

Minimizing and balancing envy among agents using Ordered Weighted Average

Parham Shams, Aurélie Beynier, Sylvain Bouveret, Nicolas Maudet

► **To cite this version:**

Parham Shams, Aurélie Beynier, Sylvain Bouveret, Nicolas Maudet. Minimizing and balancing envy among agents using Ordered Weighted Average. 13èmes Journées d'Intelligence Artificielle Fondamentale (JIAF 2019), Jul 2019, Toulouse, France. pp.175-182. hal-02302957

HAL Id: hal-02302957

<https://hal.archives-ouvertes.fr/hal-02302957>

Submitted on 1 Oct 2019

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.

Minimizing and balancing envy among agents using Ordered Weighted Average

Parham Shams¹ Aurélie Beynier¹ Sylvain Bouveret² Nicolas Maudet¹

¹ LIP6, Sorbonne Université, CNRS, LIP6, F-75005 Paris, France

² LIG, Univ. Grenoble Alpes, CNRS, Grenoble INP, Grenoble, France

{aurelie.beynier,nicolas.maudet,shams}@lip6.fr sylvain.bouveret@imag.fr

Abstract

In the problem of fair resource allocation, envy freeness is one of the most interesting fairness criterion as it ensures that no agent prefers the bundle of another agent. However, when considering indivisible goods, an envy-free allocation may not exist. In this paper, we investigate a new relaxation of envy freeness consisting in minimizing the Ordered Weighted Average (OWA) of the envy vector. The idea is to choose the allocation that is fair in the sense of the distribution of the envy among agents. The OWA aggregator is a well-known tool to express fairness in multiagent optimization. In this paper, we focus on fair OWA operators where the weights of the OWA are decreasing. When an envy-free allocation exists, minimizing OWA will return this allocation. However, when no envy-free allocation exists, one may wonder how fair min OWA allocations are.

After some definitions and description of the model, we show how to formulate the computation of such a min OWA allocation as a Mixed Integer Program. Then, we investigate the link between the min OWA allocation and other well-known fairness measures such as max min share and envy freeness up to one good or to any good.

1 Introduction

In this paper, we investigate fair division of indivisible goods. In this context, several approaches have been proposed to model fairness. Amongst these models, one prominent solution concept is to look for *envy-free* allocations [9]. These allocations are such that no agent would like to swap her own bundle with the bundle of any other agent.

Envy-freeness is a very attractive criterion : the fact that each agent is better off with her own share than with any other share is a guarantee of social stability. Besides, since this criterion is only based on personal comparisons, it does not require any interpersonal comparability. Unfortunately, envy-freeness is a very demanding notion, and it is a

well-known fact that in many situations, no such allocation exists (consider for instance the simple situation where the number of items to allocate is strictly less than the number of agents at stake). Hence several relaxations of the envy-freeness notion have been studied in recent years. Two orthogonal approaches have been considered. A first possibility is to “forget” some items when comparing the agents’ shares. This leads to the definition of envy-freeness up to one good [10] and envy-freeness up to any good [6]. Recently, Amanatidis *et al.* [1] explored how different relaxations of envy-freeness relate to each other. Another possible approach is to relax the Boolean notion of envy and to introduce a quantity of envy that we seek to minimize. This is the path followed by Lipton *et al.* [10] or Endriss *et al.* [7] for instance. Several approximation algorithms dedicated to minimize these measures were subsequently designed, see e.g. [11].

In this paper, we elaborate on the idea of minimizing the degree of envy. More precisely, we explore the idea of finding allocations where envy is “fairly balanced” amongst agents. For that purpose, we start from the notion of individual degree of envy and use a *fair* Ordered Weighted Average operator¹ to aggregate these individual envies into a collective one, that we try to minimize. After giving some preliminary definitions in Section 2, we formally introduce our fairness MinOWA envy criterion (Section 3) and we show of OWA minimization problem can be formulated as a linear program. We then relate this criterion to other fairness notions and study properties of the allocations obtained by minimizing the OWA of the envy vector (Section 4). Finally, we present some experimental results investigating the fairness of min OWA solutions (Section 5).

1. By “fair”, we mean an OWA where weights are non-increasing.

2 Model and Definitions

We will consider a classic multiagent resource allocation setting, where a finite set of *objects* $\mathcal{O} = \{o_1, \dots, o_m\}$ has to be allocated to a finite set of *agents* $\mathcal{N} = \{1, \dots, n\}$. In this setting, an *allocation* is a vector $\vec{\pi} = \langle \pi_1, \dots, \pi_n \rangle$ of *bundles* of objects, such that $\forall i, \forall j$ with $i \neq j : \pi_i \cap \pi_j = \emptyset$ (exclusion : a given object cannot be allocated to more than one agent) and $\bigcup_{i \in \mathcal{N}} \pi_i = \mathcal{O}$ (no free-disposal : all the objects are allocated). $\pi_i \subseteq \mathcal{O}$ is called agent i 's *share*.

Any satisfactory allocation must take into account the agents' preferences on the objects. Here, we will make the assumption that these preferences are *numerically additive*. Each agent i has a *utility function* $u_i : 2^{\mathcal{O}} \rightarrow \mathbb{R}^+$ measuring her satisfaction $u_i(\pi)$ when she obtains share π , which is defined as follows :

$$u_i(\pi) \stackrel{\text{def}}{=} \sum_{o_k \in \pi} w(i, o_k),$$

where $w(i, o_k)$ is the weight given by agent i to object o_k . This assumption, as restrictive as it may seem, is made by a lot of authors [10, 2, for instance] and is considered a good compromise between expressivity and conciseness.

Definition 1 *An instance of the additive multiagent resource allocation problem (add-MARA instance for short) $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ is a tuple with \mathcal{N} and \mathcal{O} as defined above and $w : \mathcal{N} \times \mathcal{O} \rightarrow \mathbb{R}^+$ is a mapping with $w(i, o_k)$ being the weight given by agent i to object o_k . We will denote by $\mathcal{P}(I)$ the set of allocations for I .*

2.1 Envy-free allocations

A prominent fairness notion in multiagent resource allocation is *envy-freeness*. Envy-freeness (EF) can be defined as follows :

Definition 2 *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\vec{\pi}$ be an allocation of I . $\vec{\pi}$ is envy-free if and only if $\forall i, j \in \mathcal{N}, u_i(\pi_i) \geq u_i(\pi_j)$.*

In other words, every agent i weakly prefers her own share to the share of any other agent j .

In the context of fair division of indivisible goods, this notion is very demanding and there exists a lot of add-MARA instances for which no envy-free allocation exists. To relax envy-freeness, a possibility is to introduce a notion of degree of envy based on pairwise envy [10].

Definition 3 *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\vec{\pi}$ be an allocation of I . The pairwise envy between agents i and j is defined as follows :*

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

In other words, the pairwise envy between i and j is 0 if i does not envy j , and is equal to the difference between the utility for agent i if she had agent j 's bundle and her actual utility in the allocation $\vec{\pi}$. It can be interpreted as how much agent i envies agent j 's bundle.

From that notion of pairwise envy, we can derive a notion of total envy of an agent, that we define as the maximal pairwise envy that this agent experiences for another agent :

Definition 4 *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\vec{\pi}$ be an allocation of I . The envy of agent i is :*

$$e(i, \vec{\pi}) \stackrel{\text{def}}{=} \max_{j \in \mathcal{N}} pe(i, j, \vec{\pi}).$$

The vector $\vec{e}(\vec{\pi}) = \langle e(1, \vec{\pi}), \dots, e(n, \vec{\pi}) \rangle$ will be called vector of envy of allocation $\vec{\pi}$.

Note that an allocation $\vec{\pi}$ is envy-free if and only if $\vec{e}(\vec{\pi}) = \langle 0, \dots, 0 \rangle$.

2.2 Relaxations of envy-freeness

Different relaxations of the envy-freeness notions have been proposed to measure the fairness of an allocation when there is no envy-free solution. They correspond to weaker solution concepts that are easier to satisfy. *Envy-freeness up to one good* (EF1) [10, 4] is one of the most studied relaxations. An allocation is said to be envy-free up to one good if, for each envious agent i , the envy of i towards an agent j can be eliminated by removing an item from the bundle of j .

Definition 5 *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\vec{\pi}$ be an allocation of I . $\vec{\pi}$ is envy-free up to one good if and only if $\forall i, j \in \mathcal{N}$, either $u_i(\pi_i) \geq u_i(\pi_j)$ or $\exists o_k \in \pi_j$ such that $u_i(\pi_i) \geq u_i(\pi_j) \setminus \{o_k\}$.*

It has been proved that an EF1 allocation always exists and, in the additive case, an EF1 allocation can be obtained using a round-robin protocol.

Caragiannis *et al.* [6] proposed another relaxation of envy-freeness that is closer to the original notion. An allocation is said to be *envy-free up to any good* (EFX) if for all envious agents i , the envy of i towards j can be eliminated by removing *any* item from the bundle of j .

Definition 6 *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\vec{\pi}$ be an allocation of I . $\vec{\pi}$ is envy-free up to one good if and only if $\forall i, j \in \mathcal{N}$, either $u_i(\pi_i) \geq u_i(\pi_j)$ or $\forall o_k \in \pi_j$ $u_i(\pi_i) \geq u_i(\pi_j \setminus \{o_k\})$.*

Clearly, we have $EF \implies EFX \implies EF1$. While an EF1 allocation can be computed in polynomial time, the guarantee of existence of an EFX allocation (where the full set of objects is allocated) remains an open issue [5].

2.3 Other fairness notions

Other notions of fairness have been introduced in the literature. Bouveret *et al.* [3] showed that some connections can be drawn between some widely used notions among which the min-max fair share, the proportional fair share and the max-min fair share.

Definition 7 For a MARA instance \mathcal{I} we define the min-max share (mMS) of agent i as follows :

$$u_i^{mMS} \stackrel{\text{def}}{=} \min_{\vec{\pi} \in \mathcal{P}(\mathcal{I})} \max_{j \in \mathcal{N}} u_j(\pi_j)$$

Besides, we say that an allocation $\vec{\pi}$ is min-max fair share if every agent gets at least her min-max share; formally : $u_i(\pi_i) \geq u_i^{mMS}$

Definition 8 For a MARA instance \mathcal{I} we define the proportional share (PS) of agent i as follows :

$$u_i^{PS} \stackrel{\text{def}}{=} \frac{1}{n} u_i(\mathcal{N})$$

Besides, we say that an allocation $\vec{\pi}$ is proportional fair share if every agent gets at least her proportional share; formally : $u_i(\pi_i) \geq u_i^{PS}$.

Definition 9 For a MARA instance \mathcal{I} we define the max-min share (MMS) of agent i as follows :

$$u_i^{MMS} \stackrel{\text{def}}{=} \max_{\vec{\pi} \in \mathcal{P}(\mathcal{I})} \min_{j \in \mathcal{N}} u_j(\pi_j)$$

Besides, we say that an allocation $\vec{\pi}$ is max-min fair share if every agent gets at least her max-min share; formally : $u_i(\pi_i) \geq u_i^{MMS}$.

Bouveret *et al.* [3] showed that these notions form a linear scale of increasing requirements where :

$$EF \implies mMS \implies PS \implies MMS$$

This scale can be used to characterize the level of fairness of a given allocation. In this hierarchy, EF is the most demanding criterion while the max-min fair share impose few restrictions.

3 MinOWA Envy

Our approach elaborates on minimizing the degree of envy of the agents while balancing the envy among the agents as suggested by [10]. The general idea would be to look for allocations that minimize this vector of envy in some sense : the lower this vector is, the less envious the agents are. This corresponds to a multiobjective optimization problem where each component of the envy vector is a different objective to minimize.

3.1 Fair OWA

There are different ways to tackle this minimization problem, each approach conveying a different definition of minimization. Our approach, guided by the egalitarian notion of fairness [13], is to ensure that the envy is as equally distributed as possible amongst agents. To this end, we will use a prominent aggregation operator that can convey fairness requirements : order weighted averages.

Ordered Weighted Averages (OWA) have been introduced by Yager [14] with the idea to build a family of aggregators that can weight the importance of objectives (or agents) according to their relative utilities, instead of their identities. In this way, we can explicitly choose to favour the poorest (or richest) agents, or to concentrate the importance of the criterion on the middle-class agents. Formally, the OWA operator is defined as follows :

Definition 10 Let $\vec{\alpha} = \langle \alpha_1, \dots, \alpha_n \rangle$ be a vector of weights. In the context of minimization, the ordered weighted average parameterized by $\vec{\alpha}$ is the function :

$$owa^{\vec{\alpha}} : \vec{x} \mapsto \sum_{i=1}^n \alpha_i \times x_i^{\downarrow},$$

where \vec{x}^{\downarrow} denotes a permutation of \vec{x} where $x_1^{\downarrow} \geq x_2^{\downarrow} \geq \dots \geq x_n^{\downarrow}$.

If we want to convey some notion of fairness, we have to give more weights to the unhappiest agents. Intuitively, it means that the weights in $\vec{\alpha}$ should be decreasing. This notion can be formalized by the following property. Let \vec{x} be a vector such that $x_j \geq x_i$ (i is better off than j) and let ε be such that $0 \leq \varepsilon \leq 2(x_j - x_i)$. Then, for any non-increasing vector $\vec{\alpha}$:

$$owa^{\vec{\alpha}}(\vec{x}) \leq owa^{\vec{\alpha}}(\langle x_1, \dots, x_i + \varepsilon, \dots, x_j - \varepsilon, \dots, x_n \rangle)$$

In other words, such an OWA favours any transfer of wealth from an happier agent to an unhappier agent. Such a transfer is called a *Pigou-Dalton* transfer, and the OWA with non-increasing weight vectors $\vec{\alpha}$ are called *fair OWAs*.

Since our motivation is to minimize the envy while equally distributing it between the agents, we propose to minimise the fair OWAs of the envy vector. Our solution concept can then be defined as follows :

Definition 11 Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\vec{\alpha}$ be a non-increasing vector. An allocation $\vec{\pi}$ is an $\vec{\alpha}$ -minOWA Envy allocation if :

$$\vec{\pi} \in \operatorname{argmin}_{\vec{\pi} \in \mathcal{P}(\mathcal{I})} (owa^{\vec{\alpha}}(\vec{e}(\vec{\pi}))).$$

It is important to note that a major advantage of this solution is that it always exists as it is the result of an optimization process. Moreover, as we will see, this optimization problem can be modeled as an Integer Linear Program,

which will enable efficient computation of optimal allocations.

3.2 Linearization of OWA minimization

By using a linearization introduced by Ogryczak in [12] we can model our problem (minimization of the OWA of the envy vector) as a linear program. This linearization smartly uses the definition of OWA with its Lorenz components. Let $e_{(k)}$ denote the k^{th} bigger envy of an agent and $L_{(k)}$ the k^{th} Lorenz component of the Lorenz vector L . Moreover we consider decreasing OWA weights (because we consider here a fair OWA) so $\alpha_1 \geq \alpha_2 \dots \geq \alpha_n$ and we denote $\vec{\alpha} = (\alpha_1 - \alpha_2, \alpha_2 - \alpha_3, \dots, \alpha_n)$. Finally, we recall the definition of the Lorenz vector $L = (e_{(1)}, e_{(1)} + e_{(2)}, \dots, \sum_{i=1}^n e_{(i)})$. The OWA of a vector can be written with its Lorenz components :

$$\begin{aligned} \min \text{OWA}(\vec{e}(\vec{\pi})) &= \min \sum_{k=1}^n \alpha_k e_{(k)} \\ &= \min \sum_{k=1}^n \alpha'_k L_k(\vec{e}) \end{aligned}$$

Besides, there is a known LP to compute $L_k(\vec{e})$:

$$L_k(\vec{e}) = \max \sum_{i=1}^n a_i^k e_i$$

$$s.t. \begin{cases} \sum_{i=1}^m a_i^k = k & *r_k \\ a_i^k \in [0, 1] & \forall i \in \llbracket 1, n \rrbracket \quad *b_i^k \end{cases}$$

However, we cannot inject this LP in the previous one as their optimization directions are not similar. This is why we use the dual (with the dual variables shown in red in the previous LP) of the LP :

$$\begin{aligned} L_k(\vec{e}) &= \min \quad kr_k - \sum_{i=1}^n b_i^k \\ s.t. \quad &\begin{cases} r_k + b_i^k \geq e_i & \forall i \in \llbracket 1, n \rrbracket \\ b_i^k \geq 0 & \forall i \in \llbracket 1, n \rrbracket \end{cases} \end{aligned}$$

The linearization of OWA is over but we still have to write three types of constraints directly related to our problem. First, we want to express the fact that any item j can be held by at most one agent $\sum_{i=1}^n z_i^j = 1$. We also should write the constraint expressing completeness (every item has to be allocated) $\sum_{i=1}^n \sum_{j=1}^m z_i^j = m$ but we note that adding this constraint is redundant as it is the sum over the items of the previous constraint. Finally, we have to make the link between the envy of agent a_i denoted by e_i and the utilities of the agents. e_i is the maximum of the envies between the agents, hence we linearize that maximum easily as we are in a minimization problem. For e_i we write n constraints

expressing that e_i is greater than or equal to how much a_i envies another agent. We can finally write the final Mixed Integer Linear Program :

$$\begin{aligned} \min \text{OWA}(\vec{e}(\vec{\pi})) &= \min \sum_{k=1}^n \alpha'_k (kr_k + \sum_{i=1}^n b_i^k) \\ &\begin{cases} r_k + b_i^k \geq e_i & \forall i, k \in \llbracket 1, n \rrbracket \\ e_i \leq r_k + b_i^k & \forall i, k \in \llbracket 1, n \rrbracket \\ e_i \geq \sum_{j=1}^m w(i, o_j) (z_h^j - z_i^j) & \forall i, h \in \llbracket 1, n \rrbracket \\ \sum_{i=1}^n z_i^j = 1 & \forall j \in \llbracket 1, m \rrbracket \\ z_i^j \in \{0, 1\} & \forall j \in \llbracket 1, m \rrbracket \quad \forall i \in \llbracket 1, n \rrbracket \\ b_i^k \geq 0 & \forall i, k \in \llbracket 1, n \rrbracket \\ e_i \geq 0 & \forall i \in \llbracket 1, n \rrbracket \end{cases} \end{aligned}$$

4 Link with other fairness measures

We focus here on the possible links between the min OWA allocation and other fairness measures such as max-min fair share, envy-freeness up to one good and envy-freeness up to any good. We recall that if an envy-free allocation exists, it will be returned by the min OWA optimization.

4.1 Commensurability issues

We first show that if the utilities of the agents are not commensurable, we can find an example for which the min OWA allocation violates EF1 and MMS. Let's consider the following MARA instance involving 3 agents and 4 objects :

Π_1	o_1	o_2	o_3	o_4
a_1	6	3	2	1
a_2	0.07	0.06	0.04	0.03
a_3	20	0	0	0

Π_2	o_1	o_2	o_3	o_4
a_1	6	3	2	1
a_2	0.07	0.06	0.04	0.03
a_3	20	0	0	0

Π_3	o_1	o_2	o_3	o_4
a_1	6	3	2	1
a_2	0.07	0.06	0.04	0.03
a_3	20	0	0	0

The allocation Π_1 is obviously EF1. However, the allocation Π_2 is the one returned by the OWA minimization and it is not EF1. Indeed, a_2 gets no item in this allocation and a_1 gets three items that a_2 values strictly positively.

Hence, removing one item from a_1 's bundle is not enough for a_2 not to envy a_1 anymore. Besides, we can notice that $u_1^{\text{MMS}} = 3$, $u_2^{\text{MMS}} = 0.06$ and $u_3^{\text{MMS}} = 0$ so the allocation Π_3 is MMS. So, we have found that the min OWA allocation violates EF1 and MMS.

Finally, by going up in the hierarchy introduced by Bouveret et al., we would like to check whether the absence of commensurability prevents the min OWA allocation from being PS or mMS. This time, we cannot conclude with the same instance as the previous example has no PS allocation.

We thus consider the following instance :

Π_1	o_1	o_2	o_3	o_4
a_1	10	6	6	1
a_1	0.1	0.06	0.06	0.01
a_1	1	6	6	10

Π_2	o_1	o_2	o_3	o_4
a_1	10	6	6	1
a_1	0.1	0.06	0.06	0.01
a_1	1	6	6	10

The allocation Π_1 is mMS as $u_1^{\text{mMS}} = u_3^{\text{mMS}} = 100 * u_2^{\text{mMS}} = 10$. It is then also PS. However, the allocation Π_2 is the one returned by the minimization of the envy vector and it is obviously not PS nor mMS. This is why from now on we will only consider MARA instances with commensurable agents (with K constant) :

$$\sum_{j=1}^m w(i, o_j) = K \quad \forall i \in \llbracket 1, n \rrbracket$$

4.2 Two-agent settings

In the special case where the allocation problem involves only two agents, we highlight strong connections between the min OWA allocation and other fairness measures.

Proposition 1 *When the MARA instance involves only 2 agents then, the min OWA allocation is a max-min fair share allocation (thus also EFX).*

Proof 1 *For MARA instances where an envy-free allocation exists, our proof is straightforward as min OWA returns the EF allocation. It is thus also MMS, EF1 and EFX.*

We now focus on MARA instances for which there is no EF allocation. In the presence of only 2 agents any min OWA allocation $\vec{\pi}$ is such that only one of the two agents is envious. Indeed, if no agent is envious then it means the MARA instance has an envy-free allocation (which is a contradiction). Similarly, if both agent are envious it means there is an envy-free allocation (which is again a contradiction) as the agents would just have to exchange their

bundles to obtain that allocation. Consequently, the sorted envy vector will be of this form $(e, 0)$. So, let us consider that such an allocation is not MMS. The agent that is envy-free obviously has her max-min share. So, under the assumption that the allocation is not MMS, the envious agent does not have her max-min share. But, in this case, the envious agent could obtain a better share and hence would have a lower envy. This leads to a contradiction because we consider the min OWA allocation. Indeed, in this case just, the min OWA allocation minimizes the envy of the envious agent as the other was not envious. It is known [6] in the pairwise setting that MMS implies EFX. \square

4.3 EF1 for the general case : $n \geq 3$

We now turn to more general settings involving at least 3 agents. Since an EF1 allocation is guaranteed to exist, we more specifically focus on the possible links between min OWA and EF1. Unfortunately, we mainly obtain negative results.

Proposition 2 *In the general case, the min OWA allocation is not necessarily EF1.*

Proof 2 *We will prove this proposition through two examples with different OWA weights :*

Example 1 *We consider a MARA instance involving 3 agents and 7 objects, with the OWA weights $\alpha = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Minimizing OWA of the envy vector using this set of weight consists in minimizing the sum of the envies of the agents.*

Π_1	o_1	o_2	o_3	o_4	o_5	o_6	o_7
a_1	5	1	1	1	1	1	1
a_2	5	1	1	1	1	1	1
a_3	5	1	1	1	1	1	1

Π_2	o_1	o_2	o_3	o_4	o_5	o_6	o_7
a_1	5	1	1	1	1	1	1
a_2	5	1	1	1	1	1	1
a_3	5	1	1	1	1	1	1

The allocation Π_1 is the one returned by the minimization of the OWA of the envy vector. a_1 is envy free whereas a_2 and a_3 both envy a_1 by respectively 3 and 1. Hence, the sorted envy vector is $(3, 1, 0)$ and the value of OWA is $\frac{4}{3}$. Moreover, this allocation is not EF1 as a_2 would still envy a_3 by 1 even if one item is removed from the latter.

The allocation Π_2 has the same OWA value ($\frac{4}{3}$) but it is EF1. In fact, in this allocation, a_2 does not envy a_3 . We could then argue that one of the min OWA allocations was EF1. This is why we show a stronger example for which the unique min OWA allocation is not EF1.

Example 2 *We consider a MARA instance involving 3 agents and 4 objects, with the OWA weights $\alpha = (1, 0, 0)$.*

Minimizing OWA of the envy vector using this set of weight consists in minimizing the maximum envy of the agents.

Π_1	o_1	o_2	o_3	o_4
a_1	14	3	2	1
a_2	7	6	4	3
a_3	20	0	0	0

Π_2	o_1	o_2	o_3	o_4
a_1	14	3	2	1
a_2	7	6	4	3
a_3	20	0	0	0

The allocation Π_1 corresponds to the min OWA one (this is the only one min OWA allocation). The OWA value of the envy vector is 9 (the envy of a_1 towards a_3). However, the allocation is not EF1. In fact, regardless of the item we remove from a_1 , a_2 will still envy her.

The allocation Π_2 is EF1 but its OWA value is 11 (it still stands for the envy of a_1 towards a_3) which is quite far from the min OWA solution. Moreover, contrary to the previous example, a small change in the OWA weight vector does not change the situation. For instance, even for $\alpha = (0.9, 0.1, 0)$ the ordering stays the same.

From the previous examples, we can conclude that the minimization of the OWA of the envy vector does not necessarily return an EF1 allocation in the general case. \square

However, we can wonder if we could always find a set of weights that returns an EF1 by minimizing the OWA. We leave this question open for the moment.

4.4 Continuity of EF1

We showed that the weights of the OWA influences the EF1 property of the min OWA solution. We can thus wonder whether there exists, among all the allocations, a kind of continuity of the EF1 solutions. It would be the case if by sorting all the allocations by their OWA value, all the EF1 allocations would be contained in a single interval of OWA value. We show that it is not the case through the following example.

Example 3 We consider a MARA instance involving 3 agents and 3 objects, with the OWA weights $\alpha = (0.44, 0.36, 0.2)$:

Π_1	o_1	o_2	o_3	Π_2	o_1	o_2	o_3
a_1	2	1	3	a_1	2	1	3
a_2	2	3	1	a_2	2	3	1
a_3	1	0	5	a_3	1	0	5

Π_3	o_1	o_2	o_3	Π_4	o_1	o_2	o_3
a_1	2	1	3	a_1	2	1	3
a_2	2	3	1	a_2	2	3	1
a_3	1	0	5	a_3	1	0	5

Π_5	o_1	o_2	o_3
a_1	2	1	3
a_2	2	3	1
a_3	1	0	5

For the purpose of the proof, we extract 5 allocations among all possible allocations. Allocations Π_1 to Π_5 are sorted by increasing values of the OWA of their envy vector. Π_1 is the best solution and Π_5 is the worst among the five selected allocations. The following table summarizes the OWA values of these allocations and specifies whether the allocations are EF1 or not :

	Π_1	Π_2	Π_3	Π_4	Π_5
OWA	0.44	1.32	1.76	2.20	2.20
EF1	True	False	True	True	False

We can easily see that there is no continuity of EF1 regarding the OWA values as Π_2 (which is between Π_1 and Π_3) is not EF1 whereas Π_1 and Π_3 are both EF1. Interestingly, it can be noticed that Π_4 and Π_5 have the same OWA values even if the allocations totally differ (not a single object has been given to the same agent). Moreover, Π_4 is EF1 whereas Π_5 is not.

5 Experimental results

We drew some experiments to compare the performances of the allocations obtained by min OWA envy with other allocation methods. The first method is a succession of sincere choices where the picking sequence is a concatenation of round robins. From all the possible allocations (for every possible round robin) generated by this method we keep the minRR (maxRR) that (respectly) is the allocation with the minimum OWA envy (maximum) while meanRR is the average of all these allocations. The max OWA util is the allocation that maximizes the OWA of the valuations of the agents. For the OWA allocation method, we used the following fair weight $(\frac{1}{2}, \dots, \frac{1}{2^n})$.

We generated 100 random add-MARA instances for each method and for each couple $(|N|, |O|)$ from $(3, 4)$ to $(10, 12)$. We considered such couples of values in order to produce settings where few EF allocations exist as recommended in [8]. We did not go beyond 10 agents because of

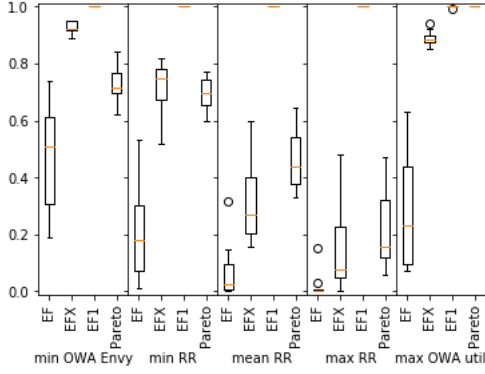


FIGURE 1 – Percentage of EF, EFX, EF1, PO for each method averaged on all $(|N|, |O|)$




FIGURE 2 – Value of the OWA of the vector of envy for each allocation protocol

the computation time of the RR methods as we had to generate every possible allocation from every possible round robin. Indeed, even for 10 agents both the OWA methods were computed in less than two seconds (via Gurobi solver).

We evaluated the performances of the different methods by measuring the percentage allocations that are EF, EFX, EF1 and Pareto optimal. The results are aggregated in Figure 1. It can first be observed that our method obtains the highest rate of EF and EFX allocations : it is not that surprising for EF as our method is a relaxation of EF but the result about EFX is very promising. Every method has returned an EF1. Regarding the number of Pareto optimal allocations, our methods places second as maximizing OWA of the utilities necessarily produces a PO allocation.

Figure 2 shows how far the other methods are in terms of OWA envy values. We notice that even the closest method (max OWA util) is relatively far from our method, none of them has similar performances. Figure 3 shows that our




FIGURE 3 – Value of the sum of the vector of utilities for each allocation protocol

method has also good performances in terms of utilitarian criteria.

We performed some experiments for the OWA weights $(1, 0, 0, 0)$ and obtained similar results. For these weights the max OWA util does not always return a PO allocation and our method returns a lower rate of EFX (but still more than the other methods). This reinforces the idea that considering the envy of all agents and not only the maximum leads to fairer allocations.

6 Conclusion

In this paper, we introduced a new fairness concept following the idea of minimizing envy. More particularly, we used an OWA to express fairness in the distribution of envy between agents. After implementing a MIP to compute min OWA allocations, we showed the connections between the min OWA allocation and other famous fairness measures. We ran some experiments to test the performances of our method and compared it with other allocation protocols.

We also left the question of the performance of our method by considering the weights of the OWA as variables and not as a fix entry of the problem. We could focus on this matter for future works and try for example to learn them.

Finally, it could also be interesting to change the definition of the envy of an agent. Indeed, instead of defining it as the maximum of the pairwise envies we could define it as other aggregations of the pairwise envies such as their sum or even their OWA.

Références

[1] Amanatidis, Georgios, Georgios Birmpas et Vange-

- lis Markakis: *Comparing approximate relaxations of envy-freeness*. Dans *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden.*, pages 42–48, 2018.
- [2] Bansal, Nikhil et Maxim Sviridenko: *The santa claus problem*. Dans *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing, STOC '06*, pages 31–40, New York, NY, USA, 2006. ACM.
- [3] Bouveret, Sylvain et Michel Lemaître: *Characterizing conflicts in fair division of indivisible goods using a scale of criteria*. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290, 2016, ISSN 1573-7454.
- [4] Budish, Eric: *The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes*. *Journal of Political Economy*, 119(6):1061–1103, dec 2011.
- [5] Caragiannis, I., N. Gravin et Xin Huang: *Envy-freeness up to any item with high nash welfare: The virtue of donating items*. note, 2019.
- [6] Caragiannis, Ioannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah et Junxing Wang: *The unreasonable fairness of maximum nash welfare*. Dans *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16*, pages 305–322, New York, NY, USA, 2016. ACM, ISBN 978-1-4503-3936-0.
- [7] Chevaleyre, Yann, Ulle Endriss et Nicolas Maudet: *Allocating goods on a graph to eliminate envy*. Dans *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI-07)*, Vancouver, British Columbia, Canada, juillet 2007.
- [8] Dickerson, John P., Jonathan Goldman, Jeremy Karp, Ariel D. Procaccia et Tuomas Sandholm: *The computational rise and fall of fairness*. Dans *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI-14)*, pages 1405–1411, Québec City, Québec, Canada, juillet 2014. AAAI Press.
- [9] Foley, Duncan K.: *Resource allocation and the public sector*. *Yale Economic Essays*, 7(1):45–98, 1967.
- [10] Lipton, Richard, Evangelos Markakis, Elchanan Mossel et Amin Saberi: *On approximately fair allocations of divisible goods*. Dans *Proceedings of the 5th ACM Conference on Electronic Commerce (EC-04)*, pages 125–131, New York, NY, mai 2004. ACM.
- [11] Nguyen, Trung Thanh et Jörg Rothe: *How to decrease the degree of envy in allocations of indivisible goods*. Dans Perny, Patrice, Marc Pirlot et Alexis Tsoukiàs (éditeurs): *Algorithmic Decision Theory*, pages 271–284, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg, ISBN 978-3-642-41575-3.
- [12] Ogryczak, Włodzimierz et Tomasz Śliwiński: *On solving linear programs with the ordered weighted averaging objective*. *European Journal of Operational Research*, 148:80–91, 2003.
- [13] Rawls, John: *A Theory of Justice*. Harvard University Press, Cambridge, Mass., 1971. Traduction française disponible aux éditions du Seuil.
- [14] Yager, Ronald R.: *On ordered weighted averaging aggregation operators in multicriteria decision making*. *IEEE Transactions on Systems, Man, and Cybernetics*, 18:183–190, 1988.