

Fair and Efficient Balanced Allocation for Indivisible Goods

Yasushi Kawase^{1,2}, Ryoga Mahara¹

¹The University of Tokyo, Tokyo, Japan

²RIKEN Center for Advanced Intelligence Project, Tokyo, Japan
kawase@mist.i.u-tokyo.ac.jp, mahara@mist.i.u-tokyo.ac.jp

Abstract

We study the problem of allocating indivisible goods among agents with additive valuation functions to achieve both fairness and efficiency under the constraint that each agent receives exactly the same number of goods (the *balanced constraint*). While this constraint is common in real-world scenarios such as team drafts or asset division, it significantly complicates the search for allocations that are both fair and efficient. Envy-freeness up to one good (EF1) is a well-established fairness notion for indivisible goods. Pareto optimality (PO) and its stronger variant, fractional Pareto optimality (fPO), are widely accepted efficiency criteria. Our main contribution establishes both the existence and polynomial-time computability of allocations that are simultaneously EF1 and fPO under balanced constraints in two fundamental cases: (1) when agents have at most two distinct types of valuation functions, and (2) when each agent has a personalized bivalued valuation. Our algorithms leverage novel applications of maximum-weight matching in bipartite graphs and duality theory, providing the first polynomial-time solutions for these cases and offering new insights for constrained fair division problems.

Introduction

The fair division of *indivisible* goods is an important problem that has been widely studied in mathematics, economics, and computer science (Brams and Taylor 1996; Brandt et al. 2016). In recent years, this topic has attracted even more interest (see surveys (Amanatidis et al. 2023; Aziz et al. 2022b; Guo, Li, and Deng 2023)). Many previous studies have discussed fairness and efficiency under the assumption that there are no constraints on allocation. However, in real-world problems, it is often necessary to make allocations under various constraints. Motivated by this, recent works have studied fair division problems under a variety of constraints (see the survey (Suksompong 2021)).

For instance, in team sports drafts, new players are assigned to teams, and it is typical for each team to receive the same number of new players to prevent any team from having a numerical advantage. Another example is the division of indivisible assets, such as inherited jewelry or artwork, among family members. Siblings often agree to take

the same number of items, but differences in market value or personal significance may still lead to envy. These examples highlight the importance of considering allocations that maintain balanced quantities while also satisfying fairness and efficiency.

In this paper, we study the problem of finding fair and efficient allocations under the constraint that each agent receives the same number of goods (the *balanced constraint*), assuming that the total number of goods is a multiple of the number of agents. We also assume that each agent has an additive valuation function.

To formalize notions of fairness in such settings, several criteria have been introduced in the literature. *Envy-freeness* (EF) (Foley 1966) is one of the most fundamental fairness concepts. It requires that no agent prefers someone else’s bundle over their own. However, with indivisible items, achieving EF can be impossible even in very simple scenarios, such as when there are only two agents and a single item. To address this issue, various relaxed notions of fairness have been proposed. The most prominent among these is *envy-freeness up to one item* (EF1) introduced by Budish (2011). EF1 requires that each agent prefers their own bundle to that of any other agent after removing at most one item from the other agent’s bundle. It is known that an EF1 allocation always exists and can be computed in polynomial time (Lipton et al. 2004). Moreover, even under the balanced constraint, an EF1 allocation can be found using the round-robin algorithm (Caragiannis et al. 2019).

On the efficiency side, *Pareto optimality* (PO) is a standard benchmark. An allocation is PO if no agent can be made strictly better off without making someone else worse off. A stronger concept, known as *fractional Pareto optimality* (fPO), extends this notion to fractional allocations. It is known that maximizing utilitarian welfare yields an fPO allocation, and thus also a PO allocation. Consequently, such an allocation always exists and can be computed in polynomial time.

While finding allocations that are either EF1 or PO individually is relatively straightforward, the problem becomes considerably more complex when both fairness and efficiency must be achieved at the same time. In what follows, we examine the main challenges in satisfying both criteria simultaneously and review previous progress on addressing this issue.

For the unconstrained case, Caragiannis et al. (2019) established the novel result that maximizing the Nash social welfare (Nash 1950; Kaneko and Nakamura 1979), defined as the geometric mean of the agents' valuations, yields an allocation that is both EF1 and PO (but may not be fPO). However, since maximizing Nash social welfare is NP-hard (Nguyen et al. 2014) and even APX-hard (Lee 2017), this approach does not directly yield an efficient algorithm. Barman, Krishnamurthy, and Vaish (2018) addressed this computational limitation by proposing a pseudo-polynomial time algorithm that computes an EF1 and PO allocation and proved the existence of an EF1 and fPO allocation. Subsequently, Mahara (2024b) developed a polynomial-time algorithm to compute an EF1 and fPO allocation when the number of agents is constant. The existence of a polynomial-time algorithm for finding an EF1 and PO (or fPO) allocation in general case remains an important open question.

There is growing interest in fair and efficient allocation under various practical constraints that extend beyond the unconstrained setting. When such constraints are present, achieving both EF1 and PO¹ becomes more challenging than in the unconstrained case. A natural first approach is to maximize Nash social welfare among allocations that satisfy the constraints, but in general, such an approach does not lead to an EF1 allocation. Furthermore, under general constraints, there may not exist any allocation that is both EF1 and PO (Wu, Li, and Gan 2021; Cookson, Ebadian, and Shah 2025). To date, the only known result guaranteeing the existence of an allocation that is both EF1 and PO under constraints is due to Shoshan, Hazon, and Segal-Halevi (2023), who studied fair division under category constraints and developed a polynomial-time algorithm to find such an allocation for the case of two agents.

Our Results

In this paper, we study the problem of fair division of indivisible goods among agents with additive valuation functions under the *balanced constraints*.

We first observe that maximizing Nash social welfare over balanced allocations does not necessarily guarantee EF1 (see Example 1). We then give a characterization that an allocation is fPO if and only if it maximizes a weighted utilitarian social welfare (Proposition 1). We also show that checking whether a given balanced allocation satisfies fPO and EF1 is polynomial-time solvable, whereas checking PO is coNP-complete. Additionally, we provide a reduction from the problem of finding an EF1 and fPO allocation under unconstrained case to balanced constraint case.

We establish both the existence and polynomial-time computability of balanced allocations that are simultaneously EF1 and fPO in two important cases: when there are

¹In the constrained case, we consider PO and fPO only with respect to feasible allocations (i.e., constrained PO and constrained fPO). It is not difficult to see that PO (or fPO) with respect to all allocations may not exist. For example, suppose there are two agents and two goods, and one agent values both goods positively while the other agent assigns zero value to both. Then, the PO allocation with respect to all allocations must assign both goods to the agent who values them, but such an allocation is not balanced.

at most two types of agents, and when each agent has a personalized bivalued valuation, i.e., each agent i assigns two distinct values a_i, b_i ($a_i > b_i \geq 0$) to the goods.

In the personalized bivalued case, we show that maximizing the weighted utilitarian social welfare with appropriately chosen weights yields an EF1 and fPO allocation.

In the two types case, we provide an algorithm that maintains primal and dual optimal solutions for the linear program that maximizes weighted social welfare. The algorithm searches for an EF1 allocation by gradually changing the weights. In this procedure, we exploit a characterization of dual optimal solutions for maximum-weight perfect matchings in bipartite graphs using the shortest path distances. It is worth mentioning that our result for the two types case implies the result for the unconstrained case.

It is worth noting that the balanced constraints are a special case of the category constraints studied in (Shoshan, Hazon, and Segal-Halevi 2023; Igarashi and Meunier 2025) and the matroid constraints considered in (Kawase and Sumita 2020; Kawase, Sumita, and Yokoi 2023; Cookson, Ebadian, and Shah 2025). Moreover, both cases analyzed in this paper have been widely studied in the recent fair division literature (see Related Work).

Due to space limitations, we omit most of the proofs, which can be found in the full paper.

Related Work

Fair and efficient allocation for divisible items The Fisher market framework, originally introduced by Irving Fisher (see (Brainard and Scarf 2005)), has long been a central object of study in both economics and computer science. This model with divisible goods is well known for exhibiting strong notions of fairness and efficiency. In particular, Varian (1974) showed that when all agents have equal budgets, the equilibrium allocation achieved in a Fisher market is both EF and PO. It is established that such market equilibria can be computed in polynomial time under additive valuation functions (Devanur et al. 2008; Orlin 2010; Végh 2012). For the constrained case, an EF and PO allocation still exists if each agent has an identical constraint (Cole and Tao 2021), but finding it is PPAD-complete even in the balanced setting where each agent receives exactly one good (Tröbst and Vazirani 2024; Caragiannis, Hansen, and Rathi 2024). Moreover, there exists an allocation that is both SD-EF (which is a stronger notion than EF and balancedness) and ordinally efficient (which is slightly weaker notion than PO) (Kojima 2009); however, such an allocation may not exist even under category constraint (Kawase, Sumita, and Yokoi 2023).

Fair and efficient allocation for indivisible chores For indivisible chores, the existence of EF1 and fPO (or PO) allocations was previously known for a few restricted cases: instances with two agents (Aziz et al. 2022a); instances with bivalued valuations for each agent (Ebadian, Peters, and Shah 2022; Garg, Murhekar, and Qin 2022); instances with two types of chores (Aziz et al. 2023); instances with three agents (Garg, Murhekar, and Qin 2023); and instances with three types of valuation functions (Garg, Murhekar, and Qin 2024). In a significant breakthrough, Mahara (2025) recently

showed the existence of EF1 and fPO allocations for general additive valuations.

Fair and efficient allocation under constraints Shoshan, Hazon, and Segal-Halevi (2023) proposed a polynomial-time algorithm for fair division under category constraints with two agents, which finds an EF1 and PO allocation when each category consists of only goods or only chores. If goods and chores are mixed, the algorithm finds an EF[1,1] (envy-free up to one good and one chore) and PO allocation. Igarashi and Meunier (2025) extended this result to general settings with n agents, proving the existence of a PO allocation in which each agent can be made envy-free by reallocating at most $n(n-1)$ items.

For budget constraints, Wu, Li, and Gan (2021) showed that any budget-feasible allocation that maximizes the Nash social welfare achieves a $1/4$ -EF1 and PO allocation for goods. They also showed that there exists an instance in which there is no $(1/2 + \varepsilon)$ -EF1 and PO allocation for any $\varepsilon > 0$. Here, α -EF1 is an approximate relaxation of EF1, which requires that each agent prefers their own bundle to α times the bundle of any other agent after the removal of at most one item from that other agent's bundle. Cookson, Ebadian, and Shah (2025) investigated fair division under matroid constraints, showing that maximizing the Nash social welfare yields a $1/2$ -EF1 and PO allocation for goods.

For a broader overview of fair division under various constraints, see the survey by Suksompong (2021).

Bivalued and two-type instances For bivalued instances², Amanatidis et al. (2021) showed that an EFX allocation can be computed in polynomial time for goods. Here, EFX (envy-freeness up to any good) is a stronger fairness concept than EF1. Garg and Murhekar (2023) further showed that both EFX and PO allocations can be computed in polynomial time for goods. In the context of chores, Garg, Murhekar, and Qin (2022) and Ebadian, Peters, and Shah (2022) proved that EF1 and PO allocations can be computed in polynomial time. Additionally, Garg, Murhekar, and Qin (2023) showed that EFX and PO allocations can be computed in polynomial time for chores when there are three agents. Akrami et al. (2022) established that maximizing the Nash social welfare is polynomial-time computable when the two values are multiples of each other, but becomes NP-hard when they are coprime.

For instances with two types of valuations, Mahara (2023, 2024a) showed that an EFX allocation always exists for goods under additive valuations, and even under more general monotone valuations. In addition, Garg, Murhekar, and Qin (2023) showed that EF1 and PO allocations can be computed in polynomial time for both goods and chores.

Model

For each natural number ℓ , we denote $[\ell] = \{1, \dots, \ell\}$.

²In a bivalued instance, each agent i assigns one of two fixed values, $a > b > 0$, to the goods. This setting is a special case of the personalized bivalued instance, where each agent may assign their own pair of distinct values.

An instance of our problem is a tuple $(N, M, (v_i)_{i \in N})$, where $N = [n]$ represents the (non-empty) set of agents and $M = [m]$ represents the (non-empty) set of *indivisible* goods. Each agent i has a *valuation function*, denoted as $v_i: M \rightarrow \mathbb{R}_+$, where \mathbb{R}_+ represents the set of nonnegative real numbers. We assume that each agent's valuation is *additive*, and write $v_i(X) := \sum_{j \in X} v_{ij}$ to denote the utility of agent i when i receives a subset of goods $X \subseteq M$. A valuation function v_i is called *personalized bivalued* if there exist $a_i, b_i \in \mathbb{R}_+$ with $a_i > b_i$ such that $v_{ij} \in \{a_i, b_i\}$ for all $j \in M$.

Throughout this paper, we assume that m is a multiple of n , and define $k := m/n$. An (integral) *allocation* is an ordered partition $\mathbf{A} = (A_1, \dots, A_n)$ of M , i.e., $\bigcup_{i \in N} A_i = M$ and $A_i \cap A_{i'} = \emptyset$ for any distinct $i, i' \in N$. Each A_i is the *bundle* allocated to agent i . An allocation \mathbf{A} is called *balanced* if $|A_i| = k$ for all $i \in N$. We sometimes represent a balanced allocation \mathbf{A} by a matrix $x \in \{0, 1\}^{N \times M}$, where $x_{ij} = 1$ if good j is allocated to agent i (i.e., $j \in A_i$), and $x_{ij} = 0$ otherwise. A *balanced fractional allocation* is a matrix $x \in \mathbb{R}_+^{N \times M}$ such that $\sum_{j \in M} x_{ij} = k$ for all $i \in N$ and $\sum_{i \in N} x_{ij} = 1$ for all $j \in M$. By the property of the total unimodularity, the set of balanced fractional allocations forms the convex hull of the balanced (integral) allocations (see, e.g., (Schrijver 2003, Sec. 21.2)). Thus, a balanced fractional allocation can be interpreted as a lottery over balanced allocations.

An allocation \mathbf{A} is called *envy-free up to one good (EF1)* if, for all $i, i' \in N$, either $A_{i'} = \emptyset$, or there exists a good $j \in A_{i'}$ such that $v_i(A_i) \geq v_i(A_{i'} \setminus \{j\})$. A balanced (fractional) allocation $x \in \mathbb{R}_+^{N \times M}$ *Pareto-dominates* a balanced (fractional) allocation $x' \in \mathbb{R}_+^{N \times M}$ if (i) $\sum_{j \in M} v_{ij} x'_{ij} \leq \sum_{j \in M} v_{ij} x_{ij}$ for all $i \in N$, and (ii) $\sum_{j \in M} v_{ij} x'_{ij} < \sum_{j \in M} v_{ij} x_{ij}$ for some $i \in N$. A balanced allocation is called *Pareto optimal (PO)* or *fractionally Pareto optimal (fPO)* if it cannot be Pareto-dominated by any other balanced integral or fractional allocation, respectively.

Observe that the set of achievable valuation vectors $(\sum_{j \in M} v_{1j} x_{1j}, \dots, \sum_{j \in M} v_{nj} x_{nj})$ arising from balanced fractional allocation x forms a polytope, which is precisely the convex hull of the set of valuation vectors $(v_1(A_1), \dots, v_n(A_n))$ corresponding to balanced allocations \mathbf{A} . Thus, the fPO allocations can be characterized within the framework of multi-objective linear programming (see, e.g., (Ehrgott 2005)).

Proposition 1. *A balanced allocation \mathbf{A} is fPO if and only if it maximizes the weighted sum (weighted utilitarian social welfare) $\sum_{i \in N} \alpha_i v_i(A_i)$ for some positive weight vector $\alpha \in \mathbb{R}_{++}^N$.*

Note that, for any given strictly positive weight vector α , the maximum value $\sum_{i \in N} \alpha_i v_i(A_i)$ taken over all balanced allocations \mathbf{A} can be computed by the following linear pro-

gramming (LP):

$$\begin{aligned} \max \quad & \sum_{i \in N} \sum_{j \in M} \alpha_i v_{ij} x_{ij} \\ \text{s.t.} \quad & \sum_{i \in N} x_{ij} = 1 \quad \forall j \in M, \\ & \sum_{j \in M} x_{ij} = k \quad \forall i \in N, \\ & x_{ij} \geq 0 \quad \forall i \in N, \forall j \in M. \end{aligned} \quad (1)$$

To illustrate the above concepts, we present a concrete example.

Example 1. Consider an instance with $N = \{1, 2\}$, $M = \{1, 2, 3, 4\}$, and the valuations are

$$\begin{aligned} (v_{11}, v_{12}, v_{13}, v_{14}) &= (10, 10, 21, 22), \text{ and} \\ (v_{21}, v_{22}, v_{23}, v_{24}) &= (0, 1, 6, 8). \end{aligned}$$

The possible valuation vectors for balanced allocations are visualized in Figure 1. This instance has a unique EF1 and fPO allocation: $(\{1, 3\}, \{2, 4\})$. Note that the allocation $(\{1, 4\}, \{2, 3\})$ is PO but not fPO. Among the balanced allocations, $(\{1, 2\}, \{3, 4\})$ maximizes the Nash social welfare, but it is not EF1. Moreover, in this instance, no EF fractional allocation (i.e., ex ante EF) can be represented as a lottery over EF1 and fPO allocations (i.e., ex post EF1 and fPO).




Figure 1: Possible valuation vectors for balanced allocations in Example 1. Square markers indicate fPO allocations and red markers denote EF1 allocations. The blue region represents the set of possible valuation vectors for fractional balanced allocations.

Properties of EF1 and fPO Allocations

In this section, we provide some properties of EF1 and fPO that will be used in subsequent discussions.

It is well known that an EF1 balanced allocation can be computed by the following *round-robin procedure* (Caragiannis et al. 2019): First, an ordering of the agents is fixed. Then, according to this ordering, the agents take turns sequentially to choose their favorite available good (breaking ties arbitrarily). This process is repeated in multiple rounds until all goods have been allocated. We call the outcome of the round-robin procedure the *round-robin allocation*.

For a balanced allocation $\mathbf{A} = (A_1, \dots, A_n)$ and a weight vector $\alpha \in \mathbb{R}_{++}^N$, define a weighted directed graph $G_{\mathbf{A}}^{(\alpha)} = (N \dot{\cup} M \dot{\cup} \{r\}; \vec{E} \cup \overleftarrow{E} \cup E^+, w^{(\alpha)})$, where $\vec{E} = N \times M$, $\overleftarrow{E} = \{(j, i) \in M \times N \mid j \in A_i\}$, $E^+ = \{r\} \times M$, $w^{(\alpha)}(i, j) = -\alpha_i v_{ij}$ for all $(i, j) \in \vec{E}$,

$w^{(\alpha)}(j, i) = \alpha_i v_{ij}$ for all $(j, i) \in \overleftarrow{E}$, and $w^{(\alpha)}(r, j) = 0$ for all $j \in M$. It is not difficult to see that \mathbf{A} maximizes $\sum_{i \in N} \alpha_i v_i(A_i)$ over all balanced allocations if and only if $G_{\mathbf{A}}^{(\alpha)}$ contains no negative-weight directed cycles.

Next, we characterize the optimal solution of LP (1). Its dual is given by:

$$\begin{aligned} \min \quad & k \cdot \sum_{i \in N} q_i + \sum_{j \in M} p_j \\ \text{s.t.} \quad & q_i + p_j \geq \alpha_i v_{ij} \quad \forall i \in N, \forall j \in M. \end{aligned} \quad (2)$$

We call the dual variables (\mathbf{q}, \mathbf{p}) the *potentials* and \mathbf{p} the *price vector*. By the complementary slackness theorem, a feasible pair \mathbf{x} and (\mathbf{q}, \mathbf{p}) forms optimal solutions to the primal and dual LPs if and only if, for all $i \in N$ and $j \in M$, either $x_{ij} = 0$ or $q_i + p_j = \alpha_i v_{ij}$.

Let \mathbf{A} be a balanced fPO allocation that maximizes the value $\sum_{i \in N} \alpha_i v_i(A_i)$ for some $\alpha \in \mathbb{R}_{++}^N$. Then, we can construct an optimal dual solution as follows (see, e.g., Schrijver (2003, Sec. 17.4)).

Lemma 1 (folklore). Let \mathbf{A} be a balanced allocation that maximizes $\sum_{i \in N} \alpha_i v_i(A_i)$ for some $\alpha \in \mathbb{R}_{++}^N$. Then, an optimal dual solution (\mathbf{q}, \mathbf{p}) to (2) can be obtained by setting q_i as the shortest length of r - i path for each $i \in N$ and $-p_j$ as the shortest length of r - j path in $G_{\mathbf{A}}^{(\alpha)}$. Here, the solution is nonnegative, i.e., $(\mathbf{q}, \mathbf{p}) \in \mathbb{R}_+^N \times \mathbb{R}_+^M$. Moreover, the potentials computed in this way do not depend on the choice of \mathbf{A} .

We introduce price envy-freeness up to one good (\mathbf{p} -EF1), which is a key concept introduced by Barman, Krishnamurthy, and Vaish (2018). This notion has been widely used in the literature for computing EF1 and PO allocations (Barman, Krishnamurthy, and Vaish 2018; Ebadian, Peters, and Shah 2022; Mahara 2024b; Garg, Murhekar, and Qin 2022, 2023, 2024; Mahara 2025).

Definition 1 (\mathbf{p} -EF1). Let (\mathbf{q}, \mathbf{p}) be potentials. For any non-empty set of goods $X \subseteq M$, define $p(X) := \sum_{j \in X} p_j$ and $\hat{p}(X) := p(X) - \max_{j \in X} p_j$. An allocation \mathbf{A} is called \mathbf{p} -EF1 if $p(A_i) \geq \hat{p}(A_{i'})$ for any pair of agents $i, i' \in N$.

For an appropriate choice of prices \mathbf{p} , \mathbf{p} -EF1 implies EF1.

Lemma 2. Given a weight vector $\alpha \in \mathbb{R}_{++}^N$, let \mathbf{A} be an optimal balanced allocation of LP (1), and let (\mathbf{q}, \mathbf{p}) be an optimal solution of the dual LP (2) such that $q_i \geq 0$ for all $i \in N$. If \mathbf{A} is \mathbf{p} -EF1, then it is also EF1.

By the LP formulation, we can check whether a given balanced allocation \mathbf{A} is fPO in polynomial time.

Theorem 1. Given a balanced allocation \mathbf{A} , whether it satisfies EF1 and fPO can be checked in polynomial time.

Next, we demonstrate that any instance of the unconstrained problem can be transformed into an equivalent instance with balanced constraint. This is achieved by introducing a set of “dummy” goods that have no value to any agent, allowing for a direct correspondence between the properties of allocations in the two settings. This equivalence is formalized in the following theorem.

Theorem 2. For an unconstrained fair allocation instance $(N, M, (v_i)_{i \in N})$, let M' be a set of dummy goods of size

$|N| \cdot (|M| - 1)$ and let $v'_i(X) = v_i(X \cap M)$ for each $i \in N$ and $X \subseteq M \cup M'$. The following equivalences hold:

- If an allocation A is PO (resp., fPO, EF1) in the unconstrained instance, then a balanced allocation $(A_1 \cup D_1, \dots, A_n \cup D_n)$ where (D_1, \dots, D_n) is an ordered partition of M' is PO (resp., fPO, EF1) in the balanced instance.
- If a balanced allocation A' is PO (resp., fPO, EF1) in the balanced instance, then an allocation $(A_1 \cap M, \dots, A_n \cap M)$ is PO (resp., fPO, EF1) in the unconstrained instance.

This theorem implies that if we construct a polynomial-time algorithm to find an EF1 and fPO balanced allocation for a specific setting, then it can also be used to find an EF1 and fPO allocation for the corresponding unconstrained setting. This result has immediate algorithmic implications. For instance, our polynomial-time algorithm for the two-types case under the balanced constraint can be directly applied to the two-types case in the unconstrained setting. This is because, after the addition of dummy goods, the agent types remain unchanged, ensuring that the resulting balanced instance continues to fall within the two-types framework.

It is known that, in the unconstrained setting, the problem of deciding whether a given allocation is PO is coNP-complete (De Keijzer et al. 2009). By combining this with Theorem 2, we obtain the following corollary.

Corollary 1. *The problem of deciding whether a given balanced allocation is PO is coNP-complete.*

Personalized Bivalued Valuations Case

In this section, we consider the case of personalized bivalued valuations, where each agent $i \in N$ assigns to every good $j \in M$ a value $v_{ij} \in \{a_i, b_i\}$ for some $a_i > b_i \geq 0$. We show that a balanced allocation that is both EF1 and fPO always exists, and that such an allocation can be computed in polynomial time.

Theorem 3. *When each agent has a personalized bivalued valuation, there always exists a balanced allocation that is both EF1 and fPO. Moreover, such an allocation can be found in polynomial time.*

Our algorithm prepare k slots for each agent, and consider perfect matching between slots and goods as a balanced allocation. We set $\varepsilon = 1/(nk(k+1))$ and the weight of each edge between sth slot of agent i and good $j \in M$ as

$$w((i, s), j) = \begin{cases} a_i/(a_i - b_i) + s \cdot \varepsilon & \text{if } v_{ij} = a_i, \\ b_i/(a_i - b_i) & \text{if } v_{ij} = b_i. \end{cases}$$

Our algorithm outputs a balanced allocation A^* corresponding to a maximum weight perfect matching X^* in this weighted bipartite graph. Intuitively, the weights without the $s \cdot \varepsilon$ term are chosen so that the objective function increases by the same amount regardless of which agent receives a high-value good. The $s \cdot \varepsilon$ term is included to ensure that high-value goods are distributed as evenly as possible among the agents.

In the personalized bivalued setting, fPO allocations can be characterized by a specific weight vector α such that

$\alpha_i = 1/(a_i - b_i)$ for each $i \in N$. This motivates the choice of this particular weight vector.

Proposition 2. *For agents with personalized bivalued valuations, a balanced allocation A is fPO if and only if it maximizes the value $\sum_{i \in N} v_i(A'_i)/(a_i - b_i)$ over all balanced allocations.*

We show that A^* computed in our algorithm satisfies both EF1 and fPO.

Lemma 3. *A^* is EF1 and fPO.*

Now we are ready to prove Theorem 3.

Proof of Theorem 3. We can compute the maximum-weight perfect matching X^* in the bipartite graph in polynomial time by the Hungarian algorithm, which runs in polynomial time (Schrijver 2003). From this matching, we construct the balanced allocation A^* defined by $A^*_i = \{j \in M \mid ((i, s), j) \in X^*\}$, which can also be obtained in polynomial time. By Lemma 3, this allocation satisfies both EF1 and fPO. Therefore, for agents with personalized bivalued valuations, a balanced allocation that is simultaneously EF1 and fPO always exists and can be computed in polynomial time. \square

Two Types Case

In this section, we consider instances in which agents are partitioned into at most two types, referred to as type-1 and type-2 agents. This structure serves as a tractable intermediate case between the setting of identical agents and fully heterogeneous populations.

We begin by revisiting the simplest case, where all agents belong to a single type; that is, every agent shares the same valuation function ($v_i = u$ for all $i \in N$). In this setting, the round-robin allocation is not only guaranteed to be EF1, but it also trivially satisfies fPO. This is because, with identical valuations, every complete allocation maximizes utilitarian social welfare.

We now turn to the case with two types. For each type $t \in \{1, 2\}$, let N_t denote the set of agents of type t , with $n_t = |N_t|$ and $n = n_1 + n_2$ the total number of agents. Without loss of generality, we assume $N_1 = \{1, \dots, n_1\}$ and $N_2 = \{n_1 + 1, \dots, n\}$. All agents of the same type share an identical valuation function: for each $i \in N_t$, $v_i = u_t$.

Our main result in this section is the following.

Theorem 4. *When the number of agent types is at most two, there always exists a balanced allocation that is both EF1 and fPO. Moreover, such an allocation can be found in polynomial time.*

To obtain such an fPO balanced allocation, we only consider weight vectors $\alpha \in \mathbb{R}_{++}^N$ such that $\alpha_i = 1$ for all $i \in N_1$ and $\alpha_i = \gamma$ for all $i \in N_2$, where $\gamma \in \mathbb{R}_{++}$.

Without loss of generality, we may assume $u_{tj} \neq u_{tj'}$ for some $t \in \{1, 2\}$ and $j, j' \in M$. Otherwise, the problem is trivial since every balanced allocation is both EF and fPO. We denote

$$\delta := \frac{\min_{t \in \{1, 2\}, j, j' \in M: u_{tj} \neq u_{tj'}} |u_{tj} - u_{tj'}|}{1 + \max_{t \in \{1, 2\}} \max_{j \in M} u_{tj}} (> 0).$$

We search γ for which the resulting balanced allocation \mathbf{A} is EF1. Recall that a balanced allocation \mathbf{A} maximizes $\sum_{i \in N} \alpha_i v_i(A_i)$ if and only if $G_{\mathbf{A}}^{(\alpha)}$ contains no negative-weight directed cycles. In what follows, we write $G_{\mathbf{A}}^{(\gamma)}$ instead of $G_{\mathbf{A}}^{(\alpha)}$ for $\alpha = (1, \dots, 1, \gamma, \dots, \gamma)$. Define

$$\mathcal{C} = \{\gamma_1, \gamma_2, \dots, \gamma_L\} \\ = \left\{ \frac{u_{1j} - u_{1j'}}{u_{2j} - u_{2j'}} \mid j, j' \in M, u_{1j} > u_{1j'}, u_{2j} > u_{2j'} \right\}$$

to be the set of *critical values* of γ , where $\delta < \gamma_1 < \gamma_2 < \dots < \gamma_L < 1/\delta$. Intuitively, a critical value γ represents a point at which a zero cycle appears in $G_{\mathbf{A}}^{(\gamma)}$. Note that $L = O(m^2)$ since there are at most $O(m^2)$ pairs of goods (j, j') . We define the intervals between these critical values as $I_1 = [\delta, \gamma_1]$, $I_2 = [\gamma_1, \gamma_2]$, \dots , $I_{L+1} = [\gamma_L, 1/\delta]$. We will denote $\gamma_0 := \delta$ and $\gamma_{L+1} := 1/\delta$. If $\mathcal{C} = \emptyset$, we define $L = 0$.

Lemma 4. *For each $\ell \in [L + 1]$, if \mathbf{A} maximizes $\sum_{i \in N_1} u_1(A_i) + \sum_{i \in N_2} \gamma u_2(A_i)$ for some $\gamma \in (\gamma_{\ell-1}, \gamma_{\ell})$, then it also maximizes the value for any $\gamma \in [\gamma_{\ell-1}, \gamma_{\ell}] = I_{\ell}$.*

For each interval I_{ℓ} , let $\mathbf{A}^{(\ell)}$ be a balanced allocation that maximizes $\sum_{i \in N_1} u_1(A_i) + \sum_{i \in N_2} \gamma u_2(A_i)$ for all $\gamma \in I_{\ell}$. Let $S^{(\ell)} := \bigcup_{i \in N_1} A_i^{(\ell)}$ and $T^{(\ell)} := \bigcup_{i \in N_2} A_i^{(\ell)}$.

For each interval I_{ℓ} and for each $\gamma \in I_{\ell}$, we determine the potentials $(\mathbf{q}^{(\gamma)}, \mathbf{p}^{(\gamma)}) \in \mathbb{R}_+^N \times \mathbb{R}_+^M$ according to the procedure described in Lemma 1. Here, we do not need to specify ℓ since the potentials do not depend on the choice of the allocation. Since agents of the same type have identical valuations and identical weights, there exists a path of length 0 in $G_{\mathbf{A}^{(\ell)}}^{(\gamma)}$ between any two agents of the same type. Thus, we have $q_i^{(\gamma)} = q_{i'}^{(\gamma)}$ for any $i, i' \in N_1$ and $q_i^{(\gamma)} = q_{i'}^{(\gamma)}$ for any $i, i' \in N_2$. We redistribute the goods in $S^{(\ell)}$ among the type-1 agents, and those in $T^{(\ell)}$ among the type-2 agents, using the round-robin procedure with respect to the price vector $\mathbf{p}^{(\gamma)}$. Let $(X_1^{(\ell, \gamma)}, \dots, X_{n_1}^{(\ell, \gamma)})$ and $(Y_1^{(\ell, \gamma)}, \dots, Y_{n_2}^{(\ell, \gamma)})$ denote the resulting allocations for type-1 and type-2 agents, respectively. We assume that the round-robin is applied in the order of indices. Then, we have the following relations.

Lemma 5. *For any $\ell \in [L + 1]$ and $\gamma \in I_{\ell}$ we have*

$$p^{(\gamma)}(X_1^{(\ell, \gamma)}) \geq \dots \geq p^{(\gamma)}(X_{n_1}^{(\ell, \gamma)}) \\ \geq \hat{p}^{(\gamma)}(X_1^{(\ell, \gamma)}) \geq \dots \geq \hat{p}^{(\gamma)}(X_{n_1}^{(\ell, \gamma)}),$$

and

$$p^{(\gamma)}(Y_1^{(\ell, \gamma)}) \geq \dots \geq p^{(\gamma)}(Y_{n_2}^{(\ell, \gamma)}) \\ \geq \hat{p}^{(\gamma)}(Y_1^{(\ell, \gamma)}) \geq \dots \geq \hat{p}^{(\gamma)}(Y_{n_2}^{(\ell, \gamma)}).$$

Let $\hat{\mathbf{A}}^{(\ell, \gamma)} = (X_1^{(\ell, \gamma)}, \dots, X_{n_1}^{(\ell, \gamma)}, Y_1^{(\ell, \gamma)}, \dots, Y_{n_2}^{(\ell, \gamma)})$. We will omit the superscripts (γ) , (ℓ) and (ℓ, γ) when they are clear from the context.

For each interval I_{ℓ} , we first check whether $\hat{\mathbf{A}}^{(\ell, \gamma)}$ is EF1 for $\gamma = \gamma_{\ell-1}$ and $\gamma = \gamma_{\ell}$. If either allocation is EF1, we immediately return it as an EF1 and fPO balanced allocation. If not, we examine the following two conditions for each $\ell \in [L + 1]$ and $\gamma \in \{\gamma_{\ell-1}, \gamma_{\ell}\}$:

- (a) $p^{(\gamma)}(X_{n_1}^{(\ell, \gamma)}) \geq \hat{p}^{(\gamma)}(Y_1^{(\ell, \gamma)})$,
- (b) $p^{(\gamma)}(Y_{n_2}^{(\ell, \gamma)}) \geq \hat{p}^{(\gamma)}(X_1^{(\ell, \gamma)})$.

We will show in Lemma 7 that at least one of these two conditions holds for any (ℓ, γ) . Furthermore, if both conditions hold for some (ℓ, γ) , then $\hat{\mathbf{A}}^{(\ell, \gamma)}$ is EF1 (Lemma 6). Since we have assumed that neither endpoint allocation is EF1, exactly one of the two conditions holds at each endpoint. Additionally, we will show in Lemma 8 that condition (a) holds at $(\ell, \gamma) = (1, \delta)$ and condition (b) holds at $(\ell, \gamma) = (L + 1, 1/\delta)$. Therefore, as we move through the intervals, at least one of the following two situations must occur:

- Case 1: For some $\ell \in [L + 1]$, condition (a) holds at $(\ell, \gamma_{\ell-1})$ and condition (b) holds at (ℓ, γ_{ℓ}) .
- Case 2: For some $\ell \in [L]$, condition (a) holds at (ℓ, γ_{ℓ}) and condition (b) holds at $(\ell + 1, \gamma_{\ell})$.

We can find such an ℓ in polynomial time by checking the conditions for each $\ell \in [L + 1]$ and $\gamma \in \{\gamma_{\ell-1}, \gamma_{\ell}\}$.

If we encounter Case 1, there exists $\gamma^* \in I_{\ell}$ such that $\hat{\mathbf{A}}^{(\ell, \gamma^*)}$ satisfies both conditions (a) and (b) (Lemma 10). Therefore, $\hat{\mathbf{A}}^{(\ell, \gamma^*)}$ is EF1 and fPO by Lemma 6. We search for such a γ^* by tracking the changes as we continuously increase γ within the interval I_{ℓ} . Since all agents of the same type share the same $q_i^{(\gamma)}$, there is a shortest path from the root r to each node in $N \cup M$ that uses at most four edges in $G_{\mathbf{A}^{(\ell)}}^{(\gamma)}$. There are at most $O(m^2)$ such paths. The length of each path is a linear function of γ , so each $-q_i^{(\gamma)}$ and $p_j^{(\gamma)}$ can be represented as the minimum of $O(m^2)$ linear functions in γ . Thus, $-q_i^{(\gamma)}$ and $p_j^{(\gamma)}$ are piecewise linear functions with at most $O(m^2)$ segments. This structure allows us to efficiently track the changes in the potentials as γ varies, enabling us to find a suitable γ^* that satisfies both conditions (a) and (b) in polynomial time.

If we encounter Case 2, we construct a sequence of allocations by successively exchanging one good from a type-1 agent with one good from a type-2 agent at each step, transitioning from $\hat{\mathbf{A}}^{(\ell, \gamma_{\ell})}$ to $\hat{\mathbf{A}}^{(\ell+1, \gamma_{\ell})}$. Specifically, at each step, we select a pair of goods (j_1, j_2) with j_1 assigned to a type-1 agent and j_2 assigned to a type-2 agent, and swap their assignments. Then, we redistribute the goods among type-1 agents and those among type-2 agents using the round-robin procedure with respect to the price vector $\mathbf{p}^{(\gamma_{\ell})}$. By repeating such single-good exchanges, we can transform $\hat{\mathbf{A}}^{(\ell, \gamma_{\ell})}$ into $\hat{\mathbf{A}}^{(\ell+1, \gamma_{\ell})}$. We will show that during this process, we reach an allocation that satisfies both conditions (a) and (b), thus yielding an EF1 and fPO balanced allocation (Lemmas 11 and 12).

We now proceed to the proofs. We first show that if $\hat{\mathbf{A}}$ satisfies the conditions (a) and (b), then it is EF1.

Lemma 6. *Let $\ell \in [L + 1]$ and $\gamma \in I_{\ell}$. If $p^{(\gamma)}(X_{n_1}^{(\ell, \gamma)}) \geq \hat{p}^{(\gamma)}(Y_1^{(\ell, \gamma)})$, then the envy from a type-1 agent to a type-2 agent can be eliminated by removing one good. Similarly, if $p^{(\gamma)}(Y_{n_2}^{(\ell, \gamma)}) \geq \hat{p}^{(\gamma)}(X_1^{(\ell, \gamma)})$, then the envy from a type-2 agent to a type-1 agent can be eliminated by removing one*

good. In addition, if both conditions hold, then the allocation is EF1.

Next, we show that at least one of the conditions (a) or (b) holds for any $\ell \in [L + 1]$ and $\gamma \in I_\ell$.

Lemma 7. *For any $\ell \in [L + 1]$ and $\gamma \in I_\ell$, at least one of $p^{(\gamma)}(X_{n_1}^{(\ell, \gamma)}) \geq \hat{p}^{(\gamma)}(Y_1^{(\ell, \gamma)})$ or $p^{(\gamma)}(Y_{n_2}^{(\ell, \gamma)}) \geq \hat{p}^{(\gamma)}(X_1^{(\ell, \gamma)})$ holds.*

We show that if $\hat{A}^{(1, \delta)}$ is not EF1, then condition (a) holds for $(\ell, \gamma) = (1, \delta)$, and if $\hat{A}^{(L+1, 1/\delta)}$ is not EF1, then condition (b) holds for $(\ell, \gamma) = (L + 1, 1/\delta)$.

Lemma 8. *If $\hat{A}^{(1, \delta)}$ is not EF1, then $p^{(\delta)}(X_{n_1}^{(1, \delta)}) \geq \hat{p}^{(\delta)}(Y_1^{(1, \delta)})$. Also, if $\hat{A}^{(L+1, 1/\delta)}$ is not EF1, then $p^{(1/\delta)}(Y_{n_2}^{(L+1, 1/\delta)}) \geq \hat{p}^{(1/\delta)}(X_1^{(L+1, 1/\delta)})$.*

We establish the existence of the desired γ^* in Case 1. As a preliminary step, we show the continuity of the prices received by each agent.

Lemma 9. *For each $\ell \in [L + 1]$, $i \in N_1$, and $i' \in N_2$, the values $p^{(\gamma)}(X_i^{(\ell, \gamma)})$, $\hat{p}^{(\gamma)}(X_i^{(\ell, \gamma)})$, $p^{(\gamma)}(Y_{i'}^{(\ell, \gamma)})$, and $\hat{p}^{(\gamma)}(Y_{i'}^{(\ell, \gamma)})$ are continuous with respect to γ over I_ℓ .*

Now we can prove the existence of γ^* in Case 1.

Lemma 10. *Take any ℓ in $[L + 1]$. Consider the case where $p^{(\gamma_{\ell-1})}(X_{n_1}^{(\ell, \gamma_{\ell-1})}) \geq \hat{p}^{(\gamma_{\ell-1})}(Y_1^{(\ell, \gamma_{\ell-1})})$ and $p^{(\gamma_\ell)}(Y_{n_2}^{(\ell, \gamma_\ell)}) \geq \hat{p}^{(\gamma_\ell)}(X_1^{(\ell, \gamma_\ell)})$. Then, there exists $\gamma^* \in I_\ell$ such that $p^{(\gamma^*)}(X_{n_1}^{(\ell, \gamma^*)}) \geq \hat{p}^{(\gamma^*)}(Y_1^{(\ell, \gamma^*)})$ and $p^{(\gamma^*)}(Y_{n_2}^{(\ell, \gamma^*)}) \geq \hat{p}^{(\gamma^*)}(X_1^{(\ell, \gamma^*)})$. Moreover, such a γ^* can be computed in polynomial time.*

We next show that, if we encounter Case 2, then we can find an EF1 and fPO balanced allocation by successively exchanging one good from a type-1 agent with one good from a type-2 agent. Let ℓ be the index chosen in the algorithm. Since both $A^{(\ell)}$ and $A^{(\ell+1)}$ maximize $\sum_{i \in N_1} u_1(A_i) + \gamma_\ell \sum_{i \in N_2} u_2(A_i)$, we have $q_i^{(\gamma_\ell)} + p_j^{(\gamma_\ell)} = u_{1j}$ for all $i \in N_1$ and $j \in S^{(\ell)} \cup S^{(\ell+1)}$, and $q_i^{(\gamma_\ell)} + p_j^{(\gamma_\ell)} = \gamma_\ell u_{2j}$ for all $i \in N_2$ and $j \in T^{(\ell)} \cup T^{(\ell+1)}$. We observe that \hat{A} obtained during the process of Case 2 is always fPO.

Lemma 11. *Fix $\ell \in [L + 1]$ and let $(\mathbf{q}, \mathbf{p}) = (\mathbf{q}^{(\gamma_\ell)}, \mathbf{p}^{(\gamma_\ell)})$. Suppose M is partitioned into S and T with $S \subseteq S^{(\ell)} \cup S^{(\ell+1)}$ and $T \subseteq T^{(\ell)} \cup T^{(\ell+1)}$. Let $\hat{A} = (X_1, \dots, X_{n_1}, Y_1, \dots, Y_{n_2})$ be the balanced allocation obtained by the round-robin procedure with respect to \mathbf{p} , assigning goods in S to type-1 agents and goods in T to type-2 agents. Then, \hat{A} is fPO.*

Next, we show that if \hat{A} does not satisfy condition (b) at some step, then \hat{A} at the next step satisfies condition (a).

Lemma 12. *Let (S, T) be a partition of M and $\mathbf{p} \in \mathbb{R}_+^N$. Let $\hat{A} = (X_1, \dots, X_{n_1}, Y_1, \dots, Y_{n_2})$ be the balanced allocation obtained by the round-robin procedure with respect to \mathbf{p} , assigning goods in S to type-1 agents and goods in T to type-2 agents. Suppose $j_1 \in S$ and $j_2 \in T$, and consider the new partition $(S', T') = (S \setminus \{j_1\} \cup \{j_2\}, T \setminus \{j_2\} \cup \{j_1\})$,*

and let $\hat{A}' = (X'_1, \dots, X'_{n_1}, Y'_1, \dots, Y'_{n_2})$ be the balanced allocation obtained by round-robin with respect to (S', T') and \mathbf{p} . Then, if $p(Y_{n_2}) < \hat{p}(X_1)$, we have $p(X'_{n_1}) \geq \hat{p}(Y'_1)$.

This lemma implies that if we encounter Case 2, we can find an EF1 and fPO balanced allocation by successively exchanging goods between type-1 and type-2 agents until conditions (a) and (b) are satisfied. This is because $\hat{A} = \hat{A}^{(\ell, \gamma_\ell)}$ at the first step satisfies condition (a), and at the end we will have $\hat{A} = \hat{A}^{(\ell+1, \gamma_{\ell+1})}$ and it satisfies condition (b). If an allocation satisfies both conditions (a) and (b), then it is EF1 by Lemma 6.

Now, we are ready to prove Theorem 4.

Proof of Theorem 4. If $\hat{A}^{(\ell, \gamma)}$ is EF1 for some $\ell \in [L + 1]$ and $\gamma \in \{\gamma_{\ell-1}, \gamma_\ell\}$, then the algorithm returns it as an EF1 and fPO balanced allocation. Otherwise, the algorithm finds an EF1 and fPO balanced allocation via Case 1 or Case 2, depending on which situation occurs. For Case 1, it finds $\gamma^* \in I_\ell$ such that $\hat{A}^{(\ell, \gamma^*)}$ satisfies both conditions (a) and (b), and returns it as an EF1 and fPO balanced allocation by Lemma 10. For Case 2, it finds an EF1 and fPO balanced allocation by Lemma 11 and Lemma 12. Thus, the algorithm always finds an EF1 and fPO balanced allocation.

For each $\ell \in [L + 1]$, the allocation $A^{(\ell)}$ can be computed in polynomial time by the Hungarian algorithm. For each $\ell \in [L + 1]$ and $\gamma \in I_\ell$, we can compute the potential $(\mathbf{q}^{(\gamma)}, \mathbf{p}^{(\gamma)})$ in polynomial time by the Bellman–Ford algorithm by Lemma 1. We can then compute the allocation $\hat{A}^{(\ell, \gamma)}$ in polynomial time since the round-robin redistribution of goods takes $O(m)$ time. Since $L = O(m^2)$ and the critical values can be computed in $O(m^2)$ time, we can check whether $\hat{A}^{(\ell, \gamma)}$ is EF1 for each $\ell \in [L + 1]$ and $\gamma \in \{\gamma_{\ell-1}, \gamma_\ell\}$ in polynomial time. For Case 1, we can find a desired γ in polynomial time by Lemma 10. Hence, the total time complexity for Case 1 is polynomial. For Case 2, we can find an EF1 and fPO balanced allocation in polynomial time since we only need to perform a sequence of single-good exchanges, and the number of such exchanges is at most $O(m)$. \square

Conclusion

In this paper, we addressed the problem of finding a fair and efficient allocation of indivisible goods under the balanced constraint, where each agent receives the same number of goods. Our main contribution is to affirmatively resolve the existence and polynomial-time computability of allocations that are simultaneously EF1 and fPO in two important cases: when agents have personalized bivalued valuations and when there are at most two types of agents. We developed novel polynomial-time algorithms for both scenarios, leveraging techniques from maximum-weight perfect matching.

A natural next step is to extend our results beyond two types of agents to the general case of three or more types, but this becomes much more complex.

Acknowledgments

This work was partially supported by the joint project of Kyoto University and Toyota Motor Corporation, titled “Advanced Mathematical Science for Mobility Society” and supported by JST ERATO Grant Number JPMJER2301, JST PRESTO Grant Number JPMJPR2122, JSPS KAKENHI Grant Numbers JP23K19956 and JP25K00137, Japan.

References

- Akrami, H.; Chaudhury, B. R.; Hoefer, M.; Mehlhorn, K.; Schmalhofer, M.; Shahkarami, G.; Varricchio, G.; Vermande, Q.; and van Wijland, E. 2022. Maximizing Nash social welfare in 2-value instances. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 4760–4767.
- 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.
- Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2021. Maximum Nash welfare and other stories about EFX. *Theoretical Computer Science*, 863: 69–85.
- 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.
- Aziz, H.; Lindsay, J.; Ritossa, A.; and Suzuki, M. 2023. Fair Allocation of Two Types of Chores. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 143–151.
- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, 557–574.
- Brainard, W. C.; and Scarf, H. E. 2005. How to compute equilibrium prices in 1891. *American Journal of Economics and Sociology*, 64(1): 57–83.
- Brams, S. J.; and Taylor, A. D. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press.
- Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D. 2016. *Handbook of computational social choice*. Cambridge University Press.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6): 1061–1103.
- Caragiannis, I.; Hansen, K. A.; and Rathi, N. 2024. On the complexity of Pareto-optimal and envy-free lotteries. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 244–252.
- 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): 12:1–32.
- Cole, R.; and Tao, Y. 2021. On the existence of Pareto efficient and envy-free allocations. *Journal of Economic Theory*, 193: 105207.
- Cookson, B.; Ebadian, S.; and Shah, N. 2025. Constrained fair and efficient allocations. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 13718–13726.
- De Keijzer, B.; Bouveret, S.; Klos, T.; and Zhang, Y. 2009. On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In *International Conference on Algorithmic Decision Theory*, 98–110. Springer.
- Devanur, N. R.; Papadimitriou, C. H.; Saberi, A.; and Vazirani, V. V. 2008. Market equilibrium via a primal–dual algorithm for a convex program. *Journal of the ACM*, 55(5): 22:1–18.
- Ebadian, S.; Peters, D.; and Shah, N. 2022. How to fairly allocate easy and difficult chores. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 372–380.
- Ehrgott, M. 2005. *Multicriteria Optimization*. Berlin; New York: Springer, 2nd edition.
- Foley, D. K. 1966. *Resource allocation and the public sector*. Yale University.
- Garg, J.; and Murhekar, A. 2023. Computing fair and efficient allocations with few utility values. *Theoretical Computer Science*, 962: 113932.
- Garg, J.; Murhekar, A.; and Qin, J. 2022. Fair and efficient allocations of chores under bivalued preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 5043–5050.
- Garg, J.; Murhekar, A.; and Qin, J. 2023. New algorithms for the fair and efficient allocation of indivisible chores. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2710–2718.
- Garg, J.; Murhekar, A.; and Qin, J. 2024. Weighted EF1 and PO allocations with few types of agents or chores. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*, 2799–2806.
- Guo, H.; Li, W.; and Deng, B. 2023. A survey on fair allocation of chores. *Mathematics*, 11(16): 3616:1–28.
- Igarashi, A.; and Meunier, F. 2025. Fair and efficient allocation of indivisible items under category constraints. *arXiv:2503.20260*.
- Kaneko, M.; and Nakamura, K. 1979. The Nash social welfare function. *Econometrica: Journal of the Econometric Society*, 423–435.
- Kawase, Y.; and Sumita, H. 2020. On the max-min fair stochastic allocation of indivisible goods. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2070–2078.
- Kawase, Y.; Sumita, H.; and Yokoi, Y. 2023. Random assignment of indivisible goods under constraints. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2792–2799.

- Kojima, F. 2009. Random assignment of multiple indivisible objects. *Mathematical Social Sciences*, 57(1): 134–142.
- 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.
- Mahara, R. 2023. Existence of EFX for two additive valuations. *Discrete Applied Mathematics*, 340: 115–122.
- Mahara, R. 2024a. Extension of additive valuations to general valuations on the existence of EFX. *Mathematics of operations research*, 49(2): 1263–1277.
- Mahara, R. 2024b. A polynomial-time algorithm for fair and efficient allocation with a fixed number of agents. *arXiv:2411.01810*.
- Mahara, R. 2025. Existence of fair and efficient allocation of indivisible chores. *arXiv:2507.09544*.
- Nash, J. F. 1950. The bargaining problem. *Econometrica*, 18(2): 155–162.
- 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.
- Orlin, J. B. 2010. Improved algorithms for computing Fisher’s market clearing prices: Computing Fisher’s market clearing prices. In *Proceedings of the forty-second ACM symposium on Theory of computing (STOC)*, 291–300.
- Schrijver, A. 2003. *Combinatorial Optimization: Polyhedra and Efficiency*. Algorithms and Combinatorics. Springer.
- Shoshan, H.; Hazan, N.; and Segal-Halevi, E. 2023. Efficient nearly-fair division with capacity constraints. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 206–214.
- Suksompong, W. 2021. Constraints in fair division. *ACM SIGecom Exchanges*, 19(2): 46–61.
- Tröbst, T.; and Vazirani, V. V. 2024. Cardinal-utility matching markets: The quest for envy-freeness, Pareto-optimality, and efficient computability. *arXiv:2402.08851*.
- Varian, H. R. 1974. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1): 63–91.
- Végh, L. A. 2012. Strongly polynomial algorithm for a class of minimum-cost flow problems with separable convex objectives. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing (STOC)*, 27–40.
- Wu, X.; Li, B.; and Gan, J. 2021. Budget-feasible maximum Nash social welfare is almost envy-free. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 465–471.