.

We first prove that the mechanism is PEF1. The boundary line and the proportions λ_1, λ_2 are determined solely by the utility ratios, independently of agent orderings. Bundles P_1 and P_2 are determined by the boundary line, which is independent of positions. For any agent $a \in N$ and agent orderings π and π' , we denote by P_a the fixed set $P_{\pi(a)} = P_{\pi'(a)}$ of goods in M^+ . Under ordering π , agent a in position $\pi(a)$ receives either $M_a \cup P_a$, $M_a \cup M_0 \cup P_a$, $M_a \cup P_a \cup \{g_k\}$, or $M_a \cup M_0 \cup P_a \cup \{g_k\}$. Since $u_a(M_0) = 0$, agent a 's utility equals either $u_a(M_a \cup P_a)$ or $u_a(M_a \cup P_a \cup \{g_k\})$, establishing PEF1.

We now show that the mechanism always returns an EF1 allocation. Fix an agent ordering π . Without loss of generality, we assume that $\pi(a_1) = 1$ and $\pi(a_2) = 2$. When $\lambda_1 \geq \lambda_2$, we have

$$\begin{aligned} u_{a_1}(\mathcal{M}(u_\pi)_{\pi(a_1)}) &= u_{a_1}(M_{a_1} \cup P_{a_1}) + u_{a_1}(g_k) \\ &\geq u_{a_1}(M_{a_2} \cup P_{a_2}) + \lambda_2 u_{a_1}(g_k) + (1 - \lambda_1) u_{a_1}(g_k) \\ &\geq u_{a_1}(M_{a_2} \cup P_{a_2}) = u_{a_1}(\mathcal{M}(u_\pi)_{\pi(a_2)}), \end{aligned}$$

where we use Lemma 9 for the first inequality, and

$$\begin{aligned} u_{a_2}(\mathcal{M}(u_\pi)_{\pi(a_2)}) &= u_{a_2}(M_{a_2} \cup P_{a_2}) \\ &\geq u_{a_2}(M_{a_1} \cup P_{a_1}) + \lambda_1 u_{a_2}(g_k) - \lambda_2 u_{a_2}(g_k) \\ &\geq u_{a_2}(M_{a_1} \cup P_{a_1}) = u_{a_2}(\mathcal{M}(u_\pi)_{\pi(a_1)} \setminus \{g_k\}), \end{aligned}$$

where we use Lemma 9 for the first inequality, and $\lambda_1 \geq \lambda_2$ for the second inequality. When $\lambda_1 < \lambda_2$, the resulting allocation can be proven to be EF1 by an analogous argument to the case of $\lambda_1 \geq \lambda_2$.

Next, we prove that the mechanism produces Pareto optimal allocations. Let $\{A_{a_1}, A_{a_2}\}$ be the allocation induced from the ordered allocation produced by the mechanism. By Lemma 9, since goods in M_a are valued only by agent a for each $a \in N$, and goods in M_0 are valued by neither agent, any allocation that differs from $\{A_{a_1}, A_{a_2}\}$ would

make at least one agent worse off. Finally, the mechanism clearly runs in polynomial time, and its scale-invariance follows from the normalization step. \square

5 Further Analysis for Round-Robin Mechanism

We now investigate the round-robin mechanism (recall Algorithm 1). We refer to each iteration of the while loop in Algorithm 1 as a *round*, in which agents select their most preferred good from the remaining goods according to a fixed agent ordering. As mentioned in Section 2, for the round-robin mechanism, more goods need to be removed to eliminate position-based envy as n increases.

Theorem 10. *When $\lceil m/n \rceil \geq \lfloor \log_2 n \rfloor$, for the round-robin mechanism, there exists a profile and two agent orderings π, π' where even after removing any $\lfloor \log_2 n \rfloor - 1$ goods from their bundle under π , an agent prefers keeping their remaining bundle to receiving their bundle under π' .*

We prove the theorem by constructing a specific profile. The detailed proof is deferred to the full version of our paper (Mahara et al. 2025).

Proof Sketch. Consider two agent orderings $\pi, \pi' \in \Pi$ such that $\pi(a_i) = i$ and $\pi'(a_i) = n + 1 - i$ for each $i \in [n]$. The key idea is to construct utilities where agent a_1 's first $\lfloor \log_2 n \rfloor$ choices under π have significantly higher value (value C) than the remaining choices, while carefully setting the utilities of other agents to ensure that under π' , agent a_1 cannot obtain any of these highly valued goods. This construction ensures that even after removing any $\lfloor \log_2 n \rfloor - 1$ goods from agent a_1 's bundle under π , at least one good of value C remains, making this bundle more valuable than their bundle under π' . \square

Theorem 10 shows that the round-robin mechanism is not PEF1 when $n \geq 4$ and $m \geq n + 1$. When $m \leq n$, each agent receives at most one good, ensuring PEF1. Moreover, we prove that the round-robin mechanism is PEF1 for two or three agents. We defer the proof to the full version of our paper (Mahara et al. 2025).

Theorem 11. *When $n \in \{2, 3\}$, the round-robin mechanism is PEF1.*

6 Envy-Cycle Mechanism

Next, we study the envy-cycle mechanism, which always produces an EF1 allocation when agents have monotone utility functions (Lipton et al. 2004). We say that u_a is *monotone* if $u_a(S) \leq u_a(T)$ for any $S \subseteq T \subseteq M$. We will show the mechanism may not satisfy PEF1. All proofs are presented in the full version (Mahara et al. 2025).

To describe the mechanism, we define several concepts. We say that $P = (P_1, P_2, \dots, P_n)$ is an *ordered partial allocation* if $\bigcup_{i \in [n]} P_i \subseteq M$ and $P_i \cap P_{i'} = \emptyset$ for all $i \neq i' \in [n]$. For an ordered profile u and a partial allocation $P = (P_1, P_2, \dots, P_n)$, *envy graph* is defined as a directed graph $G_P = ([n], E)$, where the vertex set is $[n]$ and the edge set is $E = \{(i, i') \mid i, i' \in [n], i \neq i', u_i(P_i) <$

Algorithm 4: Envy-Cycle Mechanism

Input: Ordered profile (u_1, u_2, \dots, u_n)

Output: Ordered allocation (A_1, A_2, \dots, A_n)

- 1: Fix indices of goods as $M = \{g_1, g_2, \dots, g_m\}$.
 - 2: Set $A_i \leftarrow \emptyset$ for all $i \in [n]$.
 - 3: **for** $j = 1, 2, \dots, m$ **do**
 - 4: For partial allocation $A = \{A_i\}_{i \in [n]}$, construct the envy graph G_A .
 - 5: **while** there exists an envy cycle in G_A **do**
 - 6: Resolve the envy cycle by transferring bundles in the opposite direction of the cycle.
 - 7: Let $i \in [n]$ be the vertex of in-degree 0 in G_A with the smallest index.
 - 8: Set $A_i \leftarrow A_i \cup \{g_j\}$.
 - 9: **return** (A_1, A_2, \dots, A_n)
-

$u_i(P_{i'})\}$. An *envy cycle* is a directed cycle in the envy graph, that is, a sequence of positions $(i_1, i_2, \dots, i_\ell)$ such that $(i_k, i_{k+1}) \in E$ for all $k \in [\ell]$, where $i_{\ell+1} = i_1$.

In the envy-cycle mechanism (Algorithm 4), we first order the goods arbitrarily. For each good, we first eliminate all envy cycles in the partial allocation and then allocate it to an unenvied position. Specifically, while there exists an envy cycle in the current allocation, we resolve it by transferring bundles in the opposite direction of the cycle. After eliminating all cycles, we allocate the current good to a position not envied by any other. If multiple such positions exist, we choose the one with the smallest index.

Similar to the round-robin mechanism, the envy-cycle mechanism does not satisfy PEF1 in general.

Theorem 12. *When $n \geq 2$, for the envy-cycle mechanism, there exists a profile and two agent orderings π, π' where even after removing any $m - \lfloor \frac{m}{n} \rfloor - 1$ goods from their bundle under π , an agent prefers keeping their remaining bundle to receiving their bundle under π' .*

Theorem 12 implies that if the envy-cycle mechanism is PEF1, then $m - \lfloor \frac{m}{n} \rfloor \leq 1$. Moreover, we establish that this condition is also sufficient for PEF1.

Theorem 13. *The envy-cycle mechanism is PEF1 if $m - \lfloor \frac{m}{n} \rfloor \leq 1$.*

7 Discussion

This paper introduces a new fairness notion, PEF1, for mechanisms. We demonstrate a PEF1 mechanism producing an EF1 allocation for agents with additive utilities. For the case of two agents, we prove the existence of a scale-invariant, PEF1 mechanism that outputs an EF1 and PO allocation.

Several questions remain open for future research. While we have shown in Theorem 6 that for $n = 2$, the utility difference between any pair of MNW allocations is bounded by some good's utility for each agent, the result for $n > 2$ remains unknown. We conjecture this bound holds for any n . Additionally, the existence of a PEF1 mechanism producing an EF1 and PO allocation for any n remains an open question.

Acknowledgments

This work was partially supported by JST ERATO Grant Number JPMJER2301. Ryoga Mahara acknowledges additional support from JSPS KAKENHI Grant Number JP23K19956. Ryuhei Mizutani acknowledges additional support from JSPS KAKENHI Grant Number JP23KJ0379 and JST SPRING Grant Number JPMJSP2108. Taihei Oki acknowledges additional support from JST FOREST Grant Number JPMJFR232L. We thank the anonymous AAAI 2026 reviewers for their valuable feedback.

References

- Amanatidis, G.; Aziz, H.; Birmpas, G.; Filos-Ratsikas, A.; Li, B.; Moulin, H.; Voudouris, A. A.; and Wu, X. 2023. Fair division of indivisible goods: Recent progress and open questions. *Artificial Intelligence*, 322: 103965.
- Aziz, H.; Brânzei, S.; Filos-Ratsikas, A.; Kristoffer, S.; and Frederiksen, S. 2015. The adjusted winner procedure: Characterizations and equilibria. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 454–460.
- Aziz, H.; Caragiannis, I.; Igarashi, A.; and Walsh, T. 2022a. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi-Agent Systems*, 36: 1–21.
- Aziz, H.; Li, B.; Moulin, H.; and Wu, X. 2022b. Algorithmic fair allocation of indivisible items: A survey and new questions. *ACM SIGecom Exchanges*, 20(1): 24–40.
- Bei, X.; Huzhang, G.; and Suksompong, W. 2020. Truthful fair division without free disposal. *Social Choice and Welfare*, 55: 523–545.
- Bogomolnaia, A.; and Moulin, H. 2001. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2): 295–328.
- Brams, S. J.; and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- Brustle, J.; Dippel, J.; Narayan, V. V.; Suzuki, M.; and Vetta, A. 2020. One dollar each eliminates envy. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 23–39.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(4): 1061–1103.
- Budish, E.; Cachon, G. P.; Kessler, J. B.; and Othman, A. 2016. Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research*, 65(2): 314–336.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3): 1–12.
- Dubins, L. E.; and Spanier, E. H. 1961. How to cut a cake fairly. *The American Mathematical Monthly*, 68(1): 1–17.
- Foley, D. K. 1967. Resource allocation and the public sector. *Yale Economic Essays*, 7: 45–98.
- Gibbard, A. 1973. Manipulation of voting schemes: A general result. *Econometrica*, 41(4): 587–601.
- Goldman, J.; and Procaccia, A. D. 2015. Spliddit: Unleashing fair division algorithms. *ACM SIGecom Exchanges*, 13(2): 41–46.
- Han, J.; and Suksompong, W. 2024. Fast & fair: A collaborative platform for fair division applications. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 23796–23798.
- Igarashi, A.; and Yokoyama, T. 2023. Kajibuntan: A house chore division app. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 16449–16451.
- Lee, E. 2017. APX-hardness of maximizing Nash social welfare with indivisible items. *Information Processing Letters*, 122: 17–20.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, 125–131.
- Lovász, L.; and Plummer, M. D. 2009. *Matching Theory*. American Mathematical Society.
- Mahara, R.; Mizutani, R.; Oki, T.; and Yokoyama, T. 2025. Position fair mechanisms allocating indivisible goods. *CoRR*, abs/2409.06423.
- Manabe, Y.; and Okamoto, T. 2012. Meta-envy-free cake-cutting and pie-cutting protocols. *Journal of Information Processing*, 20(3): 686–693.
- Moulin, H. 2004. *Fair Division and Collective Welfare*. MIT press.
- Nguyen, N.-T.; Nguyen, T. T.; Roos, M.; and Rothe, J. 2014. Computational complexity and approximability of social welfare optimization in multiagent resource allocation. *Autonomous Agents and Multi-Agent Systems*, 28(2): 256–289.
- Roth, A. E.; Sönmez, T.; and Ünver, M. U. 2005. Pairwise kidney exchange. *Journal of Economic Theory*, 125(2): 151–188.
- Satterthwaite, M. A. 1975. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2): 187–217.
- Segal-Halevi, E.; and Sziklai, B. R. 2019. Monotonicity and competitive equilibrium in cake-cutting. *Economic Theory*, 68(2): 363–401.
- Shapley, L.; and Scarf, H. 1974. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1): 23–37.
- Sönmez, T. 1999. Strategy-proofness and essentially single-valued cores. *Econometrica*, 67(3): 677–689.
- Walsh, T. 2020. Fair division: The computer scientist’s perspective. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 4966–4972.
- Zhou, L. 1990. On a conjecture by Gale about one-sided matching problems. *Journal of Economic Theory*, 52(1): 123–135.