



HAL
open science

Fair in the Eyes of Others

Parham Shams, Aurélie Beynier, Sylvain Bouveret, Nicolas Maudet

► **To cite this version:**

Parham Shams, Aurélie Beynier, Sylvain Bouveret, Nicolas Maudet. Fair in the Eyes of Others. 24th European Conference on Artificial Intelligence, Jun 2020, Santiago de Compostela, Spain. hal-02502545

HAL Id: hal-02502545

<https://hal.science/hal-02502545>

Submitted on 9 Mar 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fair in the Eyes of Others

Parham Shams¹ and Aurélie Beynier¹ and Sylvain Bouveret² and Nicolas Maudet¹

Abstract. Envy-freeness is a widely studied notion in resource allocation, capturing some aspects of fairness. The notion of envy being inherently subjective though, it might be the case that an agent envies another agent, but that she objectively has no reason to do so. The difficulty here is to define the notion of objectivity, since no ground-truth can properly serve as a basis of this definition. A natural approach is to consider the judgement of the other agents as a proxy for objectivity. Building on previous work by Parijs (who introduced “unanimous envy”) we propose the notion of approval envy: an agent a_i experiences approval envy towards a_j if she is envious of a_j , and sufficiently many agents agree that this should be the case, from their own perspectives. Some interesting properties of this notion are put forward. Computing the minimal threshold guaranteeing approval envy clearly inherits well-known intractable results from envy-freeness, but (i) we identify some tractable cases such as house allocation; and (ii) we provide a general method based on a mixed integer programming encoding of the problem, which proves to be efficient in practice. This allows us in particular to show experimentally that existence of such allocations, with a rather small threshold, is very often observed.

1 INTRODUCTION

Fair division is an ubiquitous problem in multiagent systems, economics [28, 20, 29], with applications ranging from allocation of schools, courses or rooms to students [1, 21], to division of goods in inheritance or divorce settlement [7]. Envy-freeness (EF), is one of the prominent notions studied in fair division [14, 6, 18, 11, 26]. An allocation of items among a set of agents is said to be envy-free if no agent prefers the share of another agent to her own share. Unfortunately, envy-freeness is a pretty demanding notion and an envy-free allocation may not exist.

Now consider a given problem where no envy-free allocation can be returned, but suppose instead that two allocations make a single agent (say, a_i) envious of some other agent a_j (for simplicity). Now assume that in allocation π , a_i is the only agent to prefer the bundle of a_j over her own, while in allocation π' all the other agents agree on the fact that a_i should indeed envy a_j . According to Parijs [22], π' exhibits *unanimous envy*, and there seems to be no situation where π' should be returned in place of π . Inspired by this notion, we introduce in this paper the notion of K -approval envy, as a way to retrieve a continuum between envy-freeness and unanimous envy. As may be clear from the name, the idea is simply to ask agents to express their own view about envy relations expressed by other agents. The objective will thus be to seek allocations minimizing social support for

the expressed envy relations, i.e. minimizing the number of agents K approving the envy. Of course, this approach may be controversial: after all, the notion of preference is inherently subjective. Introducing this flavour of objectivity may lead to undesirable consequences. At the extreme, one may simply replace individual preferences by some unanimous “mean” profile, thus profoundly changing the very nature of the notion. We believe there are several justifications to investigate this new approach:

- First, note that we only seek the approval of other agents in the case the agent herself explicitly expresses envy: absence of envy thus remains completely subjective. While a symmetrical treatment may also be justifiable in some situations, there is an obvious reason which motivates us to start with the proposed definition, namely the fact that the notion would no longer be a relaxation of envy-freeness.
- Secondly, all other things being equal, we believe an allocation minimizing K is socially more desirable. We do not necessarily regard this notion as a compelling choice, but we think this can enrich the picture of fallback allocations when no envy-free allocation exists, as other relaxations do [2].
- Finally, one further motivation of our work that we would like to emphasize is that our approach can be seen as providing guidance regarding agents and more specifically agents’ preferences which could be focused on, in order to progress towards envy-freeness. In particular, if we envision systems integrating deliberative phases in the collective decision-making process, our model could be used to set the agenda of such deliberations. If a vast majority of agents contradict an agent on her envy towards another agent, it may indicate for instance that she lacks information regarding the actual value of (some items of) her share. Initiating a discussion might help to solve such “objectively unjustified” envies when they occur.

Outline of the paper. The remainder of this paper is as follows. Section 2 recalls some basic notions in fair division. Our notion of K -approval envy is presented in Section 3. Some properties of this notion are then studied in Section 4: it is shown in particular, that if the hypothetical situation of allocation π described at the beginning of the introduction occurs, then an EF allocation must also exist. We also show that our notion inherits from the complexity of related problems. This motivates the MIP formulation that we detail in Section 5. We next turn to the House Allocation setting and we show that if each agent exactly holds a single item, then an efficient algorithm allows for returning an allocation minimizing the value of K . One caveat of our notion is that (unlike other relaxations) it is not guaranteed to exist, as intuitively observed in the case of unanimous envy. We thus consider greatly important to provide empirical evidence showing that both in different synthetic cultures as well as with real datasets, allocations with reasonable values of K exist.

¹ LIP6, Sorbonne Université, CNRS, F-75005 Paris, France, email: {parham.shams,aurelie.beynier,nicolas.maudet}@lip6.fr

² LIG, Université Grenoble Alpes, CNRS, Grenoble INP, Grenoble, France, email: sylvain.bouveret@imag.fr

2 MODEL AND DEFINITIONS

We consider MultiAgent Resource Allocation problems (MARA) where we aim at fairly dividing a set of indivisible goods (also called items or objects) among a set of agents. A MARA instance I is defined as a finite set of *objects* $\mathcal{O} = \{o_1, \dots, o_m\}$, a finite set of *agents* $\mathcal{N} = \{a_1, \dots, a_n\}$ and a profile \mathcal{P} of preferences representing the interest of each agent $a_i \in \mathcal{N}$ towards the objects. An allocation π is a mapping of the objects in \mathcal{O} to the agents in \mathcal{N} . In the following, π_i will denote the set of objects (the share) held by agent a_i . An allocation is such that $\forall a_i, \forall a_j$ with $i \neq j : \pi_i \cap \pi_j = \emptyset$ (a given object cannot be allocated to more than one agent) and $\bigcup_{a_i \in \mathcal{N}} \pi_i = \mathcal{O}$ (all the objects from \mathcal{O} are allocated).

In this paper, we consider cardinal preference profiles so, the preferences of an agent a_i over bundles of objects is defined by a *utility function* $u_i : 2^{\mathcal{O}} \rightarrow \mathbb{Q}^+$ measuring her satisfaction $u_i(\pi_i)$ when she obtains share π_i . We make the assumption that utility functions are additive *i.e.* the utility of an agent a_i over a share π_i is defined as the sum of the utilities over the objects forming π_i :

$$u_i(\pi_i) \stackrel{\text{def}}{=} \sum_{o_k \in \pi_i} u(i, k),$$

where $u(i, k)$ is the utility given by agent a_i to object o_k . This assumption is commonly considered in MARA [18, 24, 12, 9, for instance] as additive utility functions provide a compact but yet expressive way to represent the preferences of the agents. MARA instances with additive utility functions are called add-MARA instances for short.

Different notions have been proposed in the literature to evaluate the fairness of an allocation. When the agents can compare their shares, the absence of envy [14, 18, 10] is a particularly relevant notion of fairness. An agent a_i would envy another agent a_j if she prefers the share of a_j over her own share. More formally, an agent a_i envies an agent a_j iff

$$u_i(\pi_j) > u_i(\pi_i)$$

A completely fair allocation would thus be an *envy-free allocation* *i.e.* an allocation where no agent envies another agent. Formally:

$$\forall a_i, a_j \in \mathcal{N}, u_i(\pi_i) \geq u_i(\pi_j)$$

The notion of envy-freeness conveys a natural concept of fairness viewed as social stability: agents are happy with their bundle and hence would not want to swap it with any other agent's (regarding their own preferences). However, as soon as it is required to allocate all the objects in \mathcal{O} , an envy-free allocation may not exist. An alternative objective may be to minimize a degree of envy of the society [18, 10], based on the notion of *pairwise envy*.

Definition 1 (Pairwise envy). *Let π be an allocation. The pairwise envy $pe(i, j, \pi)$ of an agent a_i towards an agent a_j in π is defined as follows:*

$$pe(i, j, \pi) \stackrel{\text{def}}{=} \max\{0, u_i(\pi_j) - u_i(\pi_i)\}.$$

The pairwise envy can be interpreted as how much agent a_i envies agent a_j 's share (this envy being 0 if a_i does not envy a_j). We can derive from this notion a collective measure of envy:

Definition 2 (Degree of envy of the society). *The degree of envy of the society for an allocation π is defined as follows:*

$$de(\pi) \stackrel{\text{def}}{=} \sum_{a_i \in \mathcal{N}} \sum_{a_j \in \mathcal{N}} pe(i, j, \pi)$$

Note that an allocation π is envy-free if and only if $de(\pi) = 0$.

To cope with the possible inexistence of an envy-free allocation, another approach is to alleviate the requirements of the fairness notion. Recently, several relaxations of envy-freeness have been proposed such as envy-freeness up to one good (EF1) [8] or envy-freeness up to any good (EFX) [9]. An allocation is said to be envy-free up to one good (resp. up to any good) if no agent a_i envies the share π_j of another agent a_j after removing from π_j *one* (resp. any) item. Existence for EF1 is guaranteed, but this is still to the best of our knowledge an open question for EFX. Amanitidis et al. [2] studied the relations between some fairness notions and their relaxations.

3 K-APPROVAL ENVY

The notion of envy being inherently subjective, it might be the case that an agent envies another agent, but that she objectively has no reason to do so. The difficulty here is to define the notion of objectivity, since no ground-truth can properly serve as a basis of this definition. In her book, Guibet-Lafaye [17] recalls the notion of *unanimous envy*, that was initially discussed in the book by Parijs [22], and that can be defined as follows: an agent a_i unanimously envies another agent a_j , if all the agents think that a_i indeed envies a_j . Here, unanimity is used as a proxy for objectivity.

As we can easily imagine, looking for allocations that are free of unanimous envy will be too weak to be interesting: as soon as one agent disagrees with the fact that a_i envies a_j , this potential envy will not be taken into account. Here, we propose an intermediate notion between envy-freeness and (unanimous envy)-freeness:

Definition 3 (K -approval envy). *Let π be an allocation, a_i, a_j be two different agents, and $1 \leq K \leq n$ be an integer. We say that a_i K -approval envies (K -app envies for short) a_j if there is a subset \mathcal{N}_K of K agents including a_i such that:*

$$\forall a_k \in \mathcal{N}_K, u_k(\pi_i) < u_k(\pi_j).$$

In other words, at least $K - 1$ agents amongst $\mathcal{N} \setminus \{a_i\}$ agree with a_i on the fact that she should actually envy a_j .

Example 1. *Let us consider the following add-MARA instance with 3 agents and 6 objects:*

| | o_1 | o_2 | o_3 | o_4 | o_5 | o_6 |
|-------|-------|-------|-------|-------|-------|-------|
| a_1 | 0 | 3 | 3 | 1 | 3 | 2 |
| a_2 | 2 | 0 | 7 | 2 | 1 | 0 |
| a_3 | 0 | 3 | 5 | 0 | 1 | 3 |

Note that there is no envy-free allocation for this instance. In the squared allocation, a_1 is not envious, a_2 envies a_3 and a_3 envies a_1 . Concerning the envy of a_2 towards a_3 , a_1 disagrees with a_2 being envious of a_3 whereas agent a_3 agrees. Hence, agent a_2 2-app envies agent a_3 . Concerning the envy of a_3 towards a_1 , agent a_1 agrees with a_3 being envious of a_1 whereas agent a_2 does not. Hence, a_3 2-app envies a_1 .

Note that in the definition, as soon as a_i does not envy a_j , then, a_i does not K -app envy a_j , no matter what the value of K is or how many agents think that a_i should actually envy a_j .

Let us start with an easy observation:

Observation 1. *Given an allocation π of an add-MARA instance, if a_i K -app envies a_j in π , then a_i $(K-1)$ -app envies a_j in π .*

Moreover, if a_i n -app envies a_j , we will say that a_i *unanimously* envies a_j . Finally, we can observe that a_i 1-app envies a_j if and only if a_i envies a_j .

We can naturally derive from Definition 3 the counterpart of EF:

Definition 4 ($(K$ -approval envy)-free allocation). *An allocation π is said to be $(K$ -app envy)-free if and only if a_i does not K -app envy a_j for all pairs of agents (a_i, a_j) .*

Definition 5 ($(K$ -approval envy)-free instance). *An add-MARA instance I will be said to be $(K$ -app envy)-free if and only if it accepts a $(K$ -app envy)-free allocation.*

Example 2. *Going back to Example 1, the squared allocation is (3-app envy)-free so the instance is (3-app envy)-free.*

A threshold of special interest is obviously $\lfloor n/2 \rfloor + 1$, since it requires a strict majority to approve the envy under inspection. A Strict Majority approval envy-free (SM-app-EF) allocation is a $(K$ -app envy)-free allocation such that $K \leq \lfloor n/2 \rfloor$, translating the fact that every time envy occurs, there is a strict majority of agents that do not agree with that envy.

Going further, it is important to notice that $(K$ -app envy)-freeness is not guaranteed to exist. Indeed, for all number of agents n and all number of objects m , there exist instances for which no $(K$ -app envy)-free allocation exists, no matter what K is. Suppose for instance that all the agents rank the same object (say o_1) first, and that for all a_i , $u(i, 1) > \sum_{k=2}^m u(i, k)$. Then obviously, everyone agrees that all the agents envy the one that will receive o_1 . Such instances will be called *unanimous envy instances*:

Definition 6 (Unanimous envy instance). *An add-MARA instance I will be said to exhibit unanimous envy if I is not $(K$ -app envy)-free for any value of K .*

Observe that for an allocation to be $(K$ -app envy)-free, for all pairs of agents a_i, a_j , either a_i or at least $n - K + 1$ agents have to think that a_i does not envy a_j . Notice that it is different from requiring that at least K agents think that this allocation is envy-free. This explains the parenthesis around $(K$ -app envy): this notion means “free of K -app envy”, which is different from “ K -app-(envy-free)”.

A useful representation, for a given allocation, is the induced envy graph: vertices are agents, and there is a directed edge from a_i to a_j if and only if a_i envies a_j [18]. An allocation is envy-free if and only if the envy graph has no arc. In our context, we can define a weighted notion of the envy graph.

Definition 7 (Weighted envy graph). *The weighted envy graph of an allocation π is defined as the weighted graph (\mathcal{N}, E) where nodes are agents, such that there is an edge $(a_i, a_j) \in E$ if a_i envies a_j , with the weight $w(a_i, a_j)$ corresponding to the number of agents (including a_i) approving this pairwise envy in π .*

Our notion of K -approval envy can be interpreted as a vote on envy, that works as follows. For each pair of agents (a_i, a_j) , if a_i declares to envy a_j , we ask the rest of the agents to vote on whether they think that a_i indeed envies a_j . Then, a voting procedure is used to determine whether a_i envies a_j according to the society of agents. Several voting procedures can be used. However, since there are only two candidates (yes / no), the most reasonable voting rules are based on quotas: a_i envies a_j if and only if there is a minimum quota of agents that think so.³ This makes a connection with related work [26] which uses voting to decide upon envy-freeness, but in the context of fair division of resources *jointly* owned by *groups* of agents.

³ More precisely, these rules exactly characterize the set of anonymous and monotonic voting rules [23].

4 SOME PROPERTIES OF K -APP ENVY

There are natural relations between the different notions of $(K$ -app envy)-freeness, for different values of K . The following observation is a direct consequence of Observation 1.

Observation 2. *Let π be an allocation, and $K \leq N$ be an integer. If π is $(K$ -app envy)-free, then π is also $((K+1)$ -app envy)-free.*

However, the converse does not hold. More precisely, the following proposition shows that the implication stated in Observation 2 is strict.

Proposition 1. *Let π be an allocation, and $3 \leq K \leq n$ be an integer. If π is $(K$ -app envy)-free, π is not necessarily $((K-1)$ -app envy)-free.*

Proof. Let $h \in \{2, \dots, n-1\}$ be an integer, and let us consider the instance with n agents and n objects defined as follows:

- $u(1, 1) = 1$;
- $u(i, 1) = u(i, i) = \frac{1}{2}$ for $i \in \{2, \dots, h-1\}$;
- $u(i, i) = 1$ for $i \in \{h, n-1\}$;
- $u(n, 1) = \frac{2}{n+1}$ and $u(n, j) = \frac{1}{n+1}$ for $j > 1$;

and $u(i, j) = \varepsilon$ for other pairs with $\varepsilon < \frac{1}{n+1}$.

Consider the allocation π where each agent a_i gets item o_i . Obviously, the only envy in this allocation concerns a_n towards a_1 . Moreover, only a_1, \dots, a_{h-1} agree on this envy. Therefore, a_n h -app envies a_1 , but does not $(h+1)$ -app envy her. Moreover, π is $((h+1)$ -app envy)-free, but not $(h$ -app envy)-free. \square

Example 3. *In order to illustrate the previous proof, let us consider the following instance with 4 agents, 4 objects (and $h=3$) and the squared allocation π :*

| | o_1 | o_2 | o_3 | o_4 |
|-------|---------------|---------------|---------------|---------------|
| a_1 | 1 | ε | ε | ε |
| a_2 | $\frac{1}{2}$ | $\frac{1}{2}$ | ε | ε |
| a_3 | ε | ε | 1 | ε |
| a_4 | $\frac{2}{5}$ | $\frac{1}{5}$ | $\frac{1}{5}$ | $\frac{1}{5}$ |

In this allocation, the only envy concerns a_4 towards a_1 . Moreover, only a_1 and a_2 agree with a_4 on her envy. Hence, π is (4-app envy)-free but is obviously not (3-app envy)-free as we can find 3 agents (a_1, a_2 and a_4) agreeing on the envy of a_4 towards a_1 (in other words a_4 3-app envies a_1).

Proposition 2. *For any $K \geq 3$, there exist instances which are $(K$ -app envy)-free but not $((K-1)$ -app envy)-free.*

Proof. Consider the same instance as in Proposition 1. We have already shown that we have an allocation π that is $((h+3)$ -app envy)-free which means that the instance is $((h+3)$ -app envy)-free. We just have to show that there is no $((h+2)$ -app envy)-free allocation in order to conclude. In that purpose, we first note that each agent has to get one and exactly one object. Indeed, if it is not the case at least one agent a_i will have no object and will thus be envious of any agent a_j that has an object. Moreover, as all agents value the empty bundle with utility 0 and every object is valued with a strictly positive utility, this envy will be unanimous. Hence, each agent has to get one and exactly one object in order to minimize the $(K$ -app envy)-freeness. Now consider objects o_j for $j \in \{h+2, n\}$. The agents a_j

that receive an object o_j and that are envious will $(h+2)$ -app envy the agent that received o_1 . Indeed, agents a_i for $i \in \{h+2, n-1\}$ value objects o_j with a utility higher than (or equal to) the one of o_1 (and thus do not approve the envy) while it is the opposite for the other agents who are exactly $h+2$ hence the $(h+2)$ -app envy. So if we want to avoid that envy, we have to give the objects o_j to agents so that they do not experience envy at all but it is not possible as such agents are agents a_p for $p \in \{h+2, n-1\}$. It means that we have $n-1-(h+2)+1$ agents that have to receive one of the $n-(h+2)+1$ objects which is obviously impossible. This means that we cannot avoid $(h+2)$ -app envy which implies that no allocation can be $((h+3)$ -app envy)-free. \square

Proposition 2 proves that the hierarchy of K -app envy instances is strict for $K \geq 3$. Rather surprisingly, we will see that it is not the case for $K = 2$.

In order to show this result, we will resort to a tool that has been proved to be really useful and powerful in many contexts dealing with envy [5, 3, 4]: the “bundle reallocation cycle technique”. This technique, originating from the seminal work of Lipton *et al.* [18], consists in performing a cyclic reallocation of *bundles* so that every agent is strictly better in the new allocation. Thus, such a reallocation corresponds to a cycle in the opposite direction of the edges in the — weighted — envy graph introduced in Definition 7. It is known that performing a reallocation cycle decreases the degree of envy [18]. Unfortunately, our first remark is that it does not necessarily decrease the level of K -app envy. Worse than that, it can actually increase it:

Proposition 3. *Let π be a $(K$ -app envy)-free allocation, for $3 \leq K \leq n-1$. After performing an improving bundle reallocation cycle (even between two agents), the resulting allocation may be $(K'$ -app envy)-free (and not $(K$ -app envy)-free) such that $K' > K$.*

Proof. Let $h \in \{0, \dots, n-4\}$ be an integer, and let us consider the instance with n agents and n objects defined by the following utility functions:

- $a_1: u(1, 1) = 1, u(1, 2) = 2, u(1, 3) = 7;$
- $a_2: u(2, 1) = 2, u(2, 2) = 1;$
- $a_3: u(3, 3) = 10;$
- a_l for $l \in \{4, h+3\}: u(l, 1) = u(l, 3) = 5, u(l, j) = 6$ for $j \geq 4;$
- a_m for $m \in \{h+4, n\}: u(m, 2) = 4, u(m, 3) = 5, u(m, i) = 6$ for $i \geq 4;$

and $u(i, j) = 0$ for other pairs.

Consider the allocation π where each agent a_i gets item o_i . Obviously, the only envy in this allocation concerns a_1 towards a_2 (approved by a_1 and agents a_m) and a_3 (approved by a_1, a_3 and agents a_m), and the envy of a_2 towards a_1 (approved by a_2 and agents a_l). Hence the allocation is $((\max\{|a_m|+3, |a_l|+2\})$ -app envy)-free. We now consider the allocation π' resulting from the improving bundle reallocation cycle between a_1 and a_2 . We note that the only envy in π' is the one of a_1 towards a_3 . Moreover, this envy is approved by herself, a_3 and agents a_l and a_m . The allocation is thus $((|a_m|+|a_l|+3)$ -app envy)-free and not $((\max\{|a_m|+3, |a_l|+2\})$ -app envy)-free as if $|a_l| \geq 1$ then $|a_m|+|a_l|+3 > \max\{|a_m|+3, |a_l|+2\}$. \square

Now consider a slight generalization of Lipton’s cycles, *weakly improving cycles* (WIC), that correspond to a reallocation of bundles where all the agents in the cycle receive a bundle they like at least as much as the one they held, with one agent at least being strictly happier. Of course, our example of Proposition 3 still applies. On the other hand, this notion suffices to guarantee the decrease of the

degree of envy (note that *identifying* the cycles themselves may not be easy any longer, but this is irrelevant for our purpose). The proof, omitted, follows directly from the arguments of Lipton [18].

Observation 3. *Let π be an allocation, and π' the allocation obtained after performing a weakly improving cycle. It holds that $de(\pi') < de(\pi)$.*

We now show that $(2$ -app envy)-freeness exhibits a special behaviour: in contrast with Proposition 3, improving cycles (in fact, even weakly improving cycles) enjoy the property of preserving the $(2$ -app envy)-freeness level of an allocation. We provide this result for *swaps* (cycles involving two agents only) as this is sufficient to establish our main result.

Lemma 1. *Let π be a $(2$ -app envy)-free allocation that is not EF. There always exists a WIC (that we can identify) between two agents such that the resulting allocation is $(K'$ -app envy)-free, with $K' \leq 2$.*

Proof. Let a_i be an envious agent (there is at least one). We identify the agent that a_i envies the most and call her a_j (if there are several agents that a_i envies the most, we can pick randomly one of them). Swapping the bundle of a_i and a_j is a WIC (as a_i envies a_j and a_j necessarily does not agree on this envy because otherwise it would contradict $(2$ -app envy)-freeness of π). We can prove that this swap leads to a $(K'$ -app envy)-free allocation with $K' \leq 2$. Let us call π' the allocation after this swap. In π' , all the agents except a_i and a_j have the same approval envy. Moreover, a_i is now EF in π' as she has received her preferred bundle. So, if π' is $(K'$ -app envy)-free with $K' > 2$, it is because a_j 2 -app envies (at least) some agent a_h (that can obviously not be a_i). Suppose for the sake of contradiction that it is the case. For this to be the case, a_j has to envy a_h and another agent a_l has to approve this envy: (1) $u_j(\pi'_j) < u_j(\pi'_h)$, (2) $u_l(\pi'_j) < u_l(\pi'_h)$. However, as a_i envies a_j in π then (3) $u_i(\pi_i) < u_i(\pi_j)$ and as π is $(2$ -app envy)-free and (3) holds, every agent a_l (except a_i of course) verifies (4) $u_l(\pi_i) \geq u_l(\pi_j)$.

Besides, π' is obtained after swapping the bundles of a_i and a_j in π so $\pi'_j = \pi_i, \pi'_i = \pi_j$ and $\pi'_h = \pi_h$; and from (2) we get: (5) $u_l(\pi_i) < u_l(\pi_h)$. By transitivity with (5) and (4), we obtain: (6) $u_l(\pi_j) < u_l(\pi_h)$. However, we know that a_j has the same utility in π and π' so $u_j(\pi'_j) = u_j(\pi_j)$. The latter combined with (1) (and the fact that $\pi'_h = \pi_h$) gives: (7) $u_j(\pi_j) < u_j(\pi_h)$. Finally, note that (6) and (7) translate the fact that a_j 2 -app envies a_h in π which contradicts the fact that π is $(2$ -app envy)-free. \square

Putting together Lemma 1 and Observation 3 allows us to prove that $(2$ -app envy)-freeness is essentially a vacuous notion, in the sense that any instance enjoying an allocation with this property will have an EF allocation as well.

Proposition 4. *If an add-MARA instance is $(2$ -app envy)-free then it is also envy-free.*

Proof. Take π as being an arbitrary $(2$ -app envy)-free allocation. First note that if there is no envious agent in π then, by definition, π is envy-free and the proposition holds. We perform a WIC leading to π' that is still $(2$ -app envy)-free (see Lemma 1). If π' is envy-free then we are done. Otherwise, thanks to Observation 3 we know the degree of envy has strictly decreased. We can repeat this process until the current allocation is EF. The process is guaranteed to stop because the degree of envy of the society is bounded below by zero and the degree of envy of the society strictly decreases at each step until it equals zero (which corresponds to an envy-free allocation). \square

Another consequence is that, for two agents, instances fall either in the envy-free or unanimous envy category:

Corollary 1. *In the special case of 2 agents, if there is no envy-free allocation in I then I is a unanimous envy instance.*

Complexity. We conclude with a few considerations on the computational complexity of the problems mentioned so far. First of all, as envy-freeness is (1-app envy)-freeness, the problem of finding the minimum K for which there exists a (K -app envy)-free allocation is at least as hard as determining whether an envy-free allocation exists.

One may also wonder how hard the problem of determining whether a given instance exhibits unanimous envy or not, *i.e.* whether a (K -app envy)-free allocation exists for *some* value of K . For this question, instances where agents all have the same preferences provide insights.

Proposition 5. *For any add-MARA instance, if all the agents have the same preferences then the notions of (1-app envy)-freeness and (n -app envy)-freeness coincide.*

Proof. We already know from Observation 2 that (1-app envy)-freeness implies (n -app envy)-freeness for any add-MARA instance. So we just have to prove that if all the agents have the same preferences then (n -app envy)-freeness implies (1-app envy)-freeness. If an allocation π is (n -app envy)-free then it means that for any pair a_i, a_j of agents, a_i does not envy a_j or there is at least one agent a_h that disagrees on the envy of a_i towards a_j . Obviously, if for every pair of agents a_i, a_j we have a_i envy-free towards a_j then the allocation π is envy-free and the proof concludes. Besides, for every pair of envious/envied agents there is at least one agent disagreeing on the envy. But all the agents have the same preferences so it means that every agent should agree with each other. Hence, no envied agent can exist and we have (1-app envy)-freeness of allocation π . \square

From Proposition 5 we get that the problem of deciding the existence of unanimous envy is at least as hard as deciding the existence of an EF allocation when agents have similar preferences which is known to be NP-hard [18]. As membership in NP is direct, we thus get as a corollary that:

Corollary 2. *Deciding whether an allocation exhibits unanimous envy is NP-Complete.*

5 A MIP FORMULATION FOR K -APP ENVY

We have seen in the previous section that the problem of determining, for a given instance I , the minimal value of K such that a (K -app envy)-free allocation exists inherited from the high complexity of determining whether an envy-free allocation exists.

To address this problem, we present in this section a Mixed Integer linear Program that returns, for a given add-MARA instance I , a (K -app envy)-free allocation with the minimal K and no solution when I is a unanimous envy instance. In this MIP, we use $n \times m$ Boolean variables \mathbf{z}_i^j (we use bold letters to denote variables) to encode an allocation: $\mathbf{z}_i^j = 1$ if and only if a_i gets item o_j . We also introduce n^3 Boolean variables \mathbf{e}_{kih} such that $\mathbf{e}_{kih} = 1$ if and only if according to a_k 's preferences a_i envies agent a_j . We also need to add n^2 Boolean variables \mathbf{x}_{ih} used to linearize the constraints on \mathbf{e}_{kih} . Finally, we use an integer variable \mathbf{K} corresponding to the K -app envy we seek to minimize.

In this section, we assume that all the utilities are integers. If they are not (recall that they are still in \mathbb{Q}^+) we can transform the instance

at stake into a new one only involving integral utilities by multiplying them by the least common multiple of their denominators.

We first need to write the constraints preventing an item from being allocated to several agents:

$$\sum_{i=1}^n \mathbf{z}_i^j = 1 \quad \forall j \in \llbracket 1, m \rrbracket \quad (1)$$

By adding these constraints we also guarantee completeness of the returned allocation (all the items have to be allocated to an agent).

Secondly, we have to write the constraints that link the variables \mathbf{e}_{kih} with the allocation variables \mathbf{z}_i^j :

$$\forall k, i, h \in \llbracket 1, n \rrbracket, \sum_{j=1}^m u(k, j)(\mathbf{z}_h^j - \mathbf{z}_i^j) > 0 \iff \mathbf{e}_{kih} = 1$$

As the utilities are integers, we can replace > 0 by ≥ 1 . In order to linearize the implication from left to right (\implies) we introduce a number M that can be arbitrarily chosen such that $M > \max_k \sum_{j=1}^m u(k, j)$:

$$M \mathbf{e}_{kih} \geq \sum_{j=1}^m u(k, j)(\mathbf{z}_h^j - \mathbf{z}_i^j) \quad \forall k, i, h \in \llbracket 1, n \rrbracket \quad (2)$$

$$\sum_{j=1}^m u(k, j)(\mathbf{z}_h^j - \mathbf{z}_i^j) \geq 1 - M(1 - \mathbf{e}_{kih}) \quad \forall k, i, h \in \llbracket 1, n \rrbracket \quad (3)$$

Note that we do not need to formalize the implication from right to left because the optimization criterion will ‘force’ the \mathbf{e}_{kih} to be as small as possible. Finally, we have to write the constraints that convey the fact that the allocation we look for is (K -app envy)-free:

$$\mathbf{e}_{i ih} = 0 \vee \sum_{k=1}^n \mathbf{e}_{kih} \leq \mathbf{K} - 1 \quad \forall i, h \in \llbracket 1, n \rrbracket$$

Since $\mathbf{e}_{i ih}$ are Boolean variables, we can replace $\mathbf{e}_{i ih} = 0$ by $\mathbf{e}_{i ih} \leq 0$. Now, this logical constraint is linearized as follows:

$$\mathbf{e}_{i ih} \leq \mathbf{x}_{ih} \quad \forall i, h \in \llbracket 1, n \rrbracket \quad (4)$$

$$\sum_{k=1}^n \mathbf{e}_{kih} \leq \mathbf{K} - 1 + n(1 - \mathbf{x}_{ih}) \quad \forall i, h \in \llbracket 1, n \rrbracket \quad (5)$$

We can now put things together. Let I be an instance. Then, we will denote by $\mathcal{M}(I)$ the MIP defined as:

$$\begin{aligned} & \text{minimize} && \mathbf{K} \\ & \text{such that} && \mathbf{z}_i^j, \mathbf{e}_{kih}, \mathbf{x}_{ih} \in \{0, 1\} \forall k, i, h \in \llbracket 1, n \rrbracket, j \in \llbracket 1, m \rrbracket \\ & && \mathbf{K} \in \llbracket 1, N \rrbracket \\ & && + \text{Constraints (1, 2, 3, 4, 5)} \end{aligned}$$

Proposition 6. *Let I be an instance. Then, there is an optimal solution with $\mathbf{K} = L$ to $\mathcal{M}(I)$ if and only if I is an (L -app envy)-free instance and not an $((L-1)$ envy)-free one. Moreover, $\mathcal{M}(I)$ does not admit any solution if and only if I is a unanimous envy instance.*

The proof of this proposition is not very involved and will thus be omitted. The key here is to show that there is a solution to the MIP $\mathcal{M}(I)$ such that $\mathbf{K} = L$ iff the corresponding allocation π such that $\mathbf{z}_i^j = 1$ if and only if $o_j \in \pi_i$ is (L -app envy)-free. The most critical point is to show that Constraints 2 and 3 are indeed a valid translation of the logical implication, and that Constraints 4 and 5 correctly encode the logical or. The rest follows easily.

As the problem is difficult in the general case, it is natural to seek special cases that could be solved efficiently.

6 HOUSE ALLOCATION

The House Allocation Problem (HAP for short) is a standard problem where there are exactly as many items as agents, and each agent receives exactly one resource. This setting is relevant in many situations and has been extensively studied [27, 25, 1, to cite a few of them]. In House Allocation Problems, computing an envy-free allocation comes down to solving a matching problem, since an envy-free allocation exists if and only if all the agents get (one of) their top item(s). It is therefore natural to wonder whether an allocation minimizing K -app envy could also be computed efficiently.

Our first observation hints in that direction. Indeed, characterizing unanimous envy becomes easy in house allocation problems.

Proposition 7. *Let I be an instance of HAP. I is an unanimous envy instance if and only if there exists at least a pair of items (o_i, o_j) such that all agents strictly prefer o_i over o_j .*

Corollary 3. *Checking whether an instance I of HAP is a unanimous envy instance or not can be done in $O(n^2)$.*

From this characterization we can also derive a result on the likelihood that unanimous envy exists when the utilities are uniformly distributed (that is, for each agent a_i and object o_j , utilities are drawn i.i.d. following the uniform distribution on some interval $[x, y]$).

Proposition 8. *Under uniformly distributed preferences, the probability of unanimous envy is upper bounded by $n(n-1)/2^n$.*

Proof. Wlog. suppose agent 1 has preferences $o_1 \succ o_2 \succ \dots \succ o_n$. The probability of the event o_i is strictly preferred to o_j by one agent is $1/2$ if preferences are strict. As preferences are not strict, this probability becomes an upper bound (think for instance if the agent values all the objects the same then the probability to have strict preference between two objects is zero). Hence, the probability of the event o_i is strictly preferred to o_j by all agents is upper bounded by $1/2^{n-1}$ as the preferences between the agents are independent. Assuming, for all pairs of items, these events to be independent (which is not the case, hence an upper bound of the upper bound), we derive our result by summing up over the $n(n-1)/2$ possible pairs. \square

Note that this value quickly tends towards 0: unanimous envy is thus already unlikely to occur for 10 agents.

We will now show here that finding an allocation that minimizes $(K$ -app envy)-freeness can be done in polynomial time. Before introducing the idea, we need an additional notation. For any pair (j, j') , let $\#_{\prec}(j, j')$ denote the number of agents strictly preferring $o_{j'}$ to o_j . For any agent a_i and object o_j , we will also define $\text{maxEnvy}[i][j]$ as follows:

$$\text{maxEnvy}[i][j] = \max_{o_{j'} \text{ s.t. } u(i, j') > u(i, j)} \#_{\prec}(j, j')$$

In other words, $\text{maxEnvy}[i][j]$ denotes the maximal value of $\#_{\prec}(j, j')$ among the objects that are strictly preferred to o_j by a_i . As we can imagine, this will exactly be the value of the K -app envy experienced by a_i if she gets item o_j (note that if o_j is among a_i 's top objects, this value will be 0).

The key to the algorithm is to see that for a given K , determining whether a $(K$ -app envy)-free allocation exists can be done in polynomial time by solving a matching problem. Namely, for each K , we build the following bipartite graph: $\mathcal{N} \cup \mathcal{O}$ is the set of nodes, and we add an edge $(a_i, o_j) \in \mathcal{N} \times \mathcal{O}$ if and only if $\text{maxEnvy}[i][j]$ is lower

than or equal to K . We can observe that any perfect matching in this graph corresponds to a $((K+1)$ -app envy)-free allocation. The only thing that remains to do is to run through all possible values of K , which can be done by dichotomous search between 0 and n . This is formalized in Algorithm 1.

Algorithm 1: Minimizing $(K$ -app envy)-freeness in the HAP

```

input :  $I = (\mathcal{N}, \mathcal{O}, w)$  a HAP instance
1  $\text{maxEnvy} \leftarrow \text{computeMaxEnvy}()$ ;
2  $\text{res} \leftarrow \text{None}$ ;
3  $\text{low} \leftarrow 0, \text{high} \leftarrow n$ ;
4 while  $\text{low} \leq \text{high}$  do
5    $i \leftarrow \lfloor (\text{low} + \text{high})/2 \rfloor$ ;
6    $G \leftarrow \text{buildBipartiteGraph}(\text{maxEnvy}, i)$ ;
7    $\pi \leftarrow \text{perfectMatching}(G)$ ;
8   if  $\pi$  is not None then
9      $\text{res} \leftarrow \pi, i + 1$ ;
10     $\text{high} \leftarrow i - 1$ ;
11  else
12     $\text{low} \leftarrow i + 1$ ;
13 return  $\text{res}$ 

```

Proposition 9. *For any HAP instance, we can find (one of) its optimal $(K$ -app envy)-free allocations in $O(n^3 \log(n))$.*

Proof. First, the computation of the matrix maxEnvy runs in $O(n^3)$. Indeed, to compute $\text{maxEnvy}[i][j]$ we first need to compute $\#_{\prec}(j, j')$ which already runs in $O(n^3)$ as we have to ask for each couple of objects (n^2 in total) the point of view of all the agents (n in total). From that, as $\text{maxEnvy}[i][j] = \max_{o_{j'} \text{ s.t. } u(i, j') > u(i, j)} \#_{\prec}(j, j')$, we can compute $\text{maxEnvy}[i][j]$ in $O(n)$. As there are n^2 different pairs (a_i, o_j) we have the final $O(n^3)$ complexity of computing maxEnvy .

Due to the dichotomous search, the algorithm needs to solve $\log(n)$ perfect matching problems, that can be solved in $O(n^3)$ [16]. The global complexity of Algorithm 1 is thus $O(n^3 \log(n))$. \square

7 EXPERIMENTAL RESULTS

We present here the results of the numerical tests we have conducted. These experiments serve two purposes: (i) evaluate the behaviour of the MIP we presented in Section 5 and of the polynomial algorithm described in Section 6, and (ii) observe how our notion of K -app envy depends on the number of agents, of items, and on the type of preferences. All the tests presented in this section have been run on an Intel(R) Core(TM) i7-2600K CPU with 16GB of RAM and using the Gurobi solver to solve the Mixed Integer Program. We have tested our methods on three types of instances: Spliddit instances [15], instances under uniformly distributed preferences and instances under an adaptation of Mallows distributions to cardinal utilities [13].

7.1 Spliddit instances

We have first experimented our MIP on real-world data from the fair division website Spliddit [15]. There is a total of 3535 instances from 2 agents to 15 agents and up to 93 items. Note that 1849 of these instances involve 3 agents and 6 objects. By running the MIP with a timeout of 10 minutes (after this duration the best current solution, if it exists, is returned) we were able to solve all the instances but

6. Among these 6 instances, 3 of them were HAP instances that we managed to solve optimally with Algorithm 1. This only leaves us with 3 instances, for which the solver did return a solution but did not prove that it is optimal (within a timeout of 10 minutes). Besides, 65% of the instances are EF while 23% of the instances exhibit unanimous envy. Moreover, 28% of the remaining instances with more than 5 agents are SM-app EF.

7.2 Uniformly distributed preferences

General setting. We also ran tests on instances under uniformly distributed preferences, with n varying from 3 to 10 and m such that we produce settings where few EF allocations exist [12]. For each problem size, we kept 60 instances that admit no EF allocation as we wanted to measure the behaviour of our notion when no such allocation exists (we know that if an EF allocation exists it will be returned by our methods). As we are in the general setting we solved the instances via the MIP with a timeout of 60 seconds.

The first three rows of Table 1 respectively represent the percentage of instances that have been solved to optimal (a solution has been returned before the timeout), the percentage of unanimous envy instances and the percentage of SM-app-EF instances. We then have the mean value of K/n . Finally, we store the mean computation time (in seconds) of the instances (solved to optimal).

Table 1. Results of the experiments as a function of the number of agents.

| n | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---------------|---------------|-------|------|------|------|-------|-------|-------|------|
| % OPT | 100 | 100 | 100 | 100 | 100 | 68.3 | 1.7 | 1.7 | 0 |
| % UEI | 100 | 21.7 | 5 | 0 | 0 | 0 | 0 | 0 | 0 |
| % SMAEF | 0 | 0 | 0 | 50 | 50 | 75 | 40 | 33.3 | 6.7 |
| mean(K/n) | NaN | 1 | 0.85 | 0.72 | 0.61 | 0.57 | 0.59 | 0.63 | 0.66 |
| time(s) | ε | 0.008 | 0.04 | 0.21 | 1.97 | 21.29 | 50.09 | 56.16 | NaN |

First note that considering 2 agents is a special case as shown in Corollary 1. Indeed, as we have removed the EF instances, all the remaining instances are unanimous envy ones. Moreover, we observe that the percentage of SM-app-EF allocations is zero for 4 agents. Indeed, an allocation is SM-app-EF for 4 agents if there exists a (K -app envy)-free allocation such that $K \leq 2$. As we have removed all the EF instances, we know (from Proposition 4) that we cannot find an SM-app-EF allocation. The same holds for 3 agents.

We can notice that the mean K/n seems to be stabilising around 0.6. Besides, without any surprise, the computation time rapidly increases while the percentage of instances solved to optimal (under a timeout of 60 seconds) starts decreasing for 7 agents. Finally, positive results can be pinpointed: the very low percentage of unanimous envy instances, and the pretty high percentage of SM-app-EF ones.

House allocation. We have also tested our polynomial algorithm on HAP instances under uniformly distributed preferences. We have generated 20 instances for each number of agents from 5 to 100 agents (and objects) by steps of 5.

First note that we have only found 5 unanimous envy instances and all of them involved 5 agents. This supports the probability of unanimous envy instance showed in Proposition 8 and the fact that it decreases very quickly towards 0. Moreover, like for the general setting, we notice a convergence of the K/n values towards 0.6. The algorithm runs, without any surprise (in light of Proposition 9) much faster than our MIP. Indeed, the mean runtime for 100 objects and agents is still around 2 seconds only whereas we already observed that our MIP cannot solve easier problems within 10 minutes.




Figure 1. Optimal K/n in the HAP as a function of m, n

7.3 Correlated preferences

In strict ordinal settings, a classical way to capture correlated preferences is to use Mallows distributions [19] allowing us to measure the impact of the similarity of the preferences between agents. In these experiments, we used a generalization of the Mallows distribution to cardinal preferences presented in [13] based on Von Mises–Fisher distributions. Similarly to the dispersion parameter in Mallows distributions, the similarity between the preferences of the agents is tuned by the *concentration* parameter: when it is zero agents’ preferences are uniformly distributed, whereas when it is infinite agents have the same preferences.

We expected that the more similar the preferences between the agents are, the higher the degree of K -app envy would be and the more likely unanimous envy would occur. The results of our experiments both in the general setting and in HAP support this: the number of EF instances is decreasing along with the concentration value, and from a given threshold, all the instances exhibit unanimous envy. However, the exact correlation between the level of (K -app) envy-freeness and the concentration deserves further study, especially for very low values of K . Intuitively, in some circumstances, correlation of preferences may indeed help to find large majorities of agents that contradict an agent envy, while this situation is unlikely under uniformly distributed preferences.

8 CONCLUSION

In this paper, we have introduced a new relaxation of envy-freeness. This relaxation uses a consensus notion, approval envy, as a proxy for objective envy between pairs of agents. We have proposed algorithms to compute an allocation minimizing the K -app envy, and we have experimentally shown that this notion makes sense in practice in situations where no envy-free allocation exists.

This work also opens to a more general study of consensus-based notions of envy. For instance, instead of focusing on approval envy between agents, one could also be interested in using consensus to determine whether a given agent should be envious in general or not. More generally, one could also look for allocations that are judged envy-free by a given quota of agents. We leave the study of these notions for future work.

REFERENCES

- [1] David J. Abraham, Katarína Cechlárová, David F. Manlove, and Kurt Mehlhorn, ‘Pareto optimality in house allocation problems’, in *Algorithms and Computation: 15th International Symposiums*, eds., Rudolf Fleischer and Gerhard Trippen, 3–15, Springer Berlin Heidelberg, Berlin, Heidelberg, (2005).
- [2] Georgios Amanatidis, Georgios Birmpas, and Vangelis Markakis, ‘Comparing approximate relaxations of envy-freeness’, in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pp. 42–48. International Joint Conferences on Artificial Intelligence Organization, (2018).
- [3] Georgios Amanatidis, Apostolos Ntokos, and Evangelos Markakis. Multiple birds with one stone: Beating $1/2$ for EFX and GMMS via envy cycle elimination. ArXiv-1909.07650, 2019.
- [4] Aurélie Beynier, Sylvain Bouveret, Michel Lemaître, Nicolas Maudet, Simon Rey, and Parham Shams, ‘Efficiency, Sequenceability and Deal-Optimality in Fair Division of Indivisible Goods’, in *18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019)*, Montreal, Canada, (May 2019).
- [5] Arpita Biswas and Siddharth Barman, ‘Fair division under cardinality constraints’, in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pp. 91–97. International Joint Conferences on Artificial Intelligence Organization, (7 2018).
- [6] Steven J. Brams and Peter C. Fishburn, ‘Fair division of indivisible items between two people with identical preferences: Envy-freeness, pareto-optimality, and equity’, *Social Choice and Welfare*, **17**, 247–267, (2002).
- [7] Steven J. Brams and Alan D. Taylor, *Fair Division — From Cake-cutting to Dispute Resolution*, Cambridge University Press, 1996.
- [8] Eric Budish, ‘The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes’, *Journal of Political Economy*, **119**(6), 1061–1103, (dec 2011).
- [9] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang, ‘The unreasonable fairness of maximum nash welfare’, in *Proceedings of the 2016 ACM Conference on Economics and Computation, EC ’16*, pp. 305–322, New York, NY, USA, (2016). ACM.
- [10] Yann Chevaleyre, Ulle Endriss, Sylvia Estvie, and Nicolas Maudet, ‘Reaching Envy-free States in Distributed Negotiation Settings’, in *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI’07*, pp. 1239–1244, San Francisco, CA, USA, (2007). Morgan Kaufmann Publishers Inc.
- [11] 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.
- [12] John P. Dickerson, Jonathan Goldman, Jeremy Karp, Ariel D. Procaccia, and Tuomas Sandholm, ‘The computational rise and fall of fairness’, in *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI-14)*, pp. 1405–1411, Québec City, Québec, Canada, (July 2014). AAAI Press.
- [13] François Durand, Fabien Mathieu, and Ludovic Noirie, ‘SVVAMP: simulator of various voting algorithms in manipulating populations’, in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pp. 4353–4354, (2016).
- [14] Duncan K. Foley, ‘Resource allocation and the public sector’, *Yale Economic Essays*, **7**(1), 45–98, (1967).
- [15] Jonathan Goldman and Ariel D. Procaccia, ‘Spliddit: Unleashing fair division algorithms’, *SIGecom Exch.*, **13**(2), 41–46, (January 2015).
- [16] Michel Gondran and Michel Minoux, *Graphs and algorithms*, Wiley-Interscience series in discrete mathematics, Wiley, 1984.
- [17] Caroline Guibet Lafaye, *Justice sociale et éthique individuelle*, Presses Université Laval, 2006.
- [18] Richard Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi, ‘On approximately fair allocations of divisible goods’, in *Proceedings of the 5th ACM Conference on Electronic Commerce (EC-04)*, pp. 125–131, New York, NY, (May 2004). ACM.
- [19] Colin L. Mallows, ‘Non-Null Ranking Models. I’, *Biometrika*, **44**(1-2), 114–130, (06 1957).
- [20] Hervé Moulin, *Fair Division and Collective Welfare*, MIT Press, 2003.
- [21] Abraham Othman, Tuomas Sandholm, and Eric Budish, ‘Finding approximate competitive equilibria: efficient and fair course allocation’, in *Proceedings of the 9th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-10)*, eds., Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen, pp. 873–880, Toronto, Canada, (May 2010). IFAAMAS.
- [22] Philippe Van Parijs, *Real Freedom for All: What (if Anything) Can Justify Capitalism?*, number 9780198293576 in OUP Catalogue, Oxford University Press, 1997.
- [23] Jonathan Perry and Robert C Powers, ‘Anonymity, monotonicity, and quota pair systems’, *Mathematical Social Sciences*, **60**(1), 57–60, (2010).
- [24] Ariel D. Procaccia and Junxing Wang, ‘Fair enough: Guaranteeing approximate maximin shares’, in *Proceedings of the 15th ACM Conference on Electronic Commerce (EC-14)*, eds., Moshe Babaioff, Vincent Conitzer, and David Easley, pp. 675–692, Palo Alto, CA, (June 2014). ACM.
- [25] Alvin Roth and Marilda Sotomayor, ‘Two-sided matching: A study in game-theoretic modeling and analysis’, *Games and Economic Behavior*, **4**(1), 161–165, (1992).
- [26] Erel Segal-Halevi and Warut Suksompong, ‘Democratic fair allocation of indivisible goods’, *Artificial Intelligence*, **277**, (2019).
- [27] Lloyd Shapley and Herbert Scarf, ‘On cores and indivisibility’, *Journal of mathematical economics*, **1**(1), 23–37, (1974).
- [28] Hugo Steinhaus, ‘The problem of fair division’, *Econometrica*, **16**(1), 101–104, (January 1948).
- [29] Hobart Peyton Young, *Equity in Theory and Practice*, Princeton University Press, 1994.