

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

Jonathan Scarlett¹, Nicholas Teh¹ and Yair Zick²

¹Department of Computer Science, National University of Singapore

²College of Information and Computer Sciences, University of Massachusetts Amherst

Abstract

An extensive range of recent works have explored the fair allocation of indivisible goods. Traditionally, research has focused on *individual fairness* (are individual agents satisfied with their allotted share?) and *group fairness* (are groups of agents treated fairly?). In this paper, we explore the co-existence of individual envy-freeness (*i*-EF) and its group counterpart, group weighted envy-freeness (*g*-WEF), in the allocation of indivisible goods. 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 a suitably-defined average sense. 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 [1, 2, 3, 4]. A classic problem in fair allocation involves the allocation of courses to students [5]. 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 algorithms that ensure both concurrently in the allocation of indivisible goods. In this work, we seek to establish the following:

efficient algorithms 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 [6, 7], 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 Theory 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) [8, 9] – 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 weighted envy-free allocations can co-exist when allocating indivisible goods. We present algorithms that

compute approximately individually envy-free (EF) and group weighted envy-free (WEF) allocations, where the approximation quality depends on the class of agents’ valuation functions.

One could view the WEF property as comparing the *average* (i.e. weighted) utility of an agent within a group to the average utilities of agents in other groups. However, a common flaw of this property is that it is susceptible to outliers: an agent who gets a good with an extremely high bundle value can potentially deprive other group members of valuable items. Hence, by imposing both group and individual fairness, we obtain more equitable outcomes.

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 a suitably-defined average sense (see Definition 9).

1.2 Related Work

Envy-freeness (EF) is a particularly important individual fairness notion when deciding how to fairly allocate indivisible goods [10, 11]. The existence of approximate EF allocations in conjunction with other individual fairness notions and welfare measures (such as proportionality [12], pareto-optimality [13], maximin share [14]) have been studied extensively.

Conitzer *et al.* [15] and Aziz *et al.* [16] introduce the notion of group fairness (applied to every partition of agents within the population), with both offering the “up to one good” relaxation of removing one good per player. Benabbou *et al.* [17] explore the relationship between metrics such as utilitarian social welfare in connection with group-wise fairness via an optimization approach.

Several works also suggest notions of group envy-freeness [18, 19]; we focus on a recently proposed notion called weighted envy-freeness (WEF) [20], which focuses on group fairness with pre-defined groups, allowing us to study guarantees with the removal of a single good per group. Conitzer *et al.* [15] raised this setting of pre-defined groups as an open question.

2 Preliminaries

In the problem of allocating indivisible goods, we are given a set of *agents* $N = \{p_1, \dots, p_n\}$ and *goods* $G = \{g_1, \dots, g_m\}$. Subsets of goods in G are referred to as *bundles*. Agents belong to predefined *groups* (or *types*) $\mathcal{T} = \{T_1, \dots, T_\ell\}$. 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: $v_i : 2^G \rightarrow \mathbb{R}_+$. We assume that v_i is additive [13, 21, 15], 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. For example, in the case of assigning reviewers to papers, reviewers first specify which groups they would like to belong to (by specifying their primary topic of interest), which implicitly allocates them to a group. Next, these reviewers are directly allocated papers, where their declared field is taken into consideration. Thus, the group allocation is not explicitly determined in the allocation process, but is induced from the individual allocations $\mathcal{A} = (A_1, \dots, A_n)$ instead. We denote $\text{Grp}_k(\mathcal{A}) = \bigcup_{i:p_i \in T_k} A_i$ as the induced *group bundle* for T_k . To keep our notations simple, for any group $T_k \in \mathcal{T}$, we will let $B_k = \text{Grp}_k(\mathcal{A})$ 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 [22] (see also Brandt *et al.* [4], and Lipton *et al.* [23]). However, complete, envy-free allocations with indivisible goods cannot always be guaranteed (e.g. with two agents and one good, the agent who did not get the good will always envy the other). Thus, we make use of two popular relaxations of EF.

An allocation $\mathcal{A} = (A_1, \dots, A_n)$ 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.* [20] 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 $\mathcal{A} = (A_1, \dots, A_n)$ is said to be *weighted envy-free up to one good* (WEF1) if for every two groups $T_k, T_{k'} \in \mathcal{T}$, there exists some good $g \in B_{k'}$ such that $\frac{v_{T_k}(B_k)}{w_k} \geq \frac{v_{T_k}(B_{k'} \setminus \{g\})}{w_{k'}}$. 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. An example illustrating these fairness concepts is included in the appendix.

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 [24] (i.e. when all agents have identical valuation functions). An interesting starting point is to study the existence of allocations that satisfy *i*-EFX and approximate *g*-WEF simultaneously. A natural question is whether the concept of “up to the least valued good” can be extended to the weighted setting and be achieved in conjunction with its individual counterpart. Unfortunately, we show that this is not possible, with the following proposition.

Proposition 1. *g*-WEFX is incompatible with approximate *i*-EF notions (*i*-EFX or *i*-EF1), even when all agents' valuation functions are identical.

Proof. Consider an all-common valuation setting in which we have two groups T_1 and T_2 , each with two agents $p_1, p_2 \in T_1$ and $p_3, p_4 \in T_2$. There are four goods, which all agents value equally: $v(g_1) = c$, $v(g_2) = v(g_3) = v(g_4) = 1$. Here, $c \gg 1$ is some very large constant value. Then, in order for an allocation to be i -EF1 or i -EFX, each agent must receive exactly one good. Without loss of generality, suppose $p_1 \in T_1$ gets g_1 , and the rest of the agents receive a good of value 1. In this case, T_2 has weighted envy towards T_1 , even if the lowest valued good (of value 1) is removed from T_1 's bundle. \square

Since g -WEFX is incompatible with both notions of approximate i -EF, we focus on the next best group fairness property: g -WEF1. Again, focusing on the case of all-common valuations, since i -EFX is arguably the strongest relaxation of i -EF [25], it is of interest to ask whether an i -EFX allocation can guarantee g -WEF1. However, the following proposition shows that this is not the case.

Proposition 2. *An i -EFX allocation is not necessarily g -WEF1, nor is a g -WEF1 allocation necessarily i -EFX.*

Proof. For the first part, consider the all-common valuation setting in the case of two groups T_1 and T_2 , with $p_1, p_2 \in T_1$ and $p_3, p_4 \in T_2$, and there are four goods, $v(g_1) = v(g_2) = c$, $v(g_3) = v(g_4) = 1$. Here, $c \gg 1$ is some very large constant value. Then, any allocation where each agent gets exactly one good is i -EFX. Consider the allocation where agent p_i is allocated good g_i for every $i \in \{1, \dots, 4\}$ (i.e. T_1 receives the two valuable goods, and T_2 receives the two least valued goods). This allocation is not g -WEF1, as T_2 envies T_1 even when one of the goods is removed.

For the second part, consider a setting where we have two groups T_1, T_2 ($w_1 = 2$ and $w_2 = 1$), and three goods g_1, g_2, g_3 with all-common valuations $v(g_1) = c$, $v(g_2) = v(g_3) = 1$. Again, $c \gg 1$ is some very large constant value. Note that the allocation \mathcal{A} with $(B_1, B_2) = (\{g_1\}, \{g_2, g_3\})$ is g -WEF1. However, since T_1 has two agents but only one good, one of the agents in T_1 will receive nothing. In particular, since the agent in T_2 receives two goods, the empty-handed agent will envy them even after removing any good. \square

Proposition 2 indicates that we can neither rely on existing i -EFX algorithms (that do not take into consideration groups) such as in Plaut and Roughgarden [24] and Aziz and Rey [16], nor can we directly make use of existing g -WEF1 algorithms such as in Chakraborty *et al.* [20] to achieve both properties. We therefore 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.

Intuitively, the SM-IWRR algorithm works by first assigning goods to agents via the SM algorithm¹, such that the resulting allocation is i -EFX (as shown in Theorem 4). 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 algorithm² using these values. Each agent then receives the bundle corresponding to the representative good it was allocated. We first mention an important property about the algorithm, whose proof is detailed in the appendix.

Proposition 3. *The SM-IWRR algorithm is guaranteed to terminate in $\mathcal{O}(m \log m + mn)$ steps.*

¹In fact, any i -EFX allocation algorithm could be used in place of the SM algorithm.

²Note that in this setting, since each agent gets exactly one good, in line 3, the algorithm is simply picking any agent without a good instead of one with the lowest bundle size. The algorithm is designed to be more general for use in later sections.

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

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

Output: allocation \mathcal{A}

- 1: Run the SM algorithm (see Algorithm 2) with inputs N , G and v , and obtain output $\mathcal{A}' = (A'_1, \dots, A'_n)$
 - 2: Let $A'_{\min} = \arg \min_{i:p_i \in N} v(A'_i)$ be the lowest-valued bundle in \mathcal{A}'
 - 3: Initialize set of *representative goods*, $R = \{\}$
 - 4: **for** each $A'_i \in \mathcal{A}'$ **do**
 - 5: Create a new good r_i , with $\hat{v}(r_i) = v(A'_i) - v(A'_{\min})$
 - 6: $R \leftarrow R \cup \{r_i\}$
 - 7: **end for**
 - 8: Run the IWRR algorithm (see Algorithm 3) with inputs N , R , T and \hat{v} , and obtain output $\mathcal{A} = (A_1, \dots, A_n)$.
 - 9: **for** each $A_j \in \mathcal{A}$ **do**
 - 10: **for** each $r_i \in R$ **do**
 - 11: **if** $r_i \in A_j$ **then**
 - 12: $A_j \leftarrow A'_i$
 - 13: **end if**
 - 14: **end for**
 - 15: **end for**
 - 16: **return** $\mathcal{A} = (A_1, \dots, A_n)$
-

Proof. For the SM algorithm, since valuations are all-common, finding the next favourite good of any agent can be made trivial via pre-processed sorting, which can be done in $\mathcal{O}(m \log m)$ time. There are $\mathcal{O}(m)$ iterations of the **while** loop; in each iteration, determining the next agent to be given a good takes $\mathcal{O}(\log n)$ time. Thus, the SM algorithm runs in $\mathcal{O}(m(\log m + \log n))$ time.

For the IWRR algorithm, since our setting is such that each agent gets exactly one good, finding the next favourite good of any agent can be made trivial via pre-processed sorting, which can be done in $\mathcal{O}(m \log m)$ time. There are $\mathcal{O}(n)$ iterations of the while loop; in each iteration, finding the next group takes $\mathcal{O}(\log \ell)$ time, and determining the next agent to be given a good is straightforward. Thus, the IWRR algorithm (one agent-one good variant) runs in $\mathcal{O}(m \log m + n \log \ell)$ time.

In the SM-IWRR algorithm, since each of the **for** loops takes $\mathcal{O}(n)$ steps, assuming $m > n$ and combining these with the bounds above, the time complexity of the SM-IWRR algorithm is $\mathcal{O}(m \log m)$, and is guaranteed to terminate since the input sets are finite. \square

The following theorem provides our first main result.

Theorem 4. *Under all-common, additive valuation functions, the SM-IWRR algorithm returns an i -EFX and g -WEF1 allocation.*

Proof. We first prove that the SM-IWRR algorithm outputs an i -EFX allocation. Consider the execution of the SM algorithm. Since at every round the algorithm selects an agent with the least valued bundle to be allocated a good $g \in G$, that agent cannot be envied prior to this allocation (because it has the least valued bundle from every other agent's perspective).

Hence, any envy that arises must be due to g , which is also the least valued good in that agent's bundle. This establishes the i -EFX property. When we allocate representative goods via the IWRR (which preserves i -EFX, since bundles are not modified), and map back into bundles, this amounts to a reallocation of bundles, so the resulting bundles that the agents receive are still i -EFX.

Algorithm 2 Sequential Maximin (SM)

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

Output: allocation \mathcal{A}

- 1: Initialize $A_i = \{\}$ for $i = 1, \dots, n$
 - 2: **while** there are unassigned goods $G_{\text{unassigned}} \subseteq G$ **do**
 - 3: Let $g = \arg \max_{j: g_j \in G_{\text{unassigned}}} v(g_j)$ be the highest-valued unassigned good
 - 4: Let $p_i \in N$ be the agent with the least-valued bundle A_i , where $A_i = \arg \min_{j: p_j \in N} v(A_j)$
 - 5: $A_i \leftarrow A_i \cup \{g\}$
 - 6: $G_{\text{unassigned}} \leftarrow G_{\text{unassigned}} \setminus \{g\}$
 - 7: **end while**
 - 8: **return** $\mathcal{A} = (A_1, \dots, A_n)$
-

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 $\{v_1, \dots, v_n\}$

Output: allocation \mathcal{A}

- 1: Initialize $A_i = \{\}$ for $i = 1, \dots, n$
 - 2: **while** there are unassigned goods $G_{\text{unassigned}} \subseteq G$ **do**
 - 3: Let $T_k \in \mathcal{T}$ be the group with the lowest weighted bundle size $\frac{|B_k|}{w_k}$
 - 4: Let $p_i \in T_k$ be the agent with the lowest bundle size $|A_i|$, with ties broken in favour of the one that has highest marginal utility from a good in $G_{\text{unassigned}}$ (i.e. tie-break in favour of the agent whose favourite unassigned good is of highest (subjective) valuation)
 - 5: Let $g \in G_{\text{unassigned}}$ be the good p_i values the most
 - 6: $A_i \leftarrow A_i \cup \{g\}$
 - 7: $G_{\text{unassigned}} \leftarrow G_{\text{unassigned}} \setminus \{g\}$
 - 8: **end while**
 - 9: **return** $\mathcal{A} = (A_1, \dots, A_n)$
-

The g -WEF1 property is established using the following result adapted from Theorem 3.3 of Chakraborty *et al.* [20].

Lemma 5. *Under all-common, additive valuation functions, the “Pick the Least Weight-Adjusted Frequent Picker” (PLWAFP) algorithm [20] returns a g -WEF1 allocation.*

The key point to note is that in IWRR, if we observe the group-level allocation, groups are allocated goods according to the least weight-adjusted frequency; thus, at the group level, IWRR is equivalent to PLWAFP, and in particular it ensures, by Lemma 5, that the g -WEF1 holds with respect to the \hat{v} (i.e. the value function with respect to the representative goods). However, more effort is needed to transfer this guarantee to the original valuation function v .

After the i -EFX allocation of goods to agents, consider the set of bundles $\{A_1, A_2, \dots, A_n\}$, where individual bundles are labelled such that $v(A_1) \geq v(A_2) \geq \dots \geq v(A_n)$. For all $p_i \in N$, define the *representative good value* to be $\hat{v}(r_i) = v(A_i) - v(A_n)$, where r_i is the representative good of A_i . Then, we make the following two claims:

Claim 1 For all $p_i \in N$, $\hat{v}(r_i)$ is upper bounded by the value of any one good in A_i .

Claim 2 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.

Claim 1 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 $\hat{v}(r_i) = v(A_i) - v(A_n) \leq v(\{g\})$. Claim 2 implies that the replacement step (of each representative good by its bundle) in the SM-IWRR algorithm preserves the ‘‘up to one good’’ guarantee. We proceed to prove claim 2.

Since we have a g -WEF1 allocation on representative goods, for any two groups $T_k, T_{k'} \in \mathcal{T}$, there exists a representative good $r_{\max} \in B_{k'}$, such that $\hat{v}(r_{\max}) = \max_{i': p_{i'} \in T_{k'}} \hat{v}(r_{i'})$ and the following holds:

$$\frac{\sum_{i: p_i \in T_k} \hat{v}(r_i)}{w_k} \geq \frac{\sum_{i': p_{i'} \in T_{k'}} \hat{v}(r_{i'}) - \hat{v}(r_{\max})}{w_{k'}} \quad (1)$$

By the definition of a representative good, for all i with $p_i \in T_k$, $\hat{v}(r_i) = v(A_i) - v(A_n)$, and recalling that A_n is the least-valued bundle in \mathcal{A} , the left-hand side of (1) becomes $\frac{\sum_{i: p_i \in T_k} v(A_i)}{w_k} - v(A_n)$, and the right-hand side of (1) becomes $\frac{\sum_{i': p_{i'} \in T_{k'}} v(A_{i'}) - \hat{v}(r_{\max})}{w_{k'}} - v(A_n)$. Combining these with (1), and $v(B_k) = \sum_{i: p_i \in T_k} v(A_i)$, we get

$$\frac{v(B_k)}{w_k} = \frac{\sum_{i: p_i \in T_k} v(A_i)}{w_k} \geq \frac{\sum_{i': p_{i'} \in T_{k'}} v(A_{i'}) - \hat{v}(r_{\max})}{w_{k'}}.$$

Then, since $v(B_{k'}) = \sum_{i': p_{i'} \in T_{k'}} v(A_{i'})$, it follows that

$$\frac{v(B_k)}{w_k} \geq \frac{v(B_{k'}) - \hat{v}(r_{\max})}{w_{k'}} \geq \frac{v(B_{k'}) - v(\{g_{\max}\})}{w_{k'}},$$

where the second inequality is a result of $\hat{v}(r_{\max})$ being upper bounded by the value of one good in bundle $B_{k'}$ by claim 1 (specifically, the good of maximum value $g_{\max} \in B_{k'}$).

Claim 2 ensures that the g -WEF1 property with respect to the i -EFX bundles (or representative goods) transfers to the original goods, which completes the proof of Theorem 4. \square

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. We refer to this setting as one where agents have *group-common 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 [24, 25, 26], we explore i -EF1 and its compatibility with g -WEF1. We first give a proposition showing that i -EF1 allocations are not guaranteed to be g -WEF1. The proof of Proposition 6 is similar to that of Proposition 2³, and is thus omitted.

Proposition 6. *An i -EF1 allocation is not necessarily g -WEF1, nor is a g -WEF1 allocation necessarily i -EF1, even when all agents’ valuation functions are identical.*

Much like the all-common valuation setting, we cannot rely on existing i -EF1 algorithms (that do not take into consideration groups) such as in [23, 12] to achieve both properties.

We therefore propose that the *Iterative Weighted Round Robin* (IWRR) algorithm (Algorithm 3), as introduced in the previous section, under group-common valuations, provably produces an allocation that is both i -EF1 and g -WEF1 in polynomial time.

³In fact, Proposition 6 is a generalization of Proposition 2.

Proposition 7. *The IWRR algorithm is guaranteed to terminate after $\mathcal{O}(\ell mn)$ steps.*

Proof. There are $\mathcal{O}(m)$ iterations of the **while** loop. In each iteration, finding the next group takes $\mathcal{O}(\ell)$ time, determining the next agent to allocate takes $\mathcal{O}(n)$ time, while letting the agent pick its favourite good takes $\mathcal{O}(m)$ time. Since the input sets are finite, the algorithm terminates in $\mathcal{O}(\ell mn)$ steps. \square

Next, we provide our main result for group-common valuations.

Theorem 8. *Under group-common, additive valuation, the IWRR algorithm returns an i -EF1 and g -WEF1 allocation.*

Proof. We first prove the i -EF1 property. Suppose we have a sequence of allocations of (agent, good) pairs for a single execution of the IWRR algorithm. We break this sequence into sub-sequences, which we will call *rounds*, indexed by $r \in [1, K]$, with $K = \lceil \frac{m}{n} \rceil$. At the end of round r , each agent has r goods in its bundle. While not all agents may receive a good in the final round, we can simply add dummy goods of value zero to complete the round. Let us denote by $g_{i,r}$ the good that agent p_i picked in round r . We will show that each agent $p_i \in N$, can only possibly envy any other agent $p_{i'}$ up to the good $p_{i'}$ picked in round 1 (i.e. up to good $g_{i',1}$).

Case 1: $p_{i'}$ selects after p_i in the ordering. For each round $r \in [1, K]$, since p_i selects its most valued good, $v_i(g_{i,r}) \geq v_i(g_{i',r})$. Therefore, $v_i(A_i) = \sum_{r=1}^K v_i(g_{i,r}) \geq \sum_{r=1}^K v_i(g_{i',r}) = v_i(A_{i'})$.

Case 2: $p_{i'}$ selects before p_i in the ordering. For each round $r \in [1, K-1]$, since p_i selects its most valued good, $v_i(g_{i,r}) \geq v_i(g_{i',r+1})$. Therefore, $v_i(A_i) = \sum_{r=1}^K v_i(g_{i,r}) \geq \sum_{r=1}^{K-1} v_i(g_{i,r}) \geq \sum_{r=2}^K v_i(g_{i',r}) = v_i(A_{i'} \setminus \{g_{i',1}\})$. Since this holds for all such agents $p_i, p_{i'} \in N$, the allocation is i -EF1.

The proof of g -WEF1 is the same as that of Theorem 4. \square

3.3 General Valuations

We now proceed to study the existence of individual and group fair allocations under general additive valuations. Under this class of valuation functions, the distinction between the g -WEF notion defined in Chakraborty *et al.* [20] and our setting is more apparent. Agents within a group can have different 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. This was not a concern in the previous two (All-Common and Group-Common) settings, as the valuation of the bundle to a group was the same regardless of which agent within the group actually received the good.

With this in mind, one could consider two methods of defining g -WEF – an allocation-based approach (where valuations of a group for another group's bundle depends on some internal allocation procedure, as in Bennabou *et al.* [17]), or a non-allocation based approach (where valuations of a group for another group's bundle are quantified without reference to a specific internal allocation algorithm).

We will only consider the non-allocation based definition for defining g -WEF. To do so, we introduce 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 procedure, we consider what the *average* utility would be if we were to allocate each item to a *uniformly random* agent.

Definition 9 (g -WEF1 in expectation). *An allocation $\mathcal{A} = (A_1, \dots, A_n)$ 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 $\frac{v_{T_k}(B_k)}{w_k} \geq \frac{\bar{v}_{T_k}(B_{k'} \setminus \{g\})}{w_{k'}}$, where $\bar{v}_{T_k}(B_{k'}) = \frac{1}{w_k} \sum_{i: p_i \in T_k} \left(\sum_{g' \in B_{k'}} v_i(g') \right)$.

Further relaxing this notion in a standard manner, we say that an allocation is g -WEF1 in expectation up to a factor of $\frac{1}{\gamma}$ for some constant γ when the condition in Definition 9 is replaced by $\frac{v_{T_k}(B_k)}{w_k} \geq \frac{1}{\gamma} \cdot \frac{\bar{v}_{T_k}(B_{k'} \setminus \{g\})}{w_{k'}}$. We proceed to show that the IWRR algorithm can help us achieve an approximate notion of g -WEF1 in expectation under this setting.

Theorem 10. *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. We have shown in Theorem 8 that the allocation returned by the IWRR algorithm is i -EF1. We will thus focus on proving the approximate g -WEF property.

Consider the sequence of allocations to any two groups $T_k, T_{k'} \in \mathcal{T}$ in a single execution of the IWRR algorithm. Break this sequence into sub-sequences called *rounds*, indexed by $r \in [1, K]$, with $K = \lceil \frac{|B_k| + |B_{k'}|}{w_k + w_{k'}} \rceil$. At the end of round r , each agent has r goods in its bundle. While not all agents may receive a good in the final round, we can simply add dummy goods of value zero to complete the round.

Each round is made up of $(w_k + w_{k'})$ iterations, whereby one good is selected by some agent at each iteration. Assume that ties are broken in favour of $T_{k'}$ (the other case is similar). Then, the first iteration of every round is an allocation of a good to an agent in $T_{k'}$. In particular, this agent from $T_{k'}$ that gets to make a selection on the very first iteration of the first round could have picked a good that is of arbitrarily high value to every agent in T_k – this is the good that we will drop, as part of the “up to one good” relaxation.

Case 1: $w_k < w_{k'}$. Here, we define a *shifted round* r , that consists of all iterations (except the first) of the original round r , and the first iteration of the next round $r + 1$. Note that if a round $r + 1$ doesn't exist, we can simply add dummy goods of value zero such that the setting is well-defined. In the first round, we mentioned above that the first good is dropped (let this be g_1); every other good is accounted for in some shifted round. Figure 1 illustrates a single shifted round r . We will argue the satisfiability of g -WEF1 in expectation up to a factor of $\frac{1}{3}$ in this one shifted round; the analysis then extends to all shifted rounds similarly.

Shifted Round r

...	T_k	$T_{k'}$...	$T_{k'}$	T_k	$T_{k'}$...	$T_{k'}$...	T_k	$T_{k'}$...	$T_{k'}$	T_k	...
...	V_1	V_1	...	V_1	V_1	V_1	...	V_1	...	V_1	V_1	...	V_1	V_1	...
...	V_1	V_1	...	V_1	V_2	V_2	...	V_2	...	V_2	V_2	...	V_2	V_2	...
...	V_1	V_1	...	V_1	V_2	V_2	...	V_2	...	V_3	V_3	...	V_3	V_3	...
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
...	V_1	V_1	...	V_1	V_2	V_2	...	V_2	...	V_{w_k}	V_{w_k}	...	V_{w_k}	V_{w_k}	...

Figure 1: T_k 's agent valuations for goods in $B_k \cup B_{k'}$ for a shifted round r

Let each entry (i, j) in the matrix illustrated in Figure 1 be the valuation that an agent $p_i \in T_k$ (row) has for good $g_j \in B_{k'}$ (column). Let the unshaded columns correspond to iterations whereby agents from T_k make a selection, whereas shaded columns represent the

iterations whereby an agent from $T_{k'}$ makes a selection. Without loss of generality, we can label goods and agents in a such way whereby the circled cells in unshaded columns represent the value of the good that was selected by the corresponding agent from T_k .

Then, we have that in every shifted round, we are considering a sequence of selections of the following form: T_k selects (unshaded column), $T_{k'}$ selects multiple consecutive times (shaded columns), T_k selects (unshaded column), $T_{k'}$ selects multiple consecutive times (shaded columns), etc. The structure applies due to the following lemma, whose proof is given in the appendix.

Lemma 11. *For any two groups $T_k, T_{k'} \in \mathcal{T}$, suppose $w_{k'} \geq w_k$. Then, when running the IWRR algorithm, T_k cannot make more than one consecutive selection, whereas $T_{k'}$ can only make either $\lfloor w_{k'}/w_k \rfloor$ or $\lceil w_{k'}/w_k \rceil$ consecutive selections, at any point in time (ignoring selections by groups other than T_k and $T_{k'}$).*

The values in uncircled cells indicate the maximum valuation that the agent in that row can have for the good in that column (as a result of how the algorithm works – at every iteration, the agent from T_k selects a good that gives the group maximum marginal valuation). The general observation we can make is that for every circled cell with value V_i , all the cells below it on the same column, and to its right – on the same row, or below it – cannot exceed V_i in value. This is because we labelled and arranged agents such that those who make a selection earlier is on a higher row, and goods are picked in order from left to right.

For each shifted round $r \in [1, K]$, let the set of circled cells be S_C^r and the set of shaded cells be S_B^r . In addition, define their respective sum of cell values as $u(S_C^r)$ and $u(S_B^r)$. Then, by applying this concept to all shifted rounds, we have from Definition 9 that

$$v_{T_k}(B_k) = \sum_{r=1}^K u(S_C^r), \text{ and } \bar{v}_{T_k}(B_{k'} \setminus \{g_1\}) = \frac{1}{w_k} \sum_{r=1}^K u(S_B^r).$$

Then, in order to show the g -WEF1 in expectation up to a factor of $\frac{1}{3}$ property, it is equivalent to show that

$$3w_{k'} \sum_{r=1}^K u(S_C^r) \geq \sum_{r=1}^K u(S_B^r). \quad (2)$$

In other words, if we can show that at every shifted round, the sum of shaded cells is upper bounded by $(3w_{k'} \times \text{sum of circled cells})$, the desired property follows. In the following, in every shifted round r , for every circled cell V_i^r ($i \in [1, w_k]$), define $V_{i,ROW}^r$ and $V_{i,BLK}^r$ as follows:

1. $V_{i,ROW}^r$ = sum of shaded cells in the same row and to the right of circled cell V_i^r ;
2. $V_{i,BLK}^r$ = sum of shaded cells in the shaded columns sandwiched between the unshaded columns containing the circled cells V_i^r and V_{i+1}^r , and starting from the row immediately below circled cell V_i^r .

An example is illustrated in Figure 2. Also, by the definition of \bar{u} , we have that $u(S_C^r) = \sum_{i=1}^{w_k} V_i^r$ and $u(S_B^r) = \sum_{i=1}^{w_k} (V_{i,ROW}^r + V_{i,BLK}^r)$. Moreover, since $V_{i,BLK}^r$ has a maximum of $\lceil w_{k'}/w_k \rceil$ columns (by Lemma 11) and a maximum of w_k rows (because there are w_k agents in group T_k), we have $V_{i,BLK}^r \leq V_i^r \times \lceil w_{k'}/w_k \rceil \times w_k \leq 2w_{k'} \times V_i^r$. The last inequality follows because of the assumption that $w_k \leq w_{k'}$ in this case.

Then, coupled with the fact that for every $i \in [1, w_k]$, $V_{i,ROW}^r \leq w_{k'} \times V_i^r$ (because there is a maximum of $w_{k'}$ shaded columns in any single shifted round), we obtain

$$V_{i,ROW}^r + V_{i,BLK}^r \leq 3w_{k'} \times V_i^r \quad (3)$$

...	T_k	$T_{k'}$...	$T_{k'}$	T_k	$T_{k'}$...	$T_{k'}$...	T_k	$T_{k'}$...	$T_{k'}$	$T_{k'}$...
...	V_1	V_1	...	V_1	V_1	V_1	...	$V_{1,ROW}^r$...	V_1	V_1	...	V_1	V_1	...
...	V_1	V_1	...	V_1	V_2	V_2	...	V_2	...	V_2	V_2	...	V_2	V_2	...
...	V_1	V_1	...	V_1	V_2	V_2	...	V_2	...	V_3	V_3	...	V_3	V_3	...
⋮	⋮	⋮	...	⋮	⋮	⋮	...	⋮	⋮	⋮	...	⋮	⋮	⋮	⋮
...	V_1	V_1	...	V_1	V_2	V_2	...	V_2	...	V_{w_k}	V_{w_k}	...	$V_{w_k,ROW}^r$	V_{w_k}	...

$V_{w_k,BLK}^r = \emptyset$

Figure 2: T_k 's agent valuations for goods in $B_k \cup B_{k'}$ for a shifted round r , with additional quantities indicated using shaded regions

By summing (3) on both sides over all $i \in [1, w_k]$ and shifted rounds $r \in [1, K]$,

$$\sum_{r=1}^K \sum_{i=1}^{w_k} (V_{i,ROW}^r + V_{i,BLK}^r) \leq \sum_{r=1}^K (3w_k \times V_i^r)$$

and since $u(S_C^r) = \sum_{i=1}^{w_k} V_i^r$ and $u(S_B^r) = \sum_{i=1}^{w_k} (V_{i,ROW}^r + V_{i,BLK}^r)$, we get (2) as desired.

Case 2: $w_k \geq w_{k'}$: This case is proven in a similar vein as the previous case. The difference is that instead of comparing the values of shaded cells with a single circled cell, we compare with the sum of cell values of a set of circled cells, and the bound follows. The details are given in the appendix. \square

The existence of a better approximation bound that can be obtained with comparable efficiency remains an open question, and largely depends on the allocation algorithm.

4 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 [17]. 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 $\frac{1}{3}$. Our results thus shed light on the difficulty in achieving existing notions of individual and group fairness concurrently in more complex settings. In the appendix, we include a discussion on 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.

Possible future research includes delving into allocation-based definitions of g -WEF1 to explore the existence of approximately fair allocations under that setting. It would also be interesting to consider non-sequential allocation mechanisms [27], and to understand whether better bounds exist for the case of general additive valuations. In addition to the possibility of extending the analysis to the setting with chores [12, 16], incorporating efficiency notions such as Pareto-optimality or exploring alternative valuation classes [28] are potential avenues for future work.

References

- [1] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
- [2] Ahmet Alkan, Gabrielle Demange, and David Gale. Fair allocation of indivisible goods and criteria of justice. *Econometrica*, 59(4):1023–1039, 1991.
- [3] Hervé Moulin. *Fair Division and Collective Welfare*. MIT Press, 2004.
- [4] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [5] 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.
- [6] 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.
- [7] Ivan Stelmakh, Nihar Shah, and Aarti Singh. Peerreview4all: Fair and accurate reviewer assignment in peer review. In *Proceedings of the 30th International Conference on Algorithmic Learning Theory (ALT)*, pages 828–856, 2019.
- [8] Parag Pathak. The mechanism design approach to student assignment. *Annual Review of Economics*, 3:513–536, 2011.
- [9] 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.
- [10] Hervé Moulin. Fair division: The no envy test. In *Cooperative Microeconomics: A Game-Theoretic Introduction*, chapter 10, pages 163–238. Princeton University Press, 1995.
- [11] Shinji Ohseto. Characterizations of strategy-proof and fair mechanisms for allocating indivisible goods. *Economic Theory*, 29(1):111–121, 2006.
- [12] 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.
- [13] 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.
- [14] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2016.
- [15] 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.
- [16] 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.

- [17] 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.
- [18] 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.
- [19] Martin Aleksandrov and Toby Walsh. Group envy freeness and group pareto efficiency in fair division with indivisible items. In *Proceedings of the 41st German Conference on AI (KI)*, pages 57–72, 2018.
- [20] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, 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.
- [21] Arpita Biswas and Siddharth Barman. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 91–97, 2018.
- [22] Duncan Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.
- [23] 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.
- [24] 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.
- [25] 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.
- [26] 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.
- [27] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 557–574, 2018.
- [28] Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, and Yair Zick. Finding fair and efficient allocations when valuations don’t add up. *arXiv preprint*, arXiv:2003.07060, 2020.
- [29] Haris Aziz, Hau Chan, and Bo Li. Weighted maxmin fair share allocation of indivisible chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 46–52, 2019.
- [30] Steven Brams and Alan Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- [31] 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.

A Example to Illustrate Approximate i -EF and g -WEF

Consider a setting in which we have two groups T_1 and T_2 , consisting of one and two agents respectively, with $p_1 \in T_1$ and $p_2, p_3 \in T_2$. Suppose that there are five goods g_1, g_2, g_3, g_4, g_5 , for which all agents have equal valuation: $v(g_1) = v(g_2) = v(g_3) = v(g_4) = v(g_5) = c > 0$.

Individual Envy-Freeness Suppose an allocation \mathcal{A} is such that p_1 has one good, and p_2, p_3 have two goods each. Then, p_2, p_3 clearly have no envy towards p_1 , since $v(A_2) = v(A_3) = 2c > c = v(A_1)$. This inequality also indicates that p_1 has envy for each of p_2 and p_3 . However, observe that if we were to remove *one* good from each of p_2 and p_3 's bundle, then any envy p_1 has for each of the other agents would disappear. Hence, we say that \mathcal{A} is i -EF1.

In the more general case of goods with different valuations, i -EF allows choosing any single good to remove in the above manner, e.g., the most valued one. In contrast, the stronger variant i -EFX requires the envy-free condition to hold no matter which good was removed, e.g., the least-valued one.

Group Weighted Envy-Freeness Considering the same allocation \mathcal{A} as above, T_1 's weighted bundle value is $v(B_1) = \frac{v(A_1)}{w_1} = \frac{c}{1} = c$ and T_2 's weighted bundle value is $v(B_2) = \frac{v(A_2)+v(A_3)}{w_2} = \frac{4c}{2} = 2c$. Then, clearly, T_2 has no weighted envy towards T_1 . However, the converse is not true. Observe that even if we remove any good (call it g) from T_2 's bundle, $\frac{v(B_2 \setminus \{g\})}{w_2} = \frac{3c}{2} > c = \frac{v(B_1)}{w_1}$, and there is still weighted envy by T_1 towards T_2 . Hence, this allocation \mathcal{A} is neither g -WEFX nor g -WEF1.

More generally, if the removal of some good from the bundle of B_2 gives $\frac{v(B_2 \setminus \{g\})}{w_2} \leq \frac{v(B_1)}{w_1}$, then we can say the allocation is g -WEF1. Similarly, if the same holds no matter which good is removed, then the allocation satisfies the more stringent g -WEFX property.

B Proof of Lemma 11: Consecutive Selections in the IWR algorithm

Recall that $w_{k'} \geq w_k$ by assumption. We begin by proving the first part of the lemma, that is, T_k cannot make more than one consecutive selection at any point in time.

Proof of first part

We want to show that a case of

$$\dots, T_{k'} \text{ picks}, T_k \text{ picks}, T_k \text{ picks}, \dots$$

cannot happen. Let $|B_k|, |B_{k'}|$ be the corresponding bundle sizes of T_k and $T_{k'}$ respectively before $T_{k'}$ makes such a selection as above.

Case 1: Ties broken in favour of $T_{k'}$. It must be that $\frac{|B_{k'}|}{w_{k'}} \leq \frac{|B_k|}{w_k}$. Then after $T_{k'}$ makes a selection, since we assume T_k chooses next, $\frac{|B_{k'}|+1}{w_{k'}} > \frac{|B_k|}{w_k}$. Suppose, for a contradiction, that T_k makes more than one consecutive selection. Then it must be that $\frac{|B_{k'}|+1}{w_{k'}} > \frac{|B_k|+1}{w_k}$, so that T_k can continue to make the second selection. However, this implies $\frac{|B_{k'}|}{w_{k'}} + \frac{1}{w_k} \geq \frac{|B_{k'}|+1}{w_{k'}} > \frac{|B_k|+1}{w_k}$, and cancelling $\frac{1}{w_k}$ on both ends gives us $\frac{|B_{k'}|}{w_{k'}} > \frac{|B_k|}{w_k}$, which is a contradiction.

Case 2: Ties broken in favour of T_k . It must be that $\frac{|B_{k'}|}{w_{k'}} < \frac{|B_k|}{w_k}$. Then after $T_{k'}$ makes a selection, since we assume T_k chooses next, $\frac{|B_{k'}|}{w_{k'}} \geq \frac{|B_k|+1}{w_k}$. Suppose, for a contradiction, that T_k makes more than one consecutive selection. Then it must be that $\frac{|B_{k'}|+1}{w_{k'}} \geq \frac{|B_k|+1}{w_k}$, so that $T_{k'}$ can continue to make the second selection. However, this implies $\frac{|B_{k'}|}{w_{k'}} + \frac{1}{w_k} \geq \frac{|B_{k'}|+1}{w_{k'}} \geq \frac{|B_k|+1}{w_k}$, and cancelling $\frac{1}{w_k}$ on both ends gives us $\frac{|B_{k'}|}{w_{k'}} \geq \frac{|B_k|}{w_k}$, which is a contradiction.

Proof of second part

We now prove the second part of the lemma, that is, $T_{k'}$ can only make either $\left\lfloor \frac{w_{k'}}{w_k} \right\rfloor$ or $\left\lceil \frac{w_{k'}}{w_k} \right\rceil$ consecutive selections at any point in time.

The above has shown that T_k can only make one selection at a time. Hence, we want to show that a case of

..., $T_{k'}$ picks, T_k picks, $T_{k'}$ picks H consecutive times, T_k picks, ...,

where $H < \left\lfloor \frac{w_{k'}}{w_k} \right\rfloor$ or $H > \left\lceil \frac{w_{k'}}{w_k} \right\rceil$, cannot happen. Let $|B_k|, |B_{k'}|$ be the bundle size of T_k and $T_{k'}$ before $T_{k'}$ makes the first such selection in the sequence above.

Case 1: Ties broken in favour of $T_{k'}$. It must be that

$$\frac{|B_{k'}|}{w_{k'}} \leq \frac{|B_k|}{w_k}. \quad (4)$$

Then, after $T_{k'}$ makes a selection, since we assume T_k chooses next,

$$\frac{|B_{k'}|+1}{w_{k'}} > \frac{|B_k|}{w_k}. \quad (5)$$

Thereafter, T_k gets to make a selection and it would be $T_{k'}$'s turn again, thus $\frac{|B_{k'}|+1}{w_{k'}} \leq \frac{|B_k|+1}{w_k}$. Suppose that $H < \left\lfloor \frac{w_{k'}}{w_k} \right\rfloor$. That means that after $T_{k'}$ makes $\left\lfloor \frac{w_{k'}}{w_k} \right\rfloor - 1$ consecutive selections, $\frac{|B_{k'}|+1+\left\lfloor \frac{w_{k'}}{w_k} \right\rfloor-1}{w_{k'}} > \frac{|B_k|+1}{w_k}$, so that $T_{k'}$ cannot make the $\left(\left\lfloor \frac{w_{k'}}{w_k} \right\rfloor\right)^{\text{th}}$ selection, because it's T_k 's turn. However, this implies

$$\begin{aligned} \frac{|B_{k'}|}{w_{k'}} &> \frac{|B_k|+1}{w_k} - \frac{\left\lfloor \frac{w_{k'}}{w_k} \right\rfloor}{w_{k'}} \geq \frac{|B_k|+1}{w_k} - \frac{\left(\frac{w_{k'}}{w_k}\right)}{w_{k'}} \\ &= \frac{|B_k|+1}{w_k} - \frac{1}{w_k} = \frac{|B_k|}{w_k}, \end{aligned}$$

which contradicts (4). Now suppose that $H > \left\lceil \frac{w_{k'}}{w_k} \right\rceil$. That means that $T_{k'}$ makes $\left\lceil \frac{w_{k'}}{w_k} \right\rceil$ consecutive selections and yet $\frac{|B_{k'}|+1+\left\lceil \frac{w_{k'}}{w_k} \right\rceil}{w_{k'}} \leq \frac{|B_k|+1}{w_k}$, so that $T_{k'}$ can continue to make the $\left(\left\lceil \frac{w_{k'}}{w_k} \right\rceil + 1\right)^{\text{th}}$ selection. However, this implies

$$\begin{aligned} \frac{|B_{k'}|+1}{w_{k'}} &\leq \frac{|B_k|+1}{w_k} - \frac{\left\lceil \frac{w_{k'}}{w_k} \right\rceil}{w_{k'}} \leq \frac{|B_k|+1}{w_k} - \frac{\left(\frac{w_{k'}}{w_k}\right)}{w_{k'}} \\ &= \frac{|B_k|+1}{w_k} - \frac{1}{w_k} = \frac{|B_k|}{w_k}, \end{aligned}$$

which contradicts (5).

Case 2: Ties broken in favour of T_k . It must be that

$$\frac{|B_{k'}|}{w_{k'}} < \frac{|B_k|}{w_k}. \quad (6)$$

Then, after $T_{k'}$ makes a selection, since we assume T_k chooses next,

$$\frac{|B_{k'}| + 1}{w_{k'}} \geq \frac{|B_k|}{w_k}. \quad (7)$$

Thereafter, T_k gets to make a selection and it would be $T_{k'}$'s turn again, thus $\frac{|B_{k'}|+1}{w_{k'}} < \frac{|B_k|+1}{w_k}$. Suppose that $H < \left\lfloor \frac{w_k}{w_{k'}} \right\rfloor$. That means that after $T_{k'}$ makes $\left\lfloor \frac{w_k}{w_{k'}} \right\rfloor - 1$ consecutive selections, $\frac{|B_{k'}|+1+\left\lfloor \frac{w_k}{w_{k'}} \right\rfloor-1}{w_{k'}} \geq \frac{|B_k|+1}{w_k}$, so that $T_{k'}$ cannot make the $\left(\left\lfloor \frac{w_k}{w_{k'}} \right\rfloor\right)^{\text{th}}$ selection, because it's T_k 's turn. However, this implies

$$\begin{aligned} \frac{|B_{k'}|}{w_{k'}} &\geq \frac{|B_k| + 1}{w_k} - \frac{\left\lfloor \frac{w_k}{w_{k'}} \right\rfloor}{w_{k'}} \geq \frac{|B_k| + 1}{w_k} - \frac{\left(\frac{w_{k'}}{w_k}\right)}{w_{k'}} \\ &= \frac{|B_k| + 1}{w_k} - \frac{1}{w_k} = \frac{|B_k|}{w_k}, \end{aligned}$$

which contradicts (6). Now suppose that $H > \left\lfloor \frac{w_k}{w_{k'}} \right\rfloor$. That means that $T_{k'}$ makes $\left\lfloor \frac{w_k}{w_{k'}} \right\rfloor$ consecutive selections and yet $\frac{|B_{k'}|+1+\left\lfloor \frac{w_k}{w_{k'}} \right\rfloor}{w_{k'}} < \frac{|B_k|+1}{w_k}$, so that $T_{k'}$ can continue to make the $\left(\left\lfloor \frac{w_k}{w_{k'}} \right\rfloor + 1\right)^{\text{th}}$ selection. However, this implies

$$\frac{|B_{k'}| + 1}{w_{k'}} < \frac{|B_k| + 1}{w_k} - \frac{\left\lfloor \frac{w_k}{w_{k'}} \right\rfloor}{w_{k'}} \leq \frac{|B_k| + 1}{w_k} - \frac{\left(\frac{w_{k'}}{w_k}\right)}{w_{k'}} = \frac{|B_k|}{w_k},$$

which contradicts (7).

C Proof of Theorem 10 (Case 2): i -EF1 and g -WEF1 property of the IWRR algorithm

Recall that we focused our analysis on any two groups $T_k, T_{k'} \in \mathcal{T}$ to show the g -WEF1 property of the IWRR algorithm. Case 1 handled the scenario where $w_k < w_{k'}$, we will handle the scenario where $w_k \geq w_{k'}$ in this case.

Similar to case 1, we define a *shifted round* r , that consists of all iterations (except the first) of the original round r , and the first iteration of the next round $r + 1$. Note that if a round $r + 1$ doesn't exist, we can simply add dummy goods of value zero such that the setting is well-defined. In the first round, we mentioned that the first good is dropped (let this be g_1); every other good is accounted for in some shifted round.

We make use of Figure 3 to aid in our argument – it illustrates a single shifted round r . We will argue the satisfiability of g -WEF1 in expectation up to a factor of $\frac{1}{3}$ in this one shifted round; the analysis then extends to all shifted rounds similarly.

Let each entry (i, j) in the matrix illustrated in Figure 3 above be the valuation that an agent $p_i \in T_k$ (row) has for good $g_j \in B_{k'}$ (column). Let the unshaded columns correspond to iterations whereby agents from T_k make a selection, whereas shaded columns represent

		S_1^r				S_2^r				S_d^r				
...	T_k	...	T_k	$T_{k'}$	T_k	...	T_k	$T_{k'}$...	T_k	...	T_k	$T_{k'}$...
...	V_{11}	...	V_{11}	V_{11}	V_{11}	...	V_{11}	V_{11}	...	V_{11}	...	V_{11}	V_{11}	...
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
...	V_{11}	...	$V_{1\alpha_1}$	$V_{1\alpha_1}$	$V_{1\alpha_1}$...	$V_{1\alpha_1}$	$V_{1\alpha_1}$...	$V_{1\alpha_1}$...	$V_{1\alpha_1}$	$V_{1\alpha_1}$...
...	V_{11}	...	$V_{1\alpha_1}$	$V_{1\alpha_1}$	V_{21}	...	V_{21}	V_{21}	...	V_{21}	...	V_{21}	V_{21}	...
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
...	V_{11}	...	$V_{1\alpha_1}$	$V_{1\alpha_1}$	V_{21}	...	$V_{2\alpha_2}$	$V_{2\alpha_2}$...	$V_{2\alpha_2}$...	$V_{2\alpha_2}$	$V_{2\alpha_2}$...
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
...	V_{11}	...	$V_{1\alpha_1}$	$V_{1\alpha_1}$	V_{21}	...	$V_{2\alpha_2}$	$V_{2\alpha_2}$...	V_{d1}	...	V_{d1}	V_{d1}	...
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
...	V_{11}	...	$V_{1\alpha_1}$	$V_{1\alpha_1}$	V_{21}	...	$V_{2\alpha_2}$	$V_{2\alpha_2}$...	V_{d1}	...	$V_{d\alpha_d}$	$V_{d\alpha_d}$...

Figure 3: T_k 's agent valuations for goods in $B_k \cup B_{k'}$ for a shifted round r

the iterations whereby an agent from $T_{k'}$ makes a selection. Without loss of generality, we can label goods and agents in a such way whereby the circled cells belonging to unshaded columns represent the value of the good that was selected by the corresponding agent from T_k .

Then, we have that in every shifted round, we are considering a sequence of selections of the following form: T_k selects multiple consecutive times (unshaded columns), $T_{k'}$ selects (shaded column), T_k selects multiple consecutive times (unshaded columns), $T_{k'}$ selects (shaded column), etc. This structure applies due to Lemma 11.

The values in uncircled cells indicate the maximum valuation that the agent in that row can have for the good in that column (as a result of how the algorithm works – at every iteration, the agent from T_k selects a good that gives the group maximum marginal valuation). The general observation we can make is that for every circled cell with value V , all the cells below it on the same column, and to its right – on the same row, or below it – cannot exceed V in value. This is because we labelled and arranged agents such that those who make a selection earlier is on a higher row, and goods are picked in order from left to right.

For each shifted round $r \in [1, K]$, let the set of circled cells be S_C^r and the set of shaded cells be S_B^r . In addition, define their respective sum of cell values as $u(S_C^r)$ and $u(S_B^r)$. Then, by applying this concept to all shifted rounds, we have from Definition 9 that

$$v_{T_k}(B_k) = \sum_{r=1}^K u(S_C^r), \text{ and } \bar{v}_{T_k}(B_{k'} \setminus \{g_1\}) = \frac{1}{w_k} \sum_{r=1}^K u(S_B^r).$$

Then, in order to show

$$\frac{v_{T_k}(B_k)}{w_k} \geq \frac{\bar{v}_{T_k}(B_{k'} \setminus \{g_1\})}{3w_{k'}}, \quad (8)$$

it is equivalent to show that

$$3w_{k'} \sum_{r=1}^K u(S_C^r) \geq \sum_{r=1}^K u(S_B^r). \quad (9)$$

In other words, if we can show that at every shifted round, the sum of shaded cells is upper bounded by $(3w_{k'} \times \text{sum of circled cells})$, the property in (8) follows.

The difference between this case and the previous is that now, there can be multiple consecutive unshaded columns, and hence circled cells are no longer isolated – so our analysis will not simply be considering for every single circled cell, but for every set of circled cells in consecutive unshaded columns. For instance, as illustrated in Figure 3, in shifted round r , the first set is $S_1^r = \{V_{11}, \dots, V_{1\alpha_1}\}$, second set is $S_2^r = \{V_{21}, \dots, V_{2\alpha_2}\}$, etc. Let there be d sets, and let $u(S_i^r)$ be the sum of values in a set (i.e. sum of circled cells' values in a set S_i^r). For all $i \in [1, d]$, α_i denotes the number of circled cells in set S_i^r .

In the following, in every shifted round r , for every set of circled cells S_i^r ($i \in [1, d]$), define $S_{i,BLK}^r$ and $S_{i,COL}^r$ as follows:

1. $S_{i,BLK}^r$ = set of shaded cells in the same row and to the right of every circled cell in S_i^r (in shifted round r);
2. $S_{i,COL}^r$ = set of shaded cells in the first shaded column to the right of S_i^r , and starting from the row immediately below circled cell $V_{i\alpha_i}$.

An example is illustrated in Figure 4.

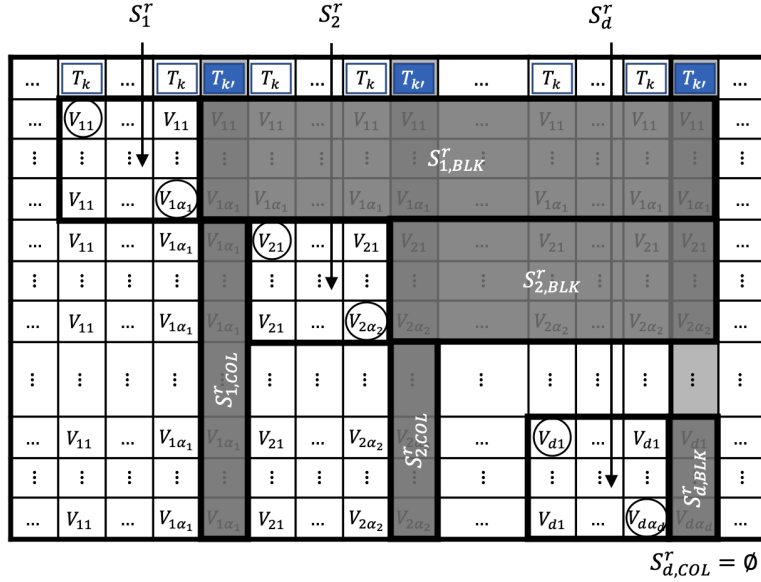


Figure 4: T_k 's agent valuations for goods in $B_k \cup B_{k'}$ for a shifted round r , with additional quantities indicated using shaded regions

By the definition of u , we have the following:

1. $u(S_C^r) = \sum_{i=1}^d u(S_i^r)$;
2. $u(S_B^r) = \sum_{i=1}^d (u(S_{i,BLK}^r) + u(S_{i,COL}^r))$.

Moreover, for any shifted round r , since $S_i^r = \{V_{i1}, \dots, V_{i\alpha_i}\}$, there are α_i circled cells forming S_i . Then,

$$u(S_{i,BLK}^r) \leq w_{k'} \times (V_{i1} + \dots + V_{i\alpha_i}) = w_{k'} \times u(S_i^r), \quad (10)$$

because there are a maximum of $w_{k'}$ shaded columns in a single shifted round. Next, for each $i \in [1, d]$, $\alpha_i \geq \left\lfloor \frac{w_k}{w_{k'}} \right\rfloor$ (by Lemma 11). Thus,

$$\begin{aligned}
u(S_{i,COL}^r) &\leq V_{i\alpha_i}(w_k - \alpha_i) \\
&\leq V_{i\alpha_i} \left(w_k - \left\lfloor \frac{w_k}{w_{k'}} \right\rfloor \right) \\
&\leq V_{i\alpha_i} \left(\frac{w_k(w_{k'} - 1)}{w_{k'}} + 1 \right) \\
&\leq V_{i\alpha_i} \left(\left(\left\lfloor \frac{w_k}{w_{k'}} \right\rfloor + 1 \right) (w_{k'} - 1) + 1 \right) \\
&\leq V_{i\alpha_i} ((\alpha_i + 1)(w_{k'} - 1) + 1) \\
&= V_{i\alpha_i} (\alpha_i w_{k'} + w_{k'} - \alpha_i) \\
&\leq V_{i\alpha_i} (2\alpha_i w_{k'}) \\
&\leq 2w_{k'} \times u(S_i^r),
\end{aligned} \tag{11}$$

where the last inequality is derived from the fact that $V_{i1} \geq V_{i2} \geq \dots \geq V_{i\alpha_i}$ and $u(S_i^r) = \sum_{j=1}^{\alpha_i} V_{ij}$, implying $\alpha_i V_{i\alpha_i} \leq u(S_i^r)$.

Then, combining (10) and (11), we obtain

$$u(S_{i,BLK}^r) + u(S_{i,COL}^r) \leq 3w_{k'} \times u(S_i^r) \tag{12}$$

By summing (12) on both sides over all $i \in [1, w_k]$ and shifted rounds $r \in [1, K]$,

$$\sum_{r=1}^K \sum_{i=1}^{w_k} (u(S_{i,BLK}^r) + u(S_{i,COL}^r)) \leq \sum_{r=1}^K (3w_{k'} \times u(S_i^r)) \tag{13}$$

and since $u(S_C^r) = \sum_{i=1}^d u(S_i^r)$ and $u(S_B^r) = \sum_{i=1}^d (u(S_{i,BLK}^r) + u(S_{i,COL}^r))$, it follows that

$$\sum_{r=1}^K u(S_B^r) \leq 3w_{k'} \sum_{r=1}^K u(S_C^r),$$

which gives (9) as desired.

D Discussion on Additional Notions of Fairness

Traditional notions of individual fairness have recently seen their group counterparts introduced [29, 20]. However, when we look at allocating to individuals in groups, new opportunities emerge for us to characterise fairness notions specific to this setting. 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.

D.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) [30] in the fair division literature, and *g*-WEF introduced in Section 2. First, we restate the definition of a relaxed version of *i*-PROP.

Definition 12 (Proportional up to one good). *An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is individually proportional up to one good (i -PROP1) if, for any agent $p_i \in N$, there exists a good $g \in G \setminus A_i$ such that $v_i(A_i \cup \{g\}) \geq \frac{v_i(G)}{n}$.*

Next, we proceed to define PEF. 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.⁴

Definition 13 (Proportionally envy-free up to one good). *An allocation $\mathcal{A} = (A_1, \dots, A_n)$ 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 [31]. Thus, a natural follow-up question would be whether i -EF1 implies PEF1, and it turns out that it is true. In fact, there is also a connection between PEF1 and i -PROP1, as the following proposition postulates.

Proposition 14. *i -EF1 implies PEF1. Additionally, when all of the group sizes (and hence weights) are equal, PEF1 implies i -PROP1.*

Proof. For the first part, we start by noting that from the definition of i -EF1, for all $p_{i'} \in T_k$, there exists some $g_{i'} \in A_{i'}$ such that $v_i(A_i) \geq v_i(A_{i'}) - v_i(g_{i'})$. Summing both sides over agents $p_{i'} \in T_k$, we obtain

$$w_k v_i(A_i) \geq \sum_{i': p_{i'} \in T_k} [v_i(A_{i'}) - v_i(\{g_{i'}\})]. \quad (14)$$

The right-hand side can be simplified as follows, with g_{\max} being the maximally valued good by p_i in the bundle $B_k \setminus A_i$:

$$\begin{aligned} & \sum_{i': p_{i'} \in T_k} v_i(A_{i'}) - \sum_{i': p_{i'} \in T_k} v_i(\{g_{i'}\}) \\ &= v_i(B_k) - \sum_{i': p_{i'} \in T_k} v_i(\{g_{i'}\}) \\ &\geq v_i(B_k) - w_k v_i(g_{\max}) \end{aligned}$$

Combining this with (14), we get

$$\begin{aligned} v_i(A_i) &\geq \frac{v_i(B_k)}{w_k} - v_i(\{g_{\max}\}) \\ \Rightarrow v_i(A_i \cup \{g_{\max}\}) &\geq \frac{v_i(B_k)}{w_k}. \end{aligned} \quad (15)$$

Thus, PEF1 is satisfied.

We now prove the second part of the proposition. Since the weights are equal, in this part, we write it as w . From the definition of PEF1, and summing both sides of (15) over all groups $T_k \in \mathcal{T}$ (recall that there are ℓ groups in total), we obtain

$$\ell \times v_i(A_i \cup \{g\}) \geq \sum_{k: T_k \in \mathcal{T}} \frac{v_i(B_k)}{w}. \quad (16)$$

⁴We adopt a similar relaxation to the traditional i -PROP property in the literature, having the good g added to the left-hand side of the equation rather than removing from the right-hand side.

Hence, since $w\ell = n$, we have that

$$\begin{aligned} v_i(A_i \cup \{g\}) &\geq \sum_{k:T_k \in \mathcal{T}} \frac{v_i(B_k)}{w\ell} \\ &= \frac{\sum_{k:T_k \in \mathcal{T}} v_i(B_k)}{w\ell} = \frac{v_i(G)}{n}. \end{aligned} \tag{17}$$

for some $g \in G \setminus A_i$. Thus, i -PROP1 is satisfied. \square

As such, the SM-IWRR and IWRR algorithms proposed in section 3 naturally satisfies PEF1 (and i -PROP1 in the case of equal-size groups) in addition to the guarantees already shown.

D.2 Approximately Group Stable Allocations

The second property that we introduce is *group stability*. 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 [9]. We introduce the notion of group stability, and consider a relaxation of the concept, which we will term group ϵ -*stability* 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 notion that discourages this behaviour.

One key thing to note here is that the notion of *stability* here is implicit, in the sense that the agent will not be able to change their group membership after being in a group. However, an allocation satisfying such a property would have more merit as agents can be assured that they could not have been much better off by misreporting their preferences.

An allocation mechanism $\mathcal{M} : N \times G \times \mathcal{T} \times V \rightarrow |N|^G$ 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 now formally introduce the relaxed notion of the group stability property.

In fact, this relaxed notion is essentially an ‘‘up to one good’’ variation, and in many real-world settings, one could argue that a single good has little utility, thereby giving rise to an *almost* stable property as defined below.

Definition 15 (Group ϵ -stability). *An allocation $\mathcal{A} = (A_1, \dots, A_n)$ returned by some mechanism $\mathcal{M}(N, G, \mathcal{T}, V)$ is group ϵ -stable if the following conditions hold:*

- (i) *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 $\mathcal{M}(N, G, \mathcal{T}', V) = \mathcal{A}' = (A'_1, \dots, A'_n)$, and \mathcal{T}' is equivalent to \mathcal{T} with the difference being that p_i is now in a group on its own.

- (ii) *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 $\mathcal{M}(N, G, \mathcal{T}^{(k)}, V) = \mathcal{A}^{(k)} = (A_1^{(k)}, \dots, A_n^{(k)})$, and $\mathcal{T}^{(k)}$ is defined by taking \mathcal{T} and moving agent p_i to group T_k .

Intuitively, (i) caters for the case whereby agents are able to choose not to join a group prior to the allocation process. Then, the property guarantees that they will not have “regretted” their decision. (ii) is similar in this regard, but the “no-regret” is with respect to reporting membership to other groups instead.

The next question that arises is whether such a property is achievable. We give two theorems that provide a positive answer.

Theorem 16. *The IWRR algorithm returns an allocation that is group ϵ -stable.*

Proof. We first prove that IWRR returns an allocation that satisfies (ii) of the group ϵ -stability property. Suppose that some agent $p_i \in N$ switches group from T_k to $T_{k'}$. By the individual round-robin nature of the IWRR, every agent gets one good per round, regardless of their group. Let $g_i^{(d)}$ and $g_i'^{(d)}$ be the d^{th} good (i.e, in round r_d) that p_i received as part of being in T_k and $T_{k'}$ respectively. Let there be a total of K rounds. Then, since each agent selects their favourite good at every round, we must have that $v_i(\{g_i^{(d)}\}) \geq v_i(\{g_i'^{(d+1)}\})$ for all $d = 1, \dots, K - 1$. Thus, for all agents $p_i \in N$ belonging to group T_k , where A_i is the bundle received by being in the group T_k , and A_i' is the bundle received by declaring membership to any other group $T_{k'}$, we have that

$$\sum_{d=1}^{K-1} v_i(\{g_i^{(d)}\}) + v_i(\{g_i^{(K)}\}) \geq \sum_{d=2}^K v_i(\{g_i'^{(d)}\})$$

where the left-hand side is equal to $v_i(A_i)$ and the right-hand side is equal to $v_i(A_i' \setminus \{g_i'^{(1)}\})$, obtained by a relabelling of the index d . The proof that IWRR returns an allocation that satisfies (i) is similar to that of (ii), where we consider $T_{k'}$ to be an empty group initially, and if p_i joins, then it becomes a singleton. The result follows. \square

Given that the IWRR algorithm is group stable up to one good, we can say the same about the SM-IWRR algorithm, with the proof being a simple combination of Theorem 16 and the representative good idea.

Theorem 17. *The SM-IWRR algorithm returns an allocation that is group ϵ -stable.*

In summary, we have shown that the SM-IWRR and IWRR algorithms also have group stability guarantees, further strengthening the fairness guarantees provided by these algorithms.