ECAI 2023
K. Gal et al. (Eds.)

doi:10.3233/FAIA230362

On the Notion of Envy Among Groups of Agents in House Allocation Problems

Nathanaël Gross–Humbert, Nawal Benabbou, Aurélie Beynier and Nicolas Maudet

LIP6 - CNRS, Sorbonne Université
4 place Jussieu, F-75005 Paris, France
ORCiD ID: Nathanaël Gross–Humbert https://orcid.org/0000-0003-3652-0393,
Nawal Benabbou https://orcid.org/0000-0002-4589-4162, Aurélie Beynier https://orcid.org/0000-0002-3812-6755,
Nicolas Maudet https://orcid.org/0000-0002-4232-069X

Abstract. Envy-freeness is one of the prominent fairness notions in multiagent resource allocation but it has been mainly studied from an individual point of view. When the agents are partitioned into groups, fairness between groups is desirable. Several notions of group envy-freeness have been proposed over the last few years in the domain of fair division. In this paper we show that when groups may have different sizes and each agent gets at most one item, existing group envy-freeness notions fail to satisfy some desirable axioms. This motivates us to propose an original notion of degree of envy-freeness among groups, based on the counterfactual comparison of subgroups of the same size. While this notion is computationally demanding, we show that it can be efficiently approximated thanks to an adapted sampling method, showing that our approach is of practical relevance.

1 Introduction

Fair multiagent resource allocation has become a widely studied topic. However, the vast majority of works on fair division have focused on individual fairness. While envy-freeness is a prominent notion, most works dealing with the absence of envy study the individual envy of one agent towards another ([9, 14, 8] among many others). In some applications, however, items are allocated to groups of agents who then divide the bundle of items they get among the members of the group according to a group-specific mechanism; for example, this is the case when allocating food between charities which redistribute items among their subscribers. In such cases, the allocation process first allocates items to groups of agents and not to a specific agent. It is then natural to consider the fairness between groups instead of individual fairness. In other contexts where agents are partitioned into groups, the interest of the group may prevail over the interest of individuals from the point of view of the resource allocation system or from the agents' point of view themselves. In these cases, the system seeks fairness between groups rather than fairness between individuals. This is even more relevant for envy-freeness in house allocation problems since, in this context, envy-freeness only holds when all the agents have a different preferred item and each agent obtains her most preferred item. If the agents are partitioned into groups, group-envy freeness may be less demanding and more suited to the context. Consider for instance an academic department assigning offices to members of different research groups; the department usually cares not just about the overall score/utility of the allocation but also about its fairness with respect to the research groups into which the faculty is divided. In a medical context, a government may have to assign medical resources (treatments, equipment) to hospitals of different sizes which then reallocate these resources to their patients or to medical departments. Here, each hospital is considered as a group and we care about fairness with respect to hospitals.

Several notions of group envy-freeness have been proposed over the last few years in the domain of fair division. When groups may have different sizes, the definition of group envy-freeness is not obvious. Indeed, the sizes of the groups have to be considered to compare the utility of the groups. Moreover, in house allocation problems, some resources may be useless to a group since each agent gets at most one item. Hence, allocating a large bundle of items to a small group may not be beneficial as a large amount of items will be wasted. In this paper we show that when groups may have different sizes and each agent gets at most one item, existing group envy-freeness notions fail to satisfy some desirable axioms. We thus propose an original notion of degree of envy among groups, based on the counterfactual comparison of subgroups of the same size.

Related work. Approaches to fair division among groups can be distinguished depending on whether they assume that groups are predefined or not, whether items are shared or assigned to individual agents within the group, and whether preferences can be aggregated at the group level or remain only individual. In this paper, we assume that groups pre-exist in the model, and that items are assigned to groups, which redistribute them in a house allocation fashion (at most one item per individual agent), with the objective to maximize the sum of agents' utilities within each group. Hence, the valuation function measuring the satisfaction of the group is non-additive. In what follows, we discuss approaches which make different assumptions, deferring the discussion of closer related work to Section 3.

Group envy-freeness has been studied by Conitzer et al. [7], based on the notion of pairwise envy between groups, which holds between G_1 and G_2 when G_1 can reallocate the items of G_2 in such a way that it Pareto-dominates its current allocation. An allocation is said to be *group fair* when no group of agents (of any size) envies any other group, under this definition. In general, envy-freeness is too difficult to satisfy in most instances so, the paper focuses on the more useful notion of *envy-freeness up to one good*. Their approach assumes additive preferences, and groups are not pre-determined.

Manurangsi et al. [15] use a model where each item is shared between the agents of a group, and group envy-freeness (up to c goods) is only reached if every agent in a group is not envious (up to c goods). A similar definition is given to proportionality up to c goods. The authors then explore what kind of value can be expected for c with regards to the size on the instance, using discrepancy theory.

In a similar vein, Kyroupoulou et al. [13] assume both predetermined and variable groups, and establish several results regarding the existence of envy-freeness (up to one good) in the case of binary valuations and in more general settings. In their setting, items are also shared within the group, and their notion of group envy does not aggregate utilities at the group level.

Aziz et al. [2] do not use a model with fixed groups but are rather interested in extending the definition of envy-freeness to account cases where agents might spontaneously form groups, and as such allocation should strive to satisfy every possible subset of agents.

Finally, Aleksandrov and Walsh [1] assume a model based on arithmetic-mean utilities for groups. Unlike us, they remain agnostic regarding the mechanism of allocation, meaning that a group may be envious because of a poor assignment of items within its members.

Outline of the paper. We start by presenting our model in Section 2. Section 3 then discuss relevant axioms for group envy-freeness, and show that approaches from the literature do not satisfy all of them. Difficulties occur for the notion of group envy when groups of unequal size must be considered, since in that case it is not obvious what thought experiment the small group (for instance) should undertake when contemplating the large group. This motivates us to propose a new notion of group envy, based on the counterfactual comparison of groups of equal size, obtained by considering subgroups of the larger group (Section 4). This notion satisfies our axioms, but has the drawback of being difficult to compute (Section 5). In Section 6 we show how to approximate this notion of envy, and provide experimental results in Section 7 suggesting that it can be used in practice.

2 Our model

We consider an extended house allocation problem involving a set $\mathcal{N}$ of n agents partitioned into k types/groups $\mathcal{T} = \{T_1, \ldots, T_k\}$ and a set $\mathcal{M}$ of m items/houses with $|\mathcal{M}| \leq |\mathcal{N}|$. An allocation A is a collection of k bundles $A_1, \ldots, A_k$ such that:

- $A_1 \cup \ldots \cup A_k = \mathcal{M}$,
- $A_p \cap A_q = \emptyset$ for all $p, q \in [k]$, with $p \neq q$,
- and $|A_p| \leq |T_p|$ for all $p \in [k]$.

For a given allocation A, A_p is the set of items allocated to group T_p. Moreover, the bundle received by agent $i \in \mathcal{N}$ will be denoted $A(i)$ in the sequel. Recall that, in our setting, each item can be assigned to at most one agent, and each agent can only receive at most one item from the bundle allocated to its group. Therefore, for any group $T_p \in \mathcal{T}$ and any agent $i \in T_p$, we have $A(i) \subseteq A_p$ and $|A(i)| \leq 1$. Moreover, for any subgroup $G \subseteq T_p$, the subset of items received by the agents in G will be denoted by $A(G)$, i.e. $A(G) = \cup_{i \in G} A(i)$. In particular, we have $A(G) = A_p$ for $G = T_p$.

For every group $T_p \in \mathcal{T}$, we are given a type-value function $V_p : 2^{\mathcal{M}} \to \mathbb{R}_+$ which quantifies some concept of overall welfare derived by the agents in T_p from any bundle of items $B \subseteq \mathcal{M}$. In this paper, we rely on the *utilitarian type-value function* as defined in [3]:

$$V_p(B) = \begin{cases} \max_{X \in \mathcal{X}_B} \sum_{i \in T_p} \sum_{j \in B} u_i(j) x_{ij}, & \text{if } B \neq \emptyset; \\ 0, & \text{otherwise.} \end{cases}$$

where $u_i(j) \in \mathbb{R}_+$ is the utility derived by agent i from item j, and $\mathcal{X}_B$ is the set of all binary matrices $X = (x_{ij})_{i \in T_p, j \in B}$ representing (maximum-cardinality) matchings in the bipartite graph $\mathcal{G} = (T_p \cup B, T_p \times B)$. In other words, we assume that every group measures the satisfaction of its own agents by the classical utilitarian social welfare, and that it uses an internal allocation procedure that distributes the items to the agents of the group optimally w.r.t. the utilitarian social welfare. Hence, for a given allocation A, group $T_p \in \mathcal{T}$ assigns the items in A_p to agents in T_p in such a way that allocation $\{A(i)\}_{i \in T_p}$ maximizes the utilitarian social welfare. We extend this definition to any subgroup $G \subseteq T_p$ of same-type agents, by simply defining V_G as V_p restricted to agents in G.

Note that, for any $G \subseteq T_p$, function V_G is not additive but submodular, and that $V_G(B)$ can be computed in polynomial time for any $B \subseteq \mathcal{M}$ as it amounts to solve the instance of the classical maximum weight assignment problem defined by the bipartite graph $\mathcal{G} = (G \cup B, G \times B)$ whose weights are given by $\{u_i(j)\}_{(i,j) \in G \times B}$.

3 Group envy-freeness: an axiomatic perspective

We are interested in assessing the fairness of allocations among groups of agents. At the individual level, envy-freeness is simply defined as follows: an agent $i \in \mathcal{N}$ envies another agent $i' \in \mathcal{N}$ when she prefers the bundle $A(i')$ to her own bundle $A(i)$, i.e. when $u_i(A(i')) < u_i(A(i))$. An allocation is envy-free when no agent envies another. Group envy-freeness extends this notion to groups of agents. Below we provide a critical discussion of some existing notions when used in our context as it poses a specific challenge due to the combination of two features: agents receive at most one item (making the value function not additive), and groups may be of different sizes. We approach the question by taking an axiomatic perspective [18], and identify three key properties we would like our notion to satisfy.

Axiom 1 (Group Envy under Unanimous Pairwise Envy) *For any allocation A and any two groups $T_p, T_q \in \mathcal{T}$:*

$$[\forall i \in T_p, \forall i' \in T_q, u_i(A(i)) < u_i(A(i'))] \Rightarrow T_p \textit{ envies } T_q$$

Axiom 2 (Group Envy-Freeness under Useless Acquisition) *For any allocation A and any two groups $T_p, T_q \in \mathcal{T}$:*

$$V_p(A_p) = V_p(A_p \cup A_q) \Rightarrow T_p \textit{ does not envy } T_q$$

Axiom 3 (Monotonicity under Unanimously Envied Agent Addition) *Let us consider an instance I and an allocation A such that group $T_p \in \mathcal{T}$ equal-sized envies group $T_q \in \mathcal{T}$. For any augmented instance I' obtained by simply adding one agent i_0 to T_q and one object j_0 allocated to T_q:*

$$[\forall i \in T_p, u_i(A(i)) < u_i(j_0)] \Rightarrow T_p \textit{ envies } T_q \textit{ in instance } I'$$

Intuitively, Axiom 1 says that if there is envy from every agent of one group to every agent of another group, then the first group must envy the second group; Axiom 2 reflects the intuition that a group which does not improve its satisfaction upon receiving all the items of another group should not be envious of this group; and finally

Axiom 3 captures the idea that adding an agent together with an item that is unanimously envied to a group that is already envied should not delete envy.

To define a *typewise* extension of individual envy-freeness, a first idea that may come to mind is to compare the bundles of items received by the different groups of agents through their type-value functions, as proposed by Benabbou et al. [3]:

Definition 1 (Typewise Envy) *For any allocation A, group $T_p \in \mathcal{T}$ is said to* typewise envy *group T_q if and only if $V_p(A_p) < V_p(A_q)$. An allocation A is said to be* typewise envy-free *(TEF) if and only if no group typewise envies another group.*

However, this definition is oblivious of the size of groups, which induces counter-intuitive results in our context:

Observation 1 (TEF does not verify Axiom 1) *Let us consider an instance with $|\mathcal{M}| = 103$ items and two groups T_1 and T_2, where T_1 is a group of 100 agents and T_2 is a group of 3 agents. Let A be an allocation such that $|A_1| = 100$ and $u_i(j) = 1$ for all agents $i \in T_1$ and all items $j \in A_1$, whereas $|A_2| = 3$ and $u_i(j) = 10$ for all agents $i \in T_1$ and all items $j \in A_2$. In that case, we have $u_i(A(i)) < u_i(A(i'))$ for all $i \in T_1$ and all $i' \in T_2$, yet T_1 does not typewise envy T_2 since $V_1(A_1) = 100 \geq V_1(A_2) = 30$.*

Chakraborty et al. introduce the notion of *weighted envy-freeness* [5] as a way to handle situations where some agents might be entitled to higher utilities than others. It can easily be adapted to the case of predetermined groups of agents by considering each group as an agent having an entitlement depending on its size [16]. The weighted envy-freeness is defined using a weight function $w : \mathcal{T} \to \mathbb{R}$ which attributes a weight to each group. In our case, this weight will be equal to the size of the group, i.e. $w(T_p) = |T_p|$ for all $T_p \in \mathcal{T}$. We will designate the weight of T_p as w_p.

Definition 2 (Weighted Envy) *For any allocation A, group $T_p \in \mathcal{T}$* weighted envies *group $T_q \in \mathcal{T}$ if and only if:*

$$\frac{V_p(A_p)}{w_p} < \frac{V_p(A_q)}{w_q}$$

Allocation A is said to be weighted envy-free *(WEF) if no group weighted envies another group.*

However, in our case, not only are the utilities not additive, but also we cannot guarantee that adding an item to the bundle of a group will strictly increase its utility (as it can only use a limited number of items at once). This will cause erratic behaviour with this definition of envy. In particular, this envy notion will verify neither Axiom 2 nor Axiom 3, as we shall see now:

Observation 2 (WEF does not satisfy Axiom 2) *Let us consider an instance with three items $\mathcal{M} = \{j_1, j_2, j_3\}$ and two groups of agents $T_1 = \{i_1, i_2\}$ and $T_2 = \{i_3\}$ with the following utilities:*

$u_i(j)$	j_1	j_2	j_3
i_1	10	0	10
i_2	0	1	0
i_3	10	0	10

For allocation A such that $A_1 = \{j_1, j_2\}$ and $A_2 = \{j_3\}$, we have:

$$\frac{V_1(A_1)}{|T_1|} = \frac{10+1}{2} < \frac{10}{1} = \frac{V_1(A_2)}{|T_2|}$$

Thus T_1 weighted envies T_2. However, we have $V_1(A_1) = V_1(A_1 \cup A_2) = 11$ since T_1 has no way to increase its overall utility using j_3.

Observation 3 (WEF does not satisfy Axiom 3) *Let us consider an instance I with two groups $T_1 = \{i_1\}$ and $T_2 = \{i_2, i_3\}$, three items $\mathcal{M} = \{j_1, j_2, j_3\}$, and utility values $u_1(j_1) = 2$ and $u_1(j_2) = u_1(j_3) = 5$. For allocation A defined by $A_1 = \{j_1\}$ and $A_2 = \{j_2, j_3\}$, group T_1 weighted envies group T_2 since we have:*

$$\frac{V_1(A_1)}{|T_1|} = \frac{2}{1} < \frac{5}{2} = \frac{V_1(A_2)}{|T_2|}$$

However, if we consider the instance obtained by adding agent i_4 to T_2 and an object j_4 allocated to T_2 with $u_1(j_4) = 5$, then T_1 does not weighted envy T_2 anymore since we now have:

$$\frac{V_1(A_1)}{|T_1|} = \frac{2}{1} > \frac{5}{3} = \frac{V_1(A_2)}{|T_2|}$$

Thus we added to T_2 an agent with an item that all agents in T_1 prefer to their own, yet T_1 no longer weighted envies T_2.

For the sake of completeness, we note that TEF satisfies Axiom 2 and Axiom 3, while WEF satisfies Axiom 1. Since the proofs are rather simple, we omit them due to space constraints.

4 A new notion of envy among groups

We now introduce a new notion of envy among groups better suited to house allocation settings, which is based on the counterfactual comparison of subgroups of the same size.

Definition 3 (Envy estimated by projection on equal-sized groups) *For any allocation A, we say that $T_p \in \mathcal{T}$* equal-sized envies *$T_q \in \mathcal{T}$ if and only if one of the following two conditions is verified:*

- *$|T_p| \leq |T_q|$ and $V_p(A_p) < V_p(A(G))$ for at least one subgroup $G \subseteq T_q$ with $|G| = |T_p|$.*
- *$|T_p| > |T_q|$ and $V_G(A(G)) < V_G(A_q)$ for at least one subgroup $G \subset T_p$ with $|G| = |T_q|$.*

Allocation A is said to be equal-sized envy-free *(ESEF) if no group in $\mathcal{T}$* equal-sized envies *another group in $\mathcal{T}$.*

The idea behind this definition is to only allow the comparison of groups of the same size, by considering all the possible subgroups of the larger group; if there is envy in at least one of these comparisons, then the group is considered to have valid reasons to be envious. Note that, whenever $|T_p| = |T_q|$, we only compare the bundles obtained by the whole groups.

We now show that this newly introduced definition of group-envy satisfies Axioms 1, 2 and 3.

Observation 4 (ESEF satisfies Axiom 1) *Let us consider an instance with two groups $T_p, T_q \in \mathcal{T}$ and some allocation A such that $u_i(A(i)) < u_i(A(i'))$ for all agents $i \in T_p$ and all agents $i' \in T_q$. Let us show that T_p equal-sized envies T_q:*

- *If $|T_p| \leq |T_q|$ then we necessarily have $V_p(A_p) < V_p(A(G))$ for any $G \subseteq T_q$ with $|G| = |T_p|$ since $u_i(A(i)) < u_i(A(i'))$ for all $i \in T_p$ and all $i' \in G$. Thus T_p equal-sized envies T_q.*
- *If $|T_p| > |T_q|$ then we necessarily have $V_G(A(G)) < V_G(A_q)$ for any $G \subset T_p$ with $|G| = |T_q|$ since $u_i(A(i)) < u_i(A(i'))$ for all $i \in G$ and all $i' \in T_q$. Thus T_p equal-sized envies T_q.*

Observation 5 (ESEF satisfies Axiom 2) *Let us consider an instance with two groups $T_p, T_q \in \mathcal{T}$ and some allocation A such that $V_p(A_p \cup A_q) = V_p(A_p)$. Let us prove by contradiction that T_p does not equal-sized envy T_q. Assume that T_p equal-sized envies T_q. Two cases may occur:*

- *If $|T_p| \leq |T_q|$ then there is at least one subgroup $G \subseteq T_q$ such that $V_p(A_p) < V_p(A(G))$, and so there exists at least one agent $i \in T_p$ and one item $j \in A(G)$ such that $u_i(A(i)) < u_i(j)$. By simply allocating j to i (instead of $A(i)$), we obtain an allocation of $A_p \cup A_q$ to agents in T_p that gives a larger overall utility than just A_p. Thus $V_p(A_p \cup A_q) > V_p(A_p)$ yielding a contradiction.*
- *If $|T_p| > |T_q|$ then there is at least one subgroup $G \subset T_p$ such that $V_G(A(G)) < V_G(A_q)$, and so there exist at least one agent $i \in G \subset T_p$ and one item $j \in A_q$ such that $u_i(A(i)) < u_i(j)$. Thus by allocating j to i (instead of $A(i)$) we obtain an allocation of $A_p \cup A_q$ to agents in T_p that provides a larger utility than A_p. Thus $V_p(A_p \cup A_q) > V_p(A_p)$ which gives a contradiction.*

Observation 6 (ESEF satisfies Axiom 3) *Let us consider an instance I and an allocation A such that group $T_p \in \mathcal{T}$ equal-sized envies group $T_q \in \mathcal{T}$. Let I' be an instance obtained from I by simply adding one agent i_0 of type T_q and one object j_0 such that $u_i(A(i)) < u_i(j_0)$ for all $i \in \mathcal{N}\backslash\{i_0\}$ (other utility values being arbitrary chosen). Let $\mathcal{T}' = \{T'_1, \dots, T'_k\}$ denote the set of types of instance I'. In instance I', let A' be the allocation defined by $A'_q = A_q \cup \{j_0\}$ and $A'_r = A_r$ for all $r \in [k]$ with $r \neq q$. Let us show that T'_p equal-sized envies T'_q in allocation A'. Two cases must be considered:*

- *If $|T_p| \leq |T_q|$, then we know that there exists $G_0 \subseteq T_q$ with $|G_0| = |T_p|$ such that $V_p(A_p) < V_p(A(G_0))$ (since T_p equal-sized envies T_q in A). For instance I' and allocation A', we have $|T'_p| < |T'_q|$, $G_0 \subseteq T_q \subset T'_q$ with $|G_0| = |T'_p|$ and $V_p(A'_p) = V_p(A_p) < V_p(A(G_0)) = V_p(A'(G_0))$. Therefore T'_p equal-sized envies T'_q in A'.*
- *If $|T_p| > |T_q|$ then there exists $G \subset T_p$ with $|G| = |T_q|$ such that $V_G(A(G)) < V_G(A_q)$ (since T_p equal-sized envies T_q in A). Let $G_0 = G \cup \{i\}$ for some $i \in T_p \backslash G$. For instance I' and allocation A', we have $V_{G_0}(A'(G_0)) = V_G(A'(G)) + u_i(A'(i)) = V_G(A(G)) + u_i(A(i)) < V_G(A_q) + u_i(j_0) \leq V_{G_0}(A'_q)$. If $|T'_p| > |T'_q|$ then we can directly conclude that T'_p equal-sized envies T'_q in A' since G_0 is a subgroup of T_p such that $|G_0| = |T'_q|$ and $V_{G_0}(A'(G_0)) < V_{G_0}(A'_q)$. Otherwise, we necessarily have $|T'_p| = |T'_q|$, and therefore $G_0 = T'_p$. In that case, $V_{G_0}(A'(G_0)) < V_{G_0}(A'_q)$ can be rewritten $V_p(A'_p) < V_p(A'_q)$ which implies that T'_p equal-sized envies T'_q in A'.*

Although ESEF verifies our three axioms, an ESEF allocation may not exist (even when considering an "up-to-one-good" relaxation), as shown by the following example.

Example 1 *Consider an instance with $|\mathcal{M}| = 4$ items and $|\mathcal{N}| = 9$ agents partitioned into two groups $\mathcal{T} = \{T_1, T_2\}$ with $|T_1| = 6$ and $|T_2| = 3$. Let us assume that $u_i(j) = 1$ for all agents $i \in \mathcal{N}$ and all items $j \in \mathcal{M}$. In any allocation A such that $|A_2| \geq 1$, at least 3 agents in T_1 are empty-handed. Therefore, there exists $G \subset T_1$ with $|G| = |T_2| = 3$ such that $V_G(A(G)) = 0 < 1 \leq V_G(A_2)$. Hence T_1 equal-sized envies T_2 whenever $|A_2| \geq 1$. In any allocation A such that $|A_2| = 0$, we have $V_2(A_2) = 0 < 1 \leq V_2(A(G))$ for any subgroup $G \subset T_1$ with $|G| = |T_2| = 3$ since only two agents in T_1 are empty-handed. Thus, T_2 equal-sized envies T_1 in that case, which shows that there is no ESEF allocation in this instance.*

Note that considering the "up to one good" relaxation which consists in removing the "best" item from the envied (sub)group before comparing utilities, does not solve the problem for this instance: T_1 equal-sized envies T_2 up to one good when $|A_2| \geq 2$, whereas T_2 equal-sized envies T_1 up to one good when $|A_2| < 2$.

Proposition 1 *The problem of determining if there exists an ESEF allocation is NP-complete, even for problems involving only two groups of equal size with identical individual utilities.*

Proof. First, let us prove that the problem is in NP. Note that determining if a given allocation A is ESEF can be done by performing (at most) a quadratic number of tests: for every pair $(T_p, T_q) \in \mathcal{T}^2$ with $T_p \neq T_q$, we need to check whether T_p equal-sized envies T_q or not. For a given pair (T_p, T_q), it can be done as follows:

- if $|T_p| \leq |T_q|$ then T_p equal-sized envies T_q iff there exists $G \subseteq T_q$ such that $|G| = |T_p|$ and $V_p(A_p) < V_p(A(G))$ (by definition). Note that there is no need to enumerate all possible subgroups G to check whether T_p equal-sized envies T_q, as we can only focus on the "best" subgroup. Indeed T_p equal-sized envies T_q iff $V_p(A_p) < \max_{G \subseteq T_q : |G| = |T_p|} V_p(A(G))$. As already mentioned before, computing $V_p(A_p)$ can be done in polynomial time as it amounts to solve an instance of the classical maximum weight assignment problem. Regarding $\max_{G \subseteq T_q : |G| = |T_p|} V_p(A(G))$, it can also be reduced to the maximum weight assignment problem, applied on the bipartite graph $g = (T_p \cup T_q, T_p \times T_q)$ with weights $\{u_i(A(i'))\}_{(i,i') \in T_p \times T_q}$ (which is solvable in polynomial time).
- If $|T_p| > |T_q|$ then T_p equal-sized envies T_q iff there exists $G \subseteq T_p$ such that $|G| = |T_q|$ and $V_G(A(G)) < V_G(A_q)$ (by definition). As in the previous case, we can only focus on the "worst" subgroup: T_p equal-sized envies T_q iff $\max_{G \subseteq T_p : |G| = |T_q|} \{V_G(A_q) - V_G(A(G))\} > 0$. This can be reduced to the maximum weight assignment problem applied on the bipartite graph $g = (T_p \cup T_q, T_p \times T_q)$ with weights $\{u_i(A(i')) - u_i(A(i))\}_{(i,i') \in T_p \times T_q}$ (which is solvable in polynomial time).

Since determining if a given allocation A is ESEF can solved in polynomial time, the existence problem of an ESEF allocation can be solved in polynomial time by a nondeterministic algorithm which creates an allocation and return *True* if it is ESEF. The problem is therefore in NP.

We now prove that the problem is NP-complete by describing a polynomial time reduction from the partition problem to our problem. Given an instance of the partition problem defined by a set of positive integers $S = \{s_1, \dots, s_f\}$, we construct an instance of our problem with only two groups (T_1 and T_2) of size f, and where $\mathcal{M}$ the set of items only includes one item j per integer $s_j \in S$ with $u_i(j) = s_j$ for any agent $i \in \mathcal{N}$. Since $|T_1| = |T_2|$, then our constructed instance is a 'yes'-instance iff there exists an allocation A such that $V_1(A_1) \geq V_1(A_2)$ and $V_2(A_2) \geq V_2(A_1)$, i.e. such that $V_1(A_1) = V_1(A_2)$ (since $V_1 = V_2$). Thus our constructed instance is a 'yes'-instance iff $\sum_{j \in A_1} s_j = \sum_{j \in A_2} s_j$ (since $u_i(j) = s_j$ for any pair $(i, j) \in \mathcal{N} \times \mathcal{M}$), i.e. iff the instance of the partition problem is a 'yes'-instance (A_1 and A_2 defining a valid partition). □

This motivates us to define a notion of degree of envy directly derived from Definition 3. A possible intuition behind the definition is as follows: suppose a small group contemplates a large group knowing that they would only obtain a number of items corresponding to their own group size. Without any prior information regarding the way these items would be picked, our degree of envy corresponds to the likelihood of being envious. The reasoning is the same when a large group contemplates a small group, except that the uncertainty lies instead on which agents would be picked to match the small group size. The formal definition is given below:

Definition 4 (Degree of envy) *For any allocation A, the* degree of envy *of $T_p \in \mathcal{T}$ towards $T_q \in \mathcal{T}$, denoted by d^A_{pq}, is defined by:*

$$d^A_{pq} = \begin{cases} \frac{|\{G \subseteq T_q : |G| = |T_p| \& V_p(A_p) < V_p(A(G))\}|}{|\{G \subseteq T_q : |G| = |T_p|\}|} & \text{if } |T_p| \leq |T_q| \\ \frac{|\{G \subseteq T_p : |G| = |T_q| \& V_G(A(G)) < V_G(A_q)\}|}{|\{G \subseteq T_p : |G| = |T_q|\}|} & \text{otherwise.} \end{cases}$$

Clearly, an allocation A is ESEF iff $d^A_{pq} = 0$ for all pairs $(T_p, T_q) \in \mathcal{T}^2$. When the degree of envy is equal to 1, we know that there is envy regardless of which equal-size subgroup is considered for comparison. Observe that when two groups are of equal size to start with, the degree of envy may only be 0 or 1.

5 On the computation of the degree of envy

We now focus on the computation of the degree of envy between two groups. To do so, let us consider the following counting problem:

NUMBER OF ENVIED SUBSETS PROBLEM : Given an allocation A and two groups $T_p, T_q \in \mathcal{T}$ with $|T_p| \leq |T_q|$, how many subsets $G \subseteq T_q$ are such that $|G| = |T_p|$ and $V_p(A_p) < V_p(A(G))$?

We now show that this counting problem is #P-complete, #P being the complexity class of the counting problems associated with NP decision problems (as first introduced by Valiant in [20]).

Proposition 2 *The number of envied subsets problem is #P-complete.*

Proof. This counting problem is obviously in #P, as we can solve the associated existence problem in polynomial time with a simple nondeterministic algorithm which picks a bundle and checks if its overall utility is high enough to be envied.

In order to show that our problem is #P-complete, we now consider the counting subset sum problem (which can be shown to be #P-complete using a reduction from the counting problem associated with 3-SAT [19]). It can be defined as follows: given a set of positive integers $S = \{s_1, \ldots, s_f\}$ and an integer T, how many subsets of S sum up to T ? Given an instance of the counting subset sum problem, we construct two instances I and I' of our allocation problem. First, let I be an instance with $|\mathcal{M}| = 2f$ items and two groups T_1 and T_2 such that $|T_1| = f$ and $|T_2| = 2f$. Moreover, $\mathcal{M}$ is partitioned into two sets M_1 and M_2 of equal-size such that, for any agent $i \in T_1$, we have $u_i(j) = T/f$ for all $j \in M_1$, and $u_i(j) = s_j$ for all $j \in M_2$. Consider the allocation A defined by $A_1 = M_1$ and $A_2 = M_2$. In that case, we have $V_1(A_1) = f \times T/f = T$, and for any subset $G \subseteq T_2$ with $|G| = |T_1| = f$, we have $V_1(A(G)) = \sum_{j \in A(G)} s_j$. Hence solving I gives the number n_I of subsets of S whose sum is strictly larger than T. Similarly, let I' be the instance obtained from I by simply changing the following utility values: $u_i(j) = T/f - 1/f$ for all $i \in T_1$ and $j \in M_1$. As a result, solving I' gives the number of subsets of S whose sum is strictly larger than $T - 1$. Since S is only composed of positive integers, then it gives the number $n_{I'}$ of subsets whose sum is larger or equal to T. Then, we can obtain the number n_T of subsets of S whose sum is exactly T as follows: $n_T = n_{I'} - n_I$. The number of envied subsets problems is thus #P-complete. □

Symmetrically one can easily prove that the following counting problem is also #P-complete:

NUMBER OF ENVIOUS SUBSETS PROBLEM : Given an allocation A and two groups $T_p, T_q \in \mathcal{T}$ with $|T_p| > |T_q|$, how many subsets $G \subset T_p$ are such that $|G| = |T_q|$ and $V_G(A(G)) < V_G(A_q)$?

Since computing the degree of envy amounts to solve either an instance of the number of envied subsets problem or an instance of the number of envious subsets problem (and then divide by the number of possible subgroups), we obtain the following complexity result:

Corollary 1 *Computing the degree of envy is #P-complete.*

6 Approximating the degree of envy

Even if computing the degree of envy is #P-complete, we now show that it can be approximated efficiently when utilities are *type-uniform* [4], i.e. when agents of the same type/group have the same utility values. We use a result from [12] which guarantees that, if we can generate random envied (or envious) subgroups following an almost uniform distribution, then we can approximate the total number of envied (or envious) subgroups by a *Fully Polynomial Randomised Approximation Scheme* (FPRAS). To simplify the presentation, we focus here on the number of envied subsets problem, as the number of envious subsets problem can be solved similarly, and we assume that $|T_p| < |T_q|$, as there is no need for an approximation algorithm when $|T_p| = |T_p|$ (only one possible envied subset). Moreover, we assume here that no agent is empty-handed in allocation A, adding dummy items with zero utility when necessary. Let $n_p = |T_p|$ and $n_q = |T_q|$.

If Ω denotes the set of envied subgroups, then given an instance I of the number of envied subsets problem, and a parameter $0 < \varepsilon < 1$, a probabilistic algorithm $\mathcal{A}$ is a FPRAS if and only if $Pr(|\mathcal{A}(I, \varepsilon) - |\Omega|| \leq \frac{\varepsilon}{2}|\Omega|) \geq \frac{3}{4}$ and its running time is polynomial in $|I|$ and ε^{-1}. We designed our FPRAS following the guidance provided in [11]. More precisely, our FPRAS computing the size of the solution set Ω is composed of:

- A poly-time sampling method generating a batch of sample solutions, following a known probability distribution over Ω.
- A family of sets Ω_i, with $i \in [\![0, n_p]\!]$, such that $\Omega_0 \subset \Omega_1 \subset \ldots \subset \Omega_{n_p} = \Omega$ and $|\Omega_0| = 1$.

The general principle is to use the sampling method to approximate the successive ratio $\frac{|\Omega_i|}{|\Omega_{i+1}|}$, and then derive an approximation of $|\Omega|$ using the following equalities:

$$\frac{|\Omega_0|}{|\Omega_1|} \times \ldots \times \frac{|\Omega_{n_p-1}|}{|\Omega_{n_p}|} = \frac{|\Omega_0|}{|\Omega_{n_p}|} = |\Omega|^{-1}$$

The successive ratios are obtained by proceeding as follows: first sample the entire set Ω to estimate $|\Omega_{n_p-1}|/|\Omega_{n_p}|$, then sample the subset Ω_{n_p-1} to estimate $|\Omega_{n_p-2}|/|\Omega_{n_p-1}|$ and so on. A study of the expectancy and variance of the results was made in [11]. Our goal is to find a good sampling method for our problem, allowing us to sample our solution set according to a uniform distribution. In order to do so, we will use a random walk on the Markov chain (Ω, P) where the state set Ω is the solution set and P is the transition matrix (defined hereafter). This can be ensured if the Markov chain satisfies the following properties:

(1) be *irreducible*, which means that we must find a neighbourhood relationship such that our solution set is connected;
(2) have a *stationary probability distribution* π over the set of states;
(3) be *rapidly mixing* [17], meaning that for any starting state $X \in \Omega$ and any state $Y \in \Omega$, $P^t(X, Y)$ is close to $\pi(Y)$ for t very small compared to $|\Omega|$.

To construct such a Markov chain, we take inspiration from a construction introduced in [6]. Here Ω corresponds to the set of envied subsets, i.e. all $G \subset T_q$ such that $V_p(A_p) < V_p(A(G))$. Since no agent is empty-handed, Ω can also be defined as the set of bundles X, with $X \subseteq A_q$ and $|X| = |T_p|$, such that $V_p(A_p) < V_p(X)$. We will say that two bundles are *neighbours* if and only if one can be obtained from the other by removing exactly one item and adding a single item. Transition matrix P is derived from the following rule: from any state $X \in \Omega$, we stay in X with probability $\frac{1}{2}$, otherwise we pick uniformly at random one item in X and one item $A_q \backslash X$, and we swap them: if the resulting bundle Y is in Ω then we move from state X to state Y (otherwise we stay in X).

Proposition 3 *The Markov chain* (Ω, P) *is irreducible, and has a stationary distribution.*

Proof. In order to show the irreducibility of the state space, we need to show that, given X and Y two bundles in Ω, we can go from X to Y by successively adding and removing one good in such a way that each bundle on the way is also in Ω. To do so, we create the weighted directed bipartite graph $\mathcal{G}_X = (N_X \cup M_X, E_X)$ such that N_X contains one vertex for each *agent* of T_p and M_X contains one vertex for each *item* of A_q. Then, for any $i \in N_X$ and any $j \in M_X$, we add:

- an arrow (i, j) in E_X iff T_p assigns item j to agent i when T_p receives bundle X, and if so then its weight $w(i, j)$ is set to $-u_i(j)$.
- an arrow (j, i) in E_X iff T_p does *not* assign item j to i when it receives bundle X, and if so then its weight $w(j, i)$ is set to $u_i(j)$.

Let S_{XY} be the set of arrows of E_X defined by:

- $(j, i) \in S_{XY}$ iff T_p assigns item j to agent i when T_p receives bundle Y, but not when it receives X.
- $(i, j) \in S_{XY}$ iff T_p assigns item j to agent i when T_p receives bundle X, but not when it receives Y.

Note that the set S_{XY} will be composed of $\mathcal{R}$ disjoint paths of even length from items that are in $Y \backslash X$ to items that are in $X \backslash Y$. Moreover, any item j in a path that is not on an extremity of the path is in $X \cap Y$. Thus, *each path corresponds to an exchange between an item only in* X *and an item only in* Y. Let us define the weight of a path as the sum of the weights of its arrows, and let us sort the paths in S_{XY} in decreasing order of weights so that w_r denotes the weight of the r^{th} path. By applying the exchange corresponding to each path taken in that order, we know that we move from X to Y. It remains to show that each intermediary bundle Z_r, obtained by applying the exchanges described by the r first paths in S_{XY}, belongs to Ω. Since $X, Y \in \Omega$, then it is sufficient to show that $V_p(Z_r) \geq V_p(X)$ or $V_p(Z_r) \geq V_p(Y)$. This is actually the case because $V_p(Z_r) \geq \max\{V_p(X) + \sum_{s=1}^{r} w_s; V_p(Y) - \sum_{s=r+1}^{\mathcal{R}} w_s\}$, which is due to the fact that applying an exchange corresponding to a path gives an allocation that is not necessarily optimal w.r.t the group utility, and we always have either $\sum_{s=1}^{r} w_s \geq 0$ or $\sum_{s=r+1}^{\mathcal{R}} w_s \leq 0$ since paths are ordered in decreasing order of weights. Thus, for any pair $(X, Y) \in \Omega \times \Omega$, there exists a sequence of neighbouring bundles in Ω between X and Y, which shows that the set of states is connected and that the Markov chain is irreducible.

To show the stationarity of the distribution, observe that the transition matrix is such that for any $X \in \Omega$, $P(X, X) \geq \frac{1}{2}$ and $P(X, Y) = 1/2n_p(n_q - n_p)$ for any $Y \in \Omega$ neighbour of X. Then, we can deduce the stationary probability distribution π from Q which is defined for all $X \in \Omega$ and all $Y \in \Omega$ neighbour of X by:

$$Q(X,Y) = \pi(X)P(X,Y) = \pi(Y)P(Y,X) = \frac{1}{|\Omega|} \times \frac{1}{2n_p(n_q - n_p)}$$

□

The proof that the Markov chain is rapidly mixing is more involved and requires to use and combine several results from the literature.

Proposition 4 *The Markov chain* (Ω, P) *is rapidly mixing.*

Proof. Formally, (Ω, P) is rapidly mixing if for any starting state $X \in \Omega$, $\tau_X(\varepsilon)$ is $O(poly(\log(|\Omega|), \log(\varepsilon^{-1}))$ where the mixing time $\tau_X(\varepsilon)$ is defined by:

$$\tau_X(\varepsilon) = \min\{t : \forall t' \geq t, \Delta_X(t') < \varepsilon\}$$

with the variation distance $\Delta_X(t)$ given by:

$$\Delta_X(t) = \frac{1}{2} \sum_{Y \in \Omega} |P^t(X,Y) - \pi(Y)|$$

Our proof here will be a slightly modified version of the proof elaborated for the knapsack problem in [10]. This proof relies on the fact that, if we can define a flow f that routes one unit of flow between every pair of unequal states, then we have:

$$\tau_X(\varepsilon) \leq 2n_q \frac{C(f)}{|\Omega|} L(f)(n_q + \ln(\varepsilon^{-1}))$$

where $C(f)$ is the maximum flow across all edges, and $L(f)$ is the length of the longest flow-carrying path. Thus, we simply need to prove that $C(f) = |\Omega| poly(n_q)$ and $L(f) = poly(n_q)$. Now our objective is to construct an appropriate flow while ensuring that we can compute the maximum flow that can go through any state $Z \in \Omega$. To do so, for any two states in $X, Y \in \Omega$, we are going to spread the flow between these two points over all the possible combinations of swaps that lead from one to the other, and then, for each combination of swaps, over a set of flow paths obtained from a family of permutations of those swaps.

Let X and Y be any two states in Ω, and let $E_{X \to Y}$ be any set of swaps in $\{(x, y) : x \in X \setminus Y, y \in Y \setminus X\}$ leading to Y from X. For any swap (x, y), let $w_{(x,y)} = u(y) - u(x)$ be the weights of swap (x, y), where u is the utility function of agents in group T_p. Following [10], we partition $E_{X \to Y}$ into two sets H and S such that H contains the 29 swaps with the largest weight (if this many exist), and S contains all the others. Let us set $m = |S|$ and let $\{w_i\}_{i=1}^m$ denote the weights of the swaps in S. We will further divide the set H into two sets H^- and H^+ corresponding to negative and non-negative weights respectively. Let a be a function defined over $\Omega \cup 2^{H \cup S}$ such that $a(Z) = V_p(Z)$ for any $Z \in \Omega$ and $a(E) = \sum_{(x,y) \in E} w_{(x,y)}$ for any $E \subseteq E_{X \to Y}$.

Let us now specify the flow paths between states X and Y. The flow will be evenly split among all families of swaps $E_{X \to Y}$ (of which there are $n_p!$), and then among the flow paths of each family; we thus divide a flow of $\frac{1}{n_p!}$ over the flow paths of each family. For a given family $E_{X \to Y}$ (which uniquely defines a set H and a set S), the flow paths will be created from a family of permutations of S as described in [10], except that we apply the swaps in H (and the corresponding reverse swaps if necessary) instead of adding and removing items. This allows us to control the amount of flow going through each path, a result from [10] that we will use later[1].

[1] This is thanks to Claim (27) from [10], which is used as a basis to generate the flow paths. The proof of this claim translates seamlessly to our case, except that H_Y must be replaced by H^+, H_X by H^- and $a(H_X)$ by $-a(H^-)$, and that we must use our own definition of function a.

We can now calculate the amount of flow that goes through each state $Z \in \Omega$. For this, we will compute the number of pairs $(X, Y) \in \Omega^2$ for which at least one flow path goes through Z. We encode each triplet (X, Y, E) by a 6-tuple (Z', h_1, h_2, U, H', E) with $Z' \in \Omega$, $h_1, h_2 \in H$, $U \subset S$, $H' \subset H$ and E a family of swaps in $\{(x, y) : x \in Z \setminus Z', y \in Z' \setminus Z\}$. The tuples are defined as in [10] with the following differences: h_1 and h_2 are swaps in E instead of objects, Z_{h_1,h_2} is the state reached from Z by applying or reversing h_1 and h_2, and E is a set of swaps from Z' to Z which is also the set of swaps from X to Y (except that some may be reversed). Due to the way the flow paths are defined, we know that there exists a universal constant C such that, for each tuple, the total flow going through Z is bounded by $\frac{Cm^2}{\binom{m}{|U|}n!}$ (this is thanks to the result of [10] we mentioned earlier). We can thus calculate the total flow by counting all possible 6-tuples: there are (i) $n_p(n_p-1)(n_q-n_p)(n_q-n_p-1)$ possibilities for h_1 and h_2, *(ii)* $|\Omega|$ possibilities for Z', *(iii)* $n_p!$ possibilities for E, *(iv)* $\binom{m}{|U|}$ possibilities for U for each $|U|$, and *(v)* 2^{29} possibilities for H'. By summing over all the flow paths, we obtain a bound on the maximum flow passing through any state Z:

$$f(Z) \leq 2^{29}.|\Omega|.Cm^3.n_p(n_p-1)(n_q-n_p)(n_q-n_p-1)$$

Finally, since $C(f) = \max_{Z\in\Omega} f(Z)$, we obtain the following bound: $\tau_X(\varepsilon) \leq poly(n_p, n_q, \log(\varepsilon^{-1}))$. □

Note that this approximation method still satisfies the axioms introduced in Section 3, as it correctly detects whether envy exists or not: indeed, the first step of the random walk is to compute an envied subgroup. The algorithm will thus stop (returning 0) if none is found, otherwise it will return a product of ratio strictly greater than 0.

7 Experiments

Our previous results have two main limitations: first, the number of steps required by the theoretical bound is very high, even if n is small, and second, it requires utility functions to be identical between agents of a same group (type-uniformity). To gain some insights regarding the behaviour of our approximation in practice, we decided to study the convergence rate on small instances, first in the case of type-uniform utilities and then with general utility functions.

The experiments were conducted considering a group T_p of size n_p (for n_p going from 4 to 8) comparing itself to a larger group T_q of size $2n_p$ (comparing a larger to a smaller group can be easily reduced to the opposite). Utility functions were generated by picking uniformly at random a value between 0 and 1, one for each item for type-uniform utilities, and one for each (agent, item) pair for general utilities (utilities are thus not correlated at all). From there, we generated the set of envied bundles Ω, and the corresponding Markov chain (Ω, P), as described in Section 6. We then calculated the evolution of the probabilities of presence $\pi(t) = \mu_0 P^t$ from a starting distribution μ_0 (equal to 1 for a single bundle and 0 for the others). We recorded the number of steps at which the variation distance goes below ε, for $\varepsilon = 0.1$ (in red) and $\varepsilon = 0.001$ (in blue), indicating convergence to the stationary distribution.

The actual number of steps required to get a precision of 0.1 or 0.001 is much lower than what the theoretical bound suggests, even for a very small ε. To give an idea, the bound is close to 1 billion steps for $n = 8$. It can be observed that in both cases convergence is reached under 200 steps on average, and that this grows only linearly with the group sizes. If we compare common or arbitrary utilities, we observe that the latter case takes s a slightly longer time on average.

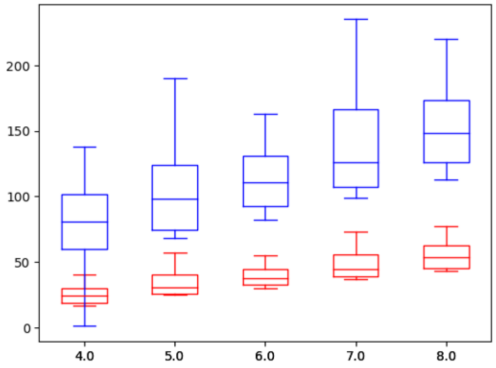

Figure 1. Number of steps before convergence as a function of n_p the group size (with type-uniform utilities).

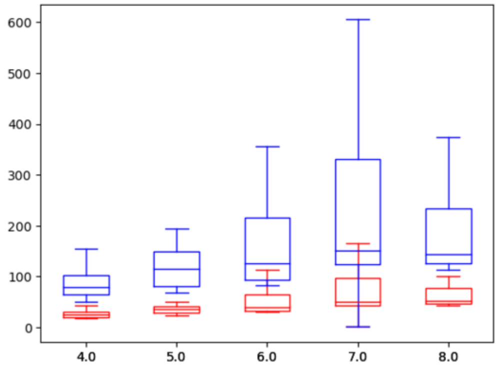

Figure 2. Number of steps before convergence as a function of n_p the group size (with general utilities).

In particular, some instances seem more challenging (note the max value), but even in the worst case, the computing time remains much below the theoretical bound. For these experiments, constructing explicitly the state space Ω is a barrier, but recall that in practice the approximation algorithm does not require it. This suggests that our algorithm can be used in practice for arbitrary utilities and groups of significant size.

8 Conclusion

We proposed a new notion of the degree of group envy that satisfies several relevant axioms in house allocation problems among groups of agents of different sizes. This notion is based on the counterfactual comparison of subgroups of the same size. Although we proved that computing the exact degree of group envy is computationally difficult, we proposed an approximation approach which is practicable as shown by our experiments. An obvious extension of this work would be to prove that a FPRAS exists beyond the case of common utilities within groups. On the axiomatic perspective, we would like to provide a characterization of our approach, which may require further axioms to be considered.

Acknowledgements

We acknowledge a financial support from the project THEMIS ANR-20-CE23-0018 of the French National Research Agency (ANR).

References

[1] Martin Aleksandrov and Toby Walsh, 'Group envy freeness and group pareto efficiency in fair division with indivisible items', in *KI 2018: Advances in Artificial Intelligence*, eds., Frank Trollmann and Anni-Yasmin Turhan, pp. 57–72, Cham, (2018). Springer International Publishing.

[2] Haris Aziz and Simon Rey, 'Almost group envy-free allocation of indivisible goods and chores', in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, ed., Christian Bessiere. International Joint Conferences on Artificial Intelligence Organization, (2020).

[3] Nawal Benabbou, Mithun Chakraborty, Edith Elkind, and Yair Zick, 'Fairness towards groups of agents in the allocation of indivisible items', in *The 28th International Joint Conference on Artificial Intelligence (IJCAI'19)*, (2019).

[4] Nawal Benabbou, Mithun Chakraborty, Xuan-Vinh Ho, Jakub Sliwinski, and Yair Zick, 'Diversity constraints in public housing allocation', in *17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2018)*, (2018).

[5] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick, 'Weighted envy-freeness in indivisible item allocation', *ACM Trans. Econ. Comput.*, **9**(3), (aug 2021).

[6] Chandra R. Chegireddy and Horst W. Hamacher, 'Algorithms for finding k-best perfect matchings', *Discrete Applied Mathematics*, **18**(2), 155–165, (1987).

[7] Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan, 'Group fairness for the allocation of indivisible goods', in *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1853–1860, (Jul. 2019).

[8] Bart de Keijzer, Sylvain Bouveret, Tomas Klos, and Yingqian Zhang, 'On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences', in *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT'09)*, Lecture Notes in Artificial Intelligence, pp. 98–110, Venice, Italy, (October 2009). Springer Verlag.

[9] Duncan K. Foley, 'Resource allocation and the public sector', *Yale Economic Essays*, **7**(1), 45–98, (1967).

[10] Venkatesan Guruswami, 'Rapidly mixing markov chains: A comparison of techniques (a survey)', *ArXiv*, **abs/1603.01512**, (2016).

[11] Mark Jerrum and Alistair Sinclair, 'The markov chain monte carlo method: An approach to approximate counting and integration', *Approximation Algorithms for NP-Hard Problems*, 482–520, (1996).

[12] Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani, 'Random generation of combinatorial structures from a uniform distribution', *Theoretical Computer Science*, **43**, 169–188, (1986).

[13] Maria Kyropoulou, Warut Suksompong, and Alexandros A. Voudouris, 'Almost envy-freeness in group resource allocation', *Theoretical Computer Science*, **841**, 110–123, (2020).

[14] Richard Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi, 'On approximately fair allocations of divisible goods', pp. 125–131, (2004).

[15] Pasin Manurangsi and Warut Suksompong, 'Almost envy-freeness for groups: Improved bounds via discrepancy theory', *Theoretical Computer Science*, **930**, 179–195, (sep 2022).

[16] Jonathan Scarlett, Nicholas Teh, and Yair Zick, 'For one and all: Individual and group fairness in the allocation of indivisible goods', *CoRR*, **abs/2302.06958**, (2023).

[17] Alistair Sinclair, 'Improved bounds for mixing rates of markov chains and multicommodity flow', *Combinatorics, Probability and Computing*, **1**(4), 351–370, (1992).

[18] William Thomson, 'On the axiomatic method and its recent applications to game theory and resource allocation', *Social Choice and Welfare*, **18**(2), 327–386, (2001).

[19] Leslie G. Valiant, 'The complexity of enumeration and reliability problems', *SIAM J. Comput.*, **8**, 410–421, (1979).

[20] L.G. Valiant, 'The complexity of computing the permanent', *Theoretical Computer Science*, **8**(2), 189–201, (1979).