For One and All: Individual and Group Fairness in the Allocation of Indivisible Goods

Jonathan Scarlett[1] , **Nicholas Teh**[1*] and **Yair Zick**[2]
[1]Department of Computer Science, National University of Singapore
[2]College of Information and Computer Sciences, University of Massachusetts Amherst
scarlett@comp.nus.edu.sg, nicholas.teh@u.nus.edu, yzick@umass.edu

Abstract

Traditionally, research into the fair allocation of indivisible goods has focused on *individual fairness* and *group fairness*. In this paper, we explore the co-existence of individual envy-freeness (i-EF) and its group counterpart, group weighted envy-freeness (g-WEF). We propose several polynomial-time algorithms that can provably achieve i-EF and g-WEF simultaneously in various degrees of approximation under three different conditions on the agents' valuation functions: (i) when agents have identical additive valuation functions, i-EFX and g-WEF1 can be achieved simultaneously; (ii) when agents within a group share a common valuation function, an allocation satisfying both i-EF1 and g-WEF1 exists; and (iii) when agents' valuations for goods within a group differ, we show that while maintaining i-EF1, we can achieve a $\frac{1}{3}$-approximation to g-WEF1 in expectation. In addition, we introduce several novel fairness characterizations that exploit inherent group structures and their relation to individuals, such as proportional envy-freeness and group stability. We show that our algorithms can guarantee these properties approximately in polynomial time. Our results thus provide a first step into connecting individual and group fairness in the allocation of indivisible goods.

1 Introduction

Fairly allocating indivisible goods is a fundamental problem at the intersection of computer science and economics [Steinhaus, 1948; Moulin, 2004; Brandt *et al.*, 2016]. A classic problem in fair allocation involves the allocation of courses to students [Budish and Cantillon, 2012; Hoshino and Raible-Clark, 2014; Budish *et al.*, 2016]. Courses have limited capacity, and therefore slots are often allocated via a centralized mechanism. Several recent works have explored a variety of *distributive justice criteria*; these broadly fall into two categories – *individual* (e.g., that individual students are not envious of their peers), and *group* (e.g. that students of certain ethnic, gender or professional groups are treated fairly overall). While both individual and group fairness have been studied extensively in recent works, to our knowledge, there have been no works proposing mechanisms that ensure both concurrently. In this work, we seek to establish the following:

> *efficient mechanisms that concurrently ensure approximate individual and group fairness, for certain classes of agent valuation functions.*

The tension between individual and group fairness exists in a variety of allocation scenarios studied in the literature; for example, when allocating reviewers (who, in this metaphor, are the goods) to papers [Charlin and Zemel, 2013], it is important to balance the individual papers' satisfaction with their allotted, and the overall quality of reviewers assigned to tracks (e.g. ensuring that the overall reviewer quality for the Learning and Adaptation track is commensurate with that of reviewers for the Robotics track). Another example is the allocation of public resources (such as housing, or slots in public schools) [Abdulkadiroğlu and Sönmez, 2003; Benabbou *et al.*, 2018] – it is important to maintain fairness towards individual recipients, as well as groups (such as ethnic or socioeconomic groups).

In this paper, we address the question of whether individually and group envy-free allocations can co-exist when allocating indivisible goods. We present mechanisms that compute approximately individually envy-free (EF) and group weighted envy-free (WEF) allocations, where the approximation quality depends on the class of agents' valuations.

1.1 Our Contributions

We design algorithms that (approximately) reconcile individual and group envy-freeness in the allocation of indivisible goods. The strength of our results naturally depends on the generality of the valuation classes we consider, with more general valuations yielding worse approximation guarantees.

Our main technical analysis is in Section 3. In Section 3.1, we show that when agents have identical valuation functions, envy-freeness up to any good (EFX) can be achieved in conjunction with group weighted envy-freeness up to one good (WEF1). In Section 3.2, when agents within each group have common valuation functions, then envy-freeness up to one good (EF1) can be satisfied together with WEF1. In Section 3.3, when valuation functions are distinct, we show that we can obtain a constant factor $\frac{1}{3}$ approximation to WEF1 in expectation. Finally, Section 4 introduces several new notions

*Contact Author

of individual/group fairness that rely on the relationship between individuals and their group structure, such as proportional envy-freeness (PEF) in Section 4.1 and group stability in Section 4.2. We supplement these definitions by proving that our algorithms can achieve these fairness notions.

1.2 Related Work

Envy-freeness (EF) is an important individual fairness notion in indivisible goods allocation [Moulin, 1995]. The existence of approximate EF allocations in conjunction with other individual fairness notions and welfare measures (such as proportionality, pareto-optimality, maximin share) have been studied extensively [Aziz *et al.*, 2019; Caragiannis *et al.*, 2016; Budish, 2016].

[Conitzer *et al.*, 2019] and [Aziz and Rey, 2020] introduce the notion of group fairness (applied to every group of agents within the population), with both offering the "up to one good" relaxation of removing one good per player.

Several works also suggest notions of group envy-freeness [Benabbou *et al.*, 2019; Kyropoulou *et al.*, 2019]; we focus on a recently proposed notion called weighted envy-freeness (WEF) [Chakraborty *et al.*, 2020].

2 Preliminaries

In the problem of allocating indivisible goods, we are given a set of *agents* $N = \{p_1, \ldots, p_n\}$ and *goods* $G = \{g_1, \ldots, g_m\}$. Subsets of goods in G are referred to as *bundles*. Agents belong to predefined *groups* (or *types*) $\mathcal{T} = \{T_1, \ldots, T_\ell\}$. We assume that $\bigcup_{k=1}^{\ell} T_k = N$, and that no two groups intersect. Furthermore, each group T_k has a *weight* w_k, corresponding to its size, i.e. $w_k = |T_k|$. Each agent $p_i \in N$ has a non-negative *valuation function* over bundles of goods: $v_i : 2^G \to R_+$. We assume that v_i is additive, i.e, that $v_i(S) = \sum_{g \in S} v_i(\{g\})$. When all agents have the same valuation, we denote their *common valuation* by v.

In our framework, we consider the direct allocation of goods to agents, whilst taking into consideration agents' group affiliation, and in the process achieving both individual and group envy-freeness. Thus, the group allocation is not explicitly determined in the allocation process, but is induced from the individual allocations $\mathcal{A} = (A_1, ..., A_n)$ instead. We denote $Grp_k(\mathcal{A}) = \bigcup_{i:p_i \in T_k} A_i$ as the induced *group bundle* for g_k. To keep our notations simple, for any group $T_k \in \mathcal{T}$, we will let $B_k = Grp_k(\mathcal{A})$ denote this induced group bundle. We also let the group utility for T_k be $v_{T_k}(B_k) = \sum_{i:p_i \in T_k} v_i(A_i)$.

Envy-freeness was introduced by [Foley, 1967] (see also [Brandt *et al.*, 2016; Budish, 2016; Lipton *et al.*, 2004]). However, complete, envy-free allocations with indivisible goods cannot always be guaranteed (e.g. with two agents and one good, the agent without the good will always envy the other). Thus, we make use of two popular relaxations of EF.

An allocation $\mathcal{A} = (A_1, \ldots, A_n)$ is individually *envy-free up to any good* (EFX) if, for every pair of agents $p_i, p_{i'} \in N$, and for all goods $g \in A_{i'}$, $v_i(A_i) \geq v_i(A_{i'} \setminus \{g\})$. Similarly, an allocation $\mathcal{A}$ is individually *envy-free up to one good* (EF1) if, for every pair of agents $p_i, p_{i'} \in N$, there is *some* good $g \in A_{i'}$ such that $v_i(A_i) \geq v_i(A_{i'} \setminus \{g\})$.

[Chakraborty *et al.*, 2020] recently introduced an extension of the EF notion to the weighted setting, known as *weighted envy-freeness* (WEF). In this setting, agents represent groups where each group has a fixed weight. We use this notion to capture inter-group envy. Similarly, we consider two relaxed notions of WEF. The definitions below rely on the assumption that the groups' valuations of a bundle are the same regardless of how goods are internally allocated according to $\mathcal{A}$; this is a valid assumption if we assume that valuation functions of agents within a group cannot differ. In Section 3.3, we introduce an extension of the WEF notion to deal with the more general case.

An allocation $\mathcal{A} = (A_1, \ldots, A_n)$ is said to be *weighted envy-free up to one good* (WEF1) if for every two groups $T_k, T_{k'}$, there exists some good $g \in B_{k'}$ such that $\frac{v_{T_k}(B_k)}{w_k} \geq \frac{v_{T_k}(B_{k'} \setminus \{g\})}{w_{k'}}$. It is weighted envy-free up to *any* good (WEFX) if this inequality holds for *any* $g \in B_{k'}$.

Note that envy-freeness and weighted envy-freeness are referred to as EF and WEF respectively in the literature, but we refer to them as i-EF and g-WEF henceforth, to highlight that the former is an individual fairness concept, and the latter is a group fairness concept.

3 Approximate i-EF and g-WEF Allocations

In this section, we analyze the existence of approximate individual EF (i-EF) and group WEF (g-WEF) allocations.

3.1 All-Common Valuations

i-EFX allocations are known to exist within the restricted setting of all-common valuations [Plaut and Roughgarden, 2018] (i.e. when all agents have identical valuation functions). However, it turns out that g-WEFX is incompatible with approximate i-EF notions (i-EFX or i-EF1), even when all agents' valuation functions are identical. Thus, we focus on the next best group fairness property: g-WEF1. We propose the *Sequential Maximin-Iterative Weighted Round Robin* (SM-IWRR) algorithm (Algorithm 1) that can, in the all-common valuation setting, provably produce an allocation that is both i-EFX and g-WEF1 in polynomial time.

Algorithm 1: Sequential Maximin-Iterative Weighted Round Robin (SM-IWRR)

Input: set of agents N, set of goods G, set of groups T, and valuation function v
Run the SM algorithm (see Algorithm 2) with inputs N, G and v, and obtain output $\mathcal{A}' = (A'_1, \ldots, A'_n)$.
Let $A'_{\text{min}} = \arg\min_{i:p_i \in N} v(A'_i)$.
Initialize set of *representative goods*, $R = \{\}$.
for *each* $A'_i \in \mathcal{A}'$ **do**
    Create a new good r_i with $\hat{v}(r_i) = v(A'_i) - v(A'_{\text{min}})$ and add r_i to R.
end
Run the IWRR algorithm (see Algorithm 3) with inputs N, R, T and $\hat{v}$, and obtain output $\mathcal{A}$.
For each $i \in \{1, ..., n\}$, if good r_i is in A_j, $A_j \leftarrow A'_i$.
return $\mathcal{A} = (A_1, \ldots, A_n)$

Algorithm 2: Sequential Maximin (SM)

Input: set of agents N, set of goods G, and valuation function v

while *there are unassigned goods* $G_{unassigned} \subseteq G$ **do**

  Let $g \in G_{\text{unassigned}}$ be the next highest valued good.

  Pick agent $p_i \in N$ with the least-valued bundle so far (with arbitrary tie-breaking).

  $A_i \leftarrow A_i \cup \{g\}$

  $G_{\text{unassigned}} \leftarrow G_{\text{unassigned}} \setminus \{g\}$

end

return $\mathcal{A} = (A_1, \ldots, A_n)$

Algorithm 3: Iterative Weighted Round Robin (IWRR)

Input: set of agents N, set of goods G, set of groups T, and set of valuation functions $\{v_1, \ldots, v_n\}$

while *there are unassigned goods* $G_{unassigned} \subseteq G$ **do**

  Select the group $T_k \in \mathcal{T}$ with the lowest weighted bundle size, $\frac{|B_k|}{w_k}$ (with any fixed tie-breaking).

  Pick an agent $p_i \in T_k$ with the lowest $|A_i|$, where ties are broken in favour of the one that has highest marginal utility from a good in $G_{\text{unassigned}}$ (if there is still a tie, use any fixed tie-breaking).

  $A_i \leftarrow A_i \cup \{g\}$, where $g \in G_{\text{unassigned}}$ is the good p_i values the most (arbitrary tie-breaking).

  $G_{\text{unassigned}} \leftarrow G_{\text{unassigned}} \setminus \{g\}$;

end

return $\mathcal{A} = (A_1, \ldots, A_n)$

Intuitively, the SM-IWRR algorithm works by first assigning goods to agents via the SM algorithm, such that the resulting allocation is i-EFX. Then, since valuations are all-common, the algorithm takes each bundle and treats it as a single good, referred to as the *representative good*. The value of each representative good is then reduced by the value of the least-valued representative good. These representative goods are then allocated to agents via the IWRR algorithm using these values. Each agent then receives the bundle corresponding to the representative good it was allocated.

Theorem 1. *Under all-common, additive valuation functions, the SM-IWRR algorithm returns an i-EFX and g-WEF1 allocation in polynomial time.*

Proof Sketch. The proof of of the i-EFX property utilizes the following two lemmas, along with the result of Theorem 3.3 in [Chakraborty *et al.*, 2020] to show that the IWRR outputs an g-WEF1 allocation.

Lemma 2. *For all $p_i \in N$, $\hat{v}(r_i)$ is upper bounded by the value of any one good in A_i.*

Lemma 3. *For any two groups $T_k, T_{k'} \in \mathcal{T}$, let B_k and $B_{k'}$ be the bundles of representative goods allocated to group T_k and $T_{k'}$ respectively. If we have a g-WEF1 allocation of representative goods to agents, then by replacing each representative good with its corresponding bundle, the allocation remains g-WEF1.*

Lemma 2 holds because the allocation $\mathcal{A}$ is i-EFX, so for all $p_i \in N$ and any $g \in A_i$, $v(A_i \setminus \{g\}) \leq v(A_n)$. Then, as valuations are additive, $v(A_i) - v(\{g\}) \leq v(A_n)$, and hence $\hat{v}(r_i) = v(A_i) - v(A_n) \leq v(\{g\})$. Lemma 3 implies that the replacement step (of each representative good by its bundle) in the SM-IWRR algorithm preserves the "up to one good" guarantee. □

3.2 Group-Common Valuations

Next, we consider the setting where agents in different groups may have different valuations, but agents *within* any given group have the same valuations. More formally, for each good $g \in G$, and any two agents $p_i, p_{i'} \in T_k$, $v_i(g) = v_{i'}(g)$.

As the existence of i-EFX allocations in this setting is still an open question [Plaut and Roughgarden, 2018; Caragiannis *et al.*, 2019; Chaudhury *et al.*, 2020], we explore i-EF1 and its compatibility with g-WEF1.

Theorem 4. *Under group-common, additive valuation, the IWRR algorithm returns an i-EF1 and g-WEF1 allocation in polynomial time.*

Proof Sketch. The i-EF1 property follows from the round-robin nature of allocating goods to individuals, whereas the proof of g-WEF1 is the same as that of Theorem 1. □

3.3 General Valuations

Under this class of valuation functions, agents within a group can have different (additive) valuations for each good, and so a key consideration in characterising g-WEF is, for any two groups $T_k, T_{k'} \in \mathcal{T}$, the valuation of a group T_k for another group's $T_{k'}$ bundle.

We will only consider the non-allocation based definition for defining g-WEF (where valuations of a group for another group's bundle are quantified without reference to a specific internal allocation mechanism). Thus, we focus on a more general, albeit weaker, notion of g-WEF, which we term *g-WEF1 in expectation*. Intuitively, instead of assuming that items are allocated to all agents by some allocation mechanism, we consider what the *average* utility would be if we were to allocate each item to a *uniformly random* agent.

Definition 5 (g-WEF1 in expectation)**.** *An allocation $\mathcal{A} = (A_1, \ldots, A_n)$ is* weighted envy-free up to one good (g-WEF1) in expectation *if, for every two groups $T_k, T_{k'} \in \mathcal{T}$, there exists some good $g \in B_{k'}$ such that $\frac{v_{T_k}(B_k)}{w_k} \geq \frac{\overline{v}_{T_k}(B_{k'} \setminus \{g\})}{w_{k'}}$, where $\overline{v}_{T_k}(B_{k'}) = \frac{1}{w_k} \sum_{i:p_i \in T_k} \left(\sum_{g' \in B_{k'}} v_i(g') \right)$.*

We say that an allocation is g-WEF1 in expectation *up to a factor of* $\frac{1}{\gamma}$ for some constant γ if the condition in Definition 5 is replaced by $\frac{v_{T_k}(B_k)}{w_k} \geq \frac{1}{\gamma} \cdot \frac{\overline{v}_{T_k}(B_{k'} \setminus \{g\})}{w_{k'}}$. Then, we have the following result.

Theorem 6. *Under general, additive valuation, the IWRR algorithm returns an i-EF1 allocation that is g-WEF1 in expectation up to a factor of $\frac{1}{3}$.*

Proof Sketch. By analyzing each *round* (each agent gets one good in every round) at a time, for any two groups T_k and

$T_{k'}$, we can show that every time an agent in T_k selects a good g, the average value (in T_k's view) of the set of goods that can be chosen by the remaining agents from $T_{k'}$ (who have not picked any good in that round) is, in the worst case, three times that of g, thereby giving rise to the approximation factor. □

4 Additional Notions of Fairness

In addition to our studies on attaining individual and group fairness simultaneously, we introduce fairness properties that rely on the relationship between individuals and their group structure. By doing so, we seek to provide further insight into the intricacies of fairness in allocation problems involving groups of agents.

4.1 Proportionally Envy-Free (PEF) Allocations

The first property we introduce, PEF, is a hybrid (and extension) of two existing notions of fairness – individual *proportionality* (i-PROP) [Brams and Taylor, 1996] in the fair division literature, and g-WEF introduced in Section 2. A PEF allocation can be interpreted as a middle-ground between i-PROP and g-WEF. It mandates that every agent value their bundle as much as they value any other group's bundle, normalized by the group size. As usual, we introduce the "up to one good" relaxation of this notion.

Definition 7 (Proportionally envy-free up to one good)**.** *An allocation $\mathcal{A} = (A_1, \ldots, A_n)$ is* proportionally envy-free up to one good *(PEF1) if, for any agent $p_i \in N$ and group $T_k \in \mathcal{T}$, there exists $g \in B_k \setminus A_i$ such that $v_i(A_i \cup \{g\}) \geq \frac{v_i(B_k)}{w_k}$.*

It is known that i-EF1 implies i-PROP1 [Conitzer *et al.*, 2017]. The following theorem provides a connection between PEF1 and these properties.

Proposition 8. *i-EF1 implies PEF1. Additionally, when all of the group sizes (and hence weights) are equal, PEF1 implies i-PROP1.*

As such, the SM-IWRR and IWRR algorithms naturally satisfies PEF1 (and i-PROP1 in the case of equal-size groups) in addition to the guarantees already shown.

4.2 Group Stable Allocations

There are scenarios whereby agents are able to declare a one-time membership to a group, and other instances where they can opt not to join any group at all, before the allocation process begins. This is in contrast to settings whereby agents inherently belong to certain groups, such as ethnic groups in housing allocation problems [Benabbou *et al.*, 2018]. We introduce the notion of *group stability*, and consider the "up to one good" relaxation of the concept for use in our allocation problem. The significance of introducing such a notion is also exemplified in settings where the strategic reporting of membership to groups may result in undesirable effects. For instance, in the conference peer review setting, authors have the option to declare a track for the paper. This may invite strategic misreporting about the most appropriate track for the paper, in a bid to improve the chances of acceptance. We would like to introduce a stability notion that discourages this behaviour.

An allocation mechanism $\mathcal{M} : N \times G \times \mathcal{T} \times V \to |N|^G$ is a function that takes in the set of agents, goods, group memberships, and valuations (where V is the set of all agents' valuation functions), and outputs an allocation of goods to agents. We only consider deterministic allocation mechanisms, but the definitions can easily be extended to consider randomized ones as well. We first introduce the two properties that contribute to such a notion of stability.

Definition 9 (Individual rationality up to one good)**.** *An allocation $\mathcal{A} = (A_1, \ldots, A_n)$ returned by some mechanism $\mathcal{M}(N, G, \mathcal{T}, V)$ is* individually rational up to one good *(IR1) if, for every agent $p_i \in N$, there exists some good $g \in A'_i$ such that $v_i(A_i) \geq v_i(A'_i \setminus \{g\})$, where $\mathcal{M}(N, G, \mathcal{T}', V) = \mathcal{A}' = (A'_1, ..., A'_n)$, and $\mathcal{T}'$ is equivalent to $\mathcal{T}$ with the difference being p_i is now in a group on its own.*

Definition 10 (Regret-free up to one good)**.** *An allocation $\mathcal{A} = (A_1, \ldots, A_n)$ returned by some mechanism $\mathcal{M}(N, G, \mathcal{T}, V)$ is* regret-free up to one good *(RF1) if, for every agent $p_i \in N$, and every group $T_k \in \mathcal{T}$, there exists some good $g \in A'_i$ such that $v_i(A_i) \geq v_i(A_i^{(k)} \setminus \{g\})$, where $\mathcal{M}(N, G, \mathcal{T}^{(k)}, V) = \mathcal{A}^{(k)} = (A_1^{(k)}, ..., A_n^{(k)})$, and $\mathcal{T}^{(k)}$ is equivalent to $\mathcal{T}$ with the difference being p_i is now in T_k.*

Then, an allocation $\mathcal{A}$ is said to be *group stable up to one good* if all agents $p_i \in N$ are IR1 and RF1. The following theorem affirms that this property is achievable, further strengthening the fairness guarantees provided by these algorithms.

Theorem 11. *The SM-IWRR and IWRR algorithms returns an allocation that is group stable up to one good.*

Proof Sketch. By examining the round-robin and weighted round-robin nature (at the individual and group level respectively) of the algorithms, we show that switching groups can introduce envy of at most up to one good, thereby giving us the IR1 and RF1 property. □

5 Conclusions and Future Work

In this work, we show that individual fairness may come at the cost of group fairness. Group fairness is a great way to ensure diversity in outcomes [Benabbou *et al.*, 2019]. Our work attempts to reconcile diversity with individual demands. We study the existence of allocations that satisfy individual and group (weighted) envy-freeness simultaneously, and show that when agents' additive valuations are identical or at least common within groups, existing approximations of envy-freeness at both individual and group levels are compatible and achievable concurrently. In the case of general, additive valuations, in mandating i-EF1, the IWRR algorithm achieves g-WEF1 in expectation up to a factor of $\frac{1}{3}$. We also introduce two new notions of fairness – PEF and Group Stability – that exploit the group structure inherent in numerous problem domains. We show that both the SM-IWRR and IWRR algorithms achieve relaxed variants of these properties in addition to their individual and group fairness guarantees.

We thus believe that our work establishes foundations for further in-depth studies into fairness notions in order to understand and reconcile individual and group fairness properties in the problem of allocating indivisible goods.

References

[Abdulkadiroğlu and Sönmez, 2003] Atila Abdulkadiroğlu and Tayfun Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.

[Aziz and Rey, 2020] Haris Aziz and Simon Rey. Almost group envy-free allocation of indivisible goods and chores. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 39–45, 2020.

[Aziz *et al.*, 2019] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 53–59, 2019.

[Benabbou *et al.*, 2018] Nawal Benabbou, Mithun Chakraborty, Xuan-Vinh Ho, Jakub Sliwinski, and Yair Zick. Diversity constraints in public housing allocation. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 973–981, 2018.

[Benabbou *et al.*, 2019] Nawal Benabbou, Mithun Chakraborty, Edith Elkind, and Yair Zick. Fairness towards groups of agents in the allocation of indivisible items. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 95–101, 2019.

[Brams and Taylor, 1996] Steven Brams and Alan Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

[Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[Budish and Cantillon, 2012] Eric Budish and Estelle Cantillon. The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *American Economic Review*, 102(5):2237–2271, May 2012.

[Budish *et al.*, 2016] Eric Budish, Gérard P Cacho, Judd B Kessler, and Abraham Othman. Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research*, 65(2):314–336, 2016.

[Budish, 2016] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2016.

[Caragiannis *et al.*, 2016] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 305–322, 2016.

[Caragiannis *et al.*, 2019] Ioannis Caragiannis, Nick Gravin, and Xin Huang. Envy-freeness up to any item with high nash welfare: The virtue of donating items. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, pages 527–545, 2019.

[Chakraborty *et al.*, 2020] Mithun Chakraborty, Ayumi Igarashi, Edith Elkind, and Yair Zick. Weighted envy-freeness in indivisible item allocation. In *Proceedings of the 19th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 231–239, 2020.

[Charlin and Zemel, 2013] Laurent Charlin and Richard Zemel. The toronto paper matching system: An automated paper-reviewer assignment system. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 2013.

[Chaudhury *et al.*, 2020] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. Efx exists for three agents. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 1–19, 2020.

[Conitzer *et al.*, 2017] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 629–646, 2017.

[Conitzer *et al.*, 2019] Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Vaughan. Group fairness for the allocation of indivisible goods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1853–1860, 2019.

[Foley, 1967] Duncan Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.

[Hoshino and Raible-Clark, 2014] Richard Hoshino and Caleb Raible-Clark. The quest draft: An automated course allocation algorithm. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, page 2906–2913, 2014.

[Kyropoulou *et al.*, 2019] Maria Kyropoulou, Warut Suksompong, and Alexandros Voudouris. Almost envy-freeness in group resource allocation. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 400–406, 2019.

[Lipton *et al.*, 2004] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 125–131, 2004.

[Moulin, 1995] Hervé Moulin. Fair division: The no envy test. In *Cooperative Microeconomics: A Game-Theoretic Introduction*, chapter 10, pages 163–238. Princeton University Press, 1995.

[Moulin, 2004] Hervé Moulin. *Fair Division and Collective Welfare*. MIT Press, 2004.

[Plaut and Roughgarden, 2018] Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2584–2603, 2018.

[Steinhaus, 1948] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.