Minimizing and balancing envy among agents using Ordered Weighted Average

Parham Shams[1], Aurélie Beynier[1], Sylvain Bouveret[2], and Nicolas Maudet[1]

[1] Sorbonne Université, CNRS, LIP6, F-75005 Paris, France
{parham.shams,aurelie.beynier,nicolas.maudet}@lip6.fr
[2] Univ. Grenoble Alpes, CNRS, LIG, Grenoble, France
sylvain.bouveret@imag.fr

Abstract. In fair resource allocation, envy freeness (EF) is one of the most interesting fairness criteria as it ensures no agent prefers the bundle of another agent. However, when considering indivisible goods, an EF allocation may not exist. In this paper, we investigate a new relaxation of EF 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 EF allocation exists, minimizing OWA envy will return this allocation. However, when no EF allocation exists, one may wonder how fair min OWA envy allocations are.
After defining the model, we show how to formulate the computation of such a min OWA envy 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 EF up to one good or to any good. Finally, we run some experiments comparing the performances of our approach with MNW (Max Nash Welfare) on several criteria such as the percentage of EF up to one good and any good.

Keywords: Social Choice, Multiagent Resource Allocation, Fair Allocation, Fair Division of Indivisible Goods

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 [12]. In such allocations, no agent would swap her bundle with the bundle of any other agent.

Envy-freeness is an 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, it does not rely on any interpersonal comparability. Unfortunately, envy-freeness is also a demanding notion as soon as we require all goods to be allocated, and it is well-known that in many situations, no such allocation exists (consider for instance the situation where the number of items to allocate is strictly less than the number of agents at stake). Hence several relaxations of envy-freeness 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 [17] and envy-freeness up to any good [7]. Recently, Amanatidis *et al.* [3] 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.* [17] or Endriss *et al.* [9] for instance. Several approximation algorithms dedicated to minimize these measures were subsequently designed – see e.g. Nguyen *et al.* [18]. Of course such approaches always rely on a specific choice to measure the degree of envy, in particular regarding the aggregation of agents' envies, which can be disputed: is it more appropriate to minimize the maximum envy experienced by some agent in the society, or to minimize the sum of agents' envies?

In this paper, we elaborate on this idea of minimizing the degree of envy but seek to offer a broader perspective. More precisely, we explore the possibility 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 operator (by "fair", we mean an OWA where weights are non-increasing.) to aggregate these individual envies into a collective one, that we try to minimize. This family of operators contains both the egalitarian and utilitarian operators mentioned previously. But doing so also sometimes allows us to draw results valid for the whole family of fair operators. Along our way, we shall for instance see that no algorithm fairly minimizing envy can be guaranteed to return an envy-free allocation up to any good, even though such allocation does exist. More generally, we provide several insights regarding the behaviour of such fair minimizing operators, comparing their outcomes with alternative approaches, either analytically or experimentally. Technically, this is made possible through to the use of linearization techniques which alleviate the burden of computing these outcomes.

The remainder of this paper is as follows. After giving some preliminary definitions in Section 2, we formally introduce our fairness minOWA envy criterion (Section 3) and we show that OWA minimization problems can be formulated as linear programs. We then investigate the link between minimizing the OWA of the envy vector and other fairness notions (Section 4). We thus study fairness guarantees of the minOWA solutions. Finally, we present some experimental results investigating the fairness of min OWA solutions (Section 5).

2 Model and Definitions

We will consider a classic multiagent resource allocation setting, where a finite set of *objects* $\mathcal{O} = \{o_1, \ldots, o_m\}$ has to be allocated to a finite set of *agents* $\mathcal{N} = \{a_1, \ldots, a_n\}$. In this setting, an *allocation* is a vector $\boldsymbol{\pi} = \langle \pi_1, \ldots, \pi_n \rangle$ of *bundles* of objects, such that $\forall a_i, a_j \in \mathcal{N}$ with $i \neq j : \pi_i \cap \pi_j = \emptyset$ (preemption:

a given object cannot be allocated to more than one agent) and $\bigcup_{a_i \in \mathcal{N}} \pi_i = \mathcal{O}$ (no free-disposal: all the objects are allocated). $\pi_i \subseteq \mathcal{O}$ is called agent a_i's *share*. The set of all the possible allocations will be denoted $\mathcal{P}(I)$.

A crucial aspect of fair division problem is how the agents express their preferences over bundles. Here, we assume that these preferences are *numerically additive*: each agent a_i has a *utility function* $u_i : 2^{\mathcal{O}} \to \mathbb{R}$ measuring her satisfaction $u_i(\pi_i)$ when she obtains share π_i, which is defined as $u_i(\pi_i) \overset{def}{=} \sum_{o_k \in \pi_i} w(a_i, o_k)$, where $w(a_i, o_k)$ is the weight given by agent a_i to object o_k. This assumption, as restrictive as it may seem, is made by a lot of authors [17, 4, 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 $w : \mathcal{N} \times \mathcal{O} \to \mathbb{R}$ is a mapping with $w(a_i, o_k)$ being the weight given by a_i to object o_k. We will denote by $\mathcal{P}(I)$ the set of allocations for I.*

In the following, we denote by $\mathcal{I}$ the set of all add-MARA instances. Furthermore, different domain restrictions will be of interest: we denote by $\mathcal{I}^p$ the set of add-MARA instances involving only two agents (pairwise instances), and by $\mathcal{I}^b$ the set of add-MARA instances where agents have binary utilities.

Unless stated otherwise, we will only consider MARA instances with *commensurable* preferences, such that: $\exists K \in \mathbb{N}$ s.t $\forall i \in [\![1, n]\!], \sum_{j=1}^{m} w(a_i, o_j) = K$.

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 $\boldsymbol{\pi}$ be an allocation of I. $\boldsymbol{\pi}$ is* envy-free *if and only if $\forall a_i, a_j \in \mathcal{N}$, $u_i(\pi_i) \geq u_i(\pi_j)$.*

In other words, every agent a_i weakly prefers her own share to the share of any other agent a_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 [17].

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

$$pe(a_i, a_j, \boldsymbol{\pi}) \overset{def}{=} \max\{0, u_i(\pi_j) - u_i(\pi_i)\}.$$

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

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

Definition 4. *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\boldsymbol{\pi}$ be an allocation of I. a_i's* envy*: $e(a_i, \boldsymbol{\pi}) \stackrel{def}{=} \max_{a_j \in \mathcal{N}} pe(a_i, a_j, \boldsymbol{\pi})$. The vector $\boldsymbol{e}(\boldsymbol{\pi}) = \langle e(a_1, \boldsymbol{\pi}), ..., e(a_n, \boldsymbol{\pi}) \rangle$ will be called* envy vector *of allocation $\boldsymbol{\pi}$.*

Here, the max operator is rather a standard choice in the context where one seeks for allocations with bounded envy [17]. Note that an allocation $\boldsymbol{\pi}$ is envy-free if and only if $\boldsymbol{e}(\boldsymbol{\pi}) = \langle 0, ..., 0 \rangle$.

2.2 Weaker notions of envy-freeness

Besides minimizing a degree of envy, different relaxations of the envy-freness notions have also been proposed to cope with situations where there is no envy-free solution. *Envy-freness up to one good* (EF1) [17, 6] is one of the most studied relaxations. An allocation is said to be envy-free up to one good if, for each envious agent a_i, the envy of a_i towards an agent a_j can be eliminated by removing an item from the bundle of a_j.

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

It has been proved that an EF1 allocation always exists and, in the additive case, can be obtained using a round-robin protocol [7].

Caragiannis *et al.* [7] proposed another relaxation of the notion of envy-freeness which is stronger than EF1. An allocation is said to be *envy-free up to any good* (EFX) if for all envious agents a_i, the envy of a_i towards a_j can be eliminated by removing *any* item from a_j's bundle.

Definition 6. *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\boldsymbol{\pi}$ be an allocation of I. $\boldsymbol{\pi}$ is* envy-free up to any (strictly positively valuated) good *if and only if $\forall a_i, a_j \in \mathcal{N}$, either $u_i(\pi_i) \geq u_i(\pi_j)$ or $\forall o_k \in \pi_j$ for which $w(a_i, o_k) > 0$, $u_i(\pi_i) \geq u_i(\pi_j \backslash \{o_k\})$.*

An even more demanding notion called EFX_0 [21, 15] differs on the fact that an agent can forget any object even the ones valued to 0:

Definition 7. *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\boldsymbol{\pi}$ be an allocation of I. $\boldsymbol{\pi}$ is* envy-free up to any good *if and only if $\forall a_i, a_j \in \mathcal{N}$, either $u_i(\pi_i) \geq u_i(\pi_j)$ or $\forall o_k \in \pi_j$, $u_i(\pi_i) \geq u_i(\pi_j \backslash \{o_k\})$.*

Clearly, we have $EF \implies EFX_0 \implies EFX \implies EF1$. While an EF1 allocation can be computed in polynomial time, the guarantee of existence of an EFX allocation remains an open issue in the general settings [7]. The existence guarantee of an EFX solutions has been proved for few agents (at most 3 agents) and specific utility functions. For instance it has been proved that an EFX_0 allocation always exists for instances with identical valuations and for instances involving two agents with general and possibly distinct valuations [21], as well as for three agents with additive valuations [8]. When the objects have only two possible valuations, Amanatidis *et al.* [2] proved that any allocation maximizing the Nash Social Welfare is EFX_0. This result provides a polynomial algorithm for computing EFX_0 allocations in the two-agent setting.

Other notions of fairness have been introduced in the literature. Bouveret *et al.* [5] for instance exhibited some connections between widely used notions among which the max-min share (MMS, also known as I cut you chose). An allocation is MMS if every agent gets at least her max-min share. As shown by Bouveret et al. [5], MMS is less demanding than EF and every EF allocation also satisfies MMS.

Example 1. Let us consider the add-MARA instance with 3 agents and 4 objects:

	o_1	o_2	o_3	o_4
a_1	2^*	$\boxed{6}$	1	1
a_2	2	5^*	$\boxed{2}$	1
a_3	$\boxed{1}$	5	2^*	$\boxed{2^*}$

Note that there is no EF allocation in this instance. The squared allocation $\boldsymbol{\pi}$ and the starred allocation $\boldsymbol{\pi'}$ are both EF1 and EFX. Both allocations satisfy MMS. Allocation $\boldsymbol{\pi}$ leads to the envy vector $\boldsymbol{e}(\boldsymbol{\pi}) = \langle 0, 3, 2\rangle$ while allocation $\boldsymbol{\pi'}$ leads to the envy vector $\boldsymbol{e}(\boldsymbol{\pi'}) = \langle 4, 0, 1\rangle$. Both allocations have the same global envy when considering the sum of the individual envies. However, the envy in $\boldsymbol{\pi'}$ is mainly supported by a_1. To promote fairness, it is natural to prefer the allocations where the envy is balanced among the agents. In this example, $\boldsymbol{\pi}$ should be considered as more fair than $\boldsymbol{\pi'}$.

Recently, the Nash social welfare (which maximizes the product of utilities) was celebrated as a particularly good trade-off between efficiency and fairness [7] because it guarantees to return an EF1 and Pareto-optimal allocation, among others. Finally, some authors have proposed to explore inequality indices in multiagent fair division settings [1, 23, 11]. However, this differs from our proposal since in these approaches inequality is (more classically) evaluated at the level of utilities, while we apply it to envies, as we detail in the next section.

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 Lipton *et al.* [17]. 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 [22], is to ensure that, while being as low as possible, the envy is also distributed as equally as possible amongst agents. To this end, we use a prominent aggregation operator that can convey fairness requirements: order weighted averages.

Ordered Weighted Averages (OWA) have been introduced by Yager [25] 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 8. *Let* $\boldsymbol{\alpha} = \langle \alpha_1, \ldots, \alpha_n \rangle$ *be a vector of weights. In the context of minimization, the* ordered weighted average *parameterized by* $\boldsymbol{\alpha}$ *is the function* $owa^{\boldsymbol{\alpha}} : \boldsymbol{x} \mapsto \sum_{i=1}^{n} \alpha_i \times x_i^{\downarrow}$*, where* $\boldsymbol{x}^{\downarrow}$ *denotes a permutation of* $\boldsymbol{x}$ *such that* $x_1^{\downarrow} \geq x_2^{\downarrow} \geq \ldots \geq x_n^{\downarrow}$.

Amongst all OWA, only those giving more weight to the unhappiest agents can be considered fair in the egalitarian sense. This property can be formalized as follows. Let $\boldsymbol{x}$ be a vector such that $x_j \geq x_i$ (a_i is better off than a_j) and let ε be such that $0 \leq \varepsilon \leq 2(x_j - x_i)$. Then, for any non-increasing vector $\boldsymbol{\alpha}$: $owa^{\boldsymbol{\alpha}}(\boldsymbol{x}) \leq owa^{\boldsymbol{\alpha}}(\langle x_1, \ldots, x_i + \varepsilon, \ldots, x_j - \varepsilon, \ldots, x_n \rangle)$.

In other words, such an OWA favours any transfer of wealth from a happier agent to an unhappier agent. Such a transfer is called a *Pigou-Dalton* transfer, and the OWA with non-increasing weight vectors $\boldsymbol{\alpha}$ are called *fair OWA*. Moreover, we have considered wlog in this paper that the weight vector sums to 1 so we will make no difference between weights $\langle 1, 1, 1 \rangle$ and $\langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \rangle$. Note that fair OWA is also referred to as *Generalized Gini Index* [24] in the literature. In matching problems [16] and multiagent allocation problems [14], fair OWA has been applied to the utility vector so as to maximize a global utility function while reducing inequalities. However, we can note that maximizing the OWA of the utility vector does not necessarily return an EF allocation even when such an allocation exists:

Example 2. Consider this add-MARA instance with 3 agents and 4 objects:

	o_1	o_2	o_3	o_4
a_1	1	2^*	3	$\boxed{4^*}$
a_2	2	$\boxed{2}$	$\boxed{5^*}$	1
a_3	$\boxed{4^*}$	0	4	2

The squared allocation is the allocation that maximizes the value of the OWA of the utility vector with weight $\langle 1, 0, 0 \rangle$. We can easily notice that this allocation is not envy free as a_1 envies a_2. Moreover, the star allocation is obviously an EF one. Note that in the context of maximization, a fair OWA is also defined with non-increasing weights but by sorting the components by decreasing value.

Since our motivation is to return an EF allocation when there is one and otherwise minimize the envy while equally distributing it between the agents, we propose to minimize the fair OWAs of the envy vector.

Definition 9. *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\boldsymbol{\alpha}$ be a non-increasing vector. An allocation $\hat{\boldsymbol{\pi}}$ is an $\boldsymbol{\alpha}$-*minOWA Envy allocation *if:*

$$\hat{\boldsymbol{\pi}} \in \arg \min_{\boldsymbol{\pi} \in \mathcal{P}(I)} (owa^{\boldsymbol{\alpha}}(\boldsymbol{e}(\boldsymbol{\pi}))).$$

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, this optimization problem can be modeled as an Integer Linear Program, which will give a way to compute optimal allocations. Keep also in mind that there can be several allocations with the same OWA envy value.

Let us now see some helpful properties of fair OWA. Note that we will consider here that we are in a minimization context.

Definition 10. *By denoting $v_k^{\downarrow}$ the k^{th} biggest component of a given vector $\boldsymbol{v}$, the Lorenz vector $\boldsymbol{L}$ of $\boldsymbol{v}$ is defined as $\boldsymbol{L}(\boldsymbol{v}) = \langle v_1^{\downarrow}, v_1^{\downarrow} + v_2^{\downarrow}, ..., \sum_{i=1}^{n} v_i^{\downarrow} \rangle$.*

Definition 11. *Let $\boldsymbol{x}$ and $\boldsymbol{y}$ be two vectors of the same size and x_i (respectively y_i) be the i^{th} component of $\boldsymbol{x}$ (respectively $\boldsymbol{y}$). We say that $\boldsymbol{x}$ Pareto dominates $\boldsymbol{y}$ iff for every component $x_i \leq y_i$ and there is one component x_j for which $x_j < y_j$ and $\boldsymbol{x}$ strongly Pareto dominates $\boldsymbol{y}$ iff for every component $x_i < y_i$.*

Definition 12. *We say $\boldsymbol{x}$ (strongly) Lorenz dominates $\boldsymbol{y}$ iff $\boldsymbol{L}(\boldsymbol{x})$ (strongly) Pareto dominates $\boldsymbol{L}(\boldsymbol{y})$.*

Theorem 1. *Perny and Spanjaard 2003*
If $\boldsymbol{x}$ Lorenz dominates $\boldsymbol{y}$ then for any non-increasing weight $\boldsymbol{\alpha}$: $owa^{\boldsymbol{\alpha}}(\boldsymbol{x}) \leq owa^{\boldsymbol{\alpha}}(\boldsymbol{y})$. Similarly if $\boldsymbol{x}$ strongly Lorenz dominates $\boldsymbol{y}$ then for any non-increasing weight $\boldsymbol{\alpha}$: $owa^{\boldsymbol{\alpha}}(\boldsymbol{x}) < owa^{\boldsymbol{\alpha}}(\boldsymbol{y})$.

This helpful property is shown in [20]. As (strong) Pareto dominance implies (strong) Lorenz dominance, the same theorem holds with (strong) Pareto dominance.

By using a linearization introduced by Ogryczak [19] we can model our problem of minimizing the OWA of the envy vector as a linear program. Moreover we consider decreasing OWA weights (fair OWA) so $\alpha_1 \geq \alpha_2 .... \geq \alpha_n$ and we denote by $\boldsymbol{\alpha'} = \langle \alpha_1 - \alpha_2, \alpha_2 - \alpha_3, ..., \alpha_n \rangle$. We introduce a set of $n \times m$ Boolean variables z_i^j: z_i^j is 1 iff o_j is allocated to a_i while r_k and b_i^k are the dual variables (of the LP computing the Lorenz components) and e_i the envy of a_i.

$$\min owa(\boldsymbol{e}(\boldsymbol{\pi})) = \min \sum_{k=1}^{n} \alpha'_k (k r_k + \sum_{i=1}^{n} b_i^k)$$

$$\begin{cases} r_k + b_i^k \geq e_i & \forall i,k \in [\![1,n]\!] \\ e_i \geq \sum_{j=1}^m w(a_i,o_j)(z_h^j - z_i^j) & \forall i,h \in [\![1,n]\!] \\ \sum_{i=1}^n z_i^j = 1 & \forall j \in [\![1,m]\!] \\ z_i^j \in \{0,1\} \quad \forall j \in [\![1,m]\!] & \forall i \in [\![1,n]\!] \\ b_i^k \geq 0, \quad e_i \geq 0 & \forall i,k \in [\![1,n]\!] \end{cases}$$

4 Link with other fairness measures

We focus here on the possible links between the min OWA allocation and other fairness measures. We recall that if an envy-free allocation exists, it will be returned by the min OWA optimization. For any instance I, we denote by PROP(I) the set of allocations satisfying PROP $\in \{EF1, EFX, EFX_0, MMS\}$. We also denote by $\boldsymbol{\alpha}$-min OWA(I) the set of all min OWA optimal allocation for the specific weight vector $\boldsymbol{\alpha}$, and by $\forall$-min OWA(I) the set of $\boldsymbol{\alpha}$-min OWA, for *all* (fair) weight vectors $\boldsymbol{\alpha}$.

4.1 Warm-up: $n = 2$

In the special case where the allocation problem involves only two agents, we highlight strong connections between min OWA allocations and other fairness measures (MMS, EF1 and EFX).

Proposition 1. $\forall I \in \mathcal{I}^p : \forall$-*min OWA*$(I) \subseteq MMS(I) \subseteq EFX(I)$

Proof. For add-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 add-MARA instances for which there is no EF allocation. In the presence of only 2 agents any min OWA allocation $\boldsymbol{\pi}$ is such that only one of the two agents is envious. Indeed, if no agent is envious then it means the add-MARA instance has an envy-free allocation (which is a contradiction). Similarly, if both agents 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 the form $(e, 0)$. Suppose for the sake of contradiction that such an allocation is not MMS. The agent that is envy-free (let us say w.l.o.g it is a_2) obviously has her max-min share. So, under the assumption that the allocation is not MMS, a_1 does not have her max-min share. It means that there is an allocation $\boldsymbol{\pi'}$ such that $\min(u_1(\pi'_1), u_1(\pi'_2)) > u_1(\pi_1)$ and a_2 is still not envious (if a_2 is envious in π', just swap her share with a_1's). Obviously, a_1's pairwise envy for a_2 has decreased in π' compared to that of π, and a_2's envy is still 0. This contradicts the fact that π is the optimal min-OWA envy allocation. Finally, it is known [7] in the two-agents setting that MMS implies EFX, which completes the proof.

However, even though an MMS allocation is EFX, this does not hold for EFX_0 even for 2 agents as we can see in Example 3.

Example 3. Consider this add-MARA instance with 2 agents and 3 objects:

	o_1	o_2	o_3
a_1	$\boxed{1}$	0	2
a_2	0	$\boxed{1}$	$\boxed{2}$

It is easy to see that the squared allocation is MMS as the max-min share of each agent is 1. Moreover, we can see that this allocation is EFX (a_1 can forget o_3) whereas it is not EFX$_0$ (because a_1 has to forget o_2 which does not make here becoming envy-free).

However, we show that we can very easily build an EFX$_0$ allocation from an arbitrary min OWA envy one.

Proposition 2. *For any instance $I \in \mathcal{I}^p$ and for any weight vector $\vec{\alpha}$: $\boldsymbol{\alpha}$-minOWA(I) $\cap$EFX$_0(I) \neq \emptyset$. Furthermore, it can be obtained from an arbitrary $\boldsymbol{\alpha}$-min-OWA envy optimal allocation in linear time.*

Proof. Let us call $\boldsymbol{\pi}$ an arbitrary min OWA allocation. If $\boldsymbol{\pi}$ is envy-free then it is obviously EFX$_0$ and the proof concludes. Note that envy-freeness is checked in $O(1)$ as we just have to check the values of both variables e_1 and e_2. Otherwise, it means that one and exactly one agent is envious, by using a same argument as in the proof of Proposition 1. W.l.o.g. we consider a_1 is the envious agent. We start from $\boldsymbol{\pi}$ and transfer to a_1 all the objects that she values with utility zero. The resulting allocation is called $\boldsymbol{\pi}'$. We show that $\boldsymbol{\pi}'$ is EFX$_0$. a_1 still envies a_2 in $\boldsymbol{\pi}'$ but is EFX by Proposition 1. By transferring all zero-valued objects to her share, she becomes EFX$_0$ in $\boldsymbol{\pi}'$. Now consider a_2. If a_2 envies a_1 in $\boldsymbol{\pi}'$ then by swapping their bundles, we can obtain an envy-free allocation. This contradicts the fact that $\boldsymbol{\pi}$ is min-OWA envy optimal. Hence, a_2 still does not envy a_1 in $\boldsymbol{\pi}'$, and thus is also EFX$_0$ obviously. Since in $\boldsymbol{\pi}'$ a_2 is still envy-free and the pairwise envy from a_1 to a_2 has not changed, $\boldsymbol{\pi}'$ is still min-OWA envy optimal. The complexity is linear in the number of objects since we have to implement the transfer of zero-valued objects to a_1's bundle.

On Example 3, this means that a_1 should receive o_2. This adjustment is inefficient: by construction, it returns an allocation which is Pareto-dominated by the original min OWA envy optimal allocation. Intuitively, it can be seen as the price to pay to get EFX$_0$: by assigning those items that the agent does not value to her, the mechanism offers the strongest possible fairness guarantees.

4.2 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 relation between min OWA and EF1. Unfortunately, we notice that in the general case these two sets can be disjoint, ie. there are instances for which no allocation is both EF1 and min-OWA, for any weight vector:

Proposition 3. $\exists I \in \mathcal{I} : EF1(I) \cap \forall\text{-}min\ OWA(I) = \emptyset$

Proof. Let us consider the add-MARA instance with 4 agents and 5 objects:

	o_1	o_2	o_3	o_4	o_5
a_1	[20]	2	2	2	4
a_2	20	[2]	[2]	[2]	4
a_3	13	1	1	1	14
a_4	0	0	0	0	[30]

In order to prove the proposition we will show that the squared allocation is the only min OWA envy allocation (for any given weight vector) and that it is (obviously) not EF1. First note that as a_1 and a_2 have similar preferences the allocation derived from the squared allocation where we swap the bundles of these agents will be the same in terms of Lorenz envy vector. The squared allocation has a vector of envy $\boldsymbol{e} = \langle 0, 14, 14, 0\rangle$ and $L(\boldsymbol{e}) = \langle 14, 28, 28, 28\rangle$. First consider the allocations in which a_4 does not possess o_5. We have $\boldsymbol{e_1} = \langle e_1, e_2, e_3, 30\rangle$ and $L(\boldsymbol{e_1}) = \langle 30, L_2, L_3, L_4\rangle$ with L_2, L_3, L_4 being greater than or equal to 30. $\boldsymbol{e_1}$ is thus strongly Lorenz dominated by $\boldsymbol{e}$. Let us now consider the other possible allocations (in which a_4 possesses o_5): if a_3 has o_1 instead of a_1 then $\boldsymbol{e_2} = \langle 20, 14, 1, 0\rangle$ and $L(\boldsymbol{e_2}) = \langle 20, 34, 35, 35\rangle$. $\boldsymbol{e_2}$ is thus strongly Lorenz dominated by $\boldsymbol{e}$. Finally, we focus on allocations in which a_3 has one to three items from the set of objects $\{o_1, o_2, o_3\}$. If a_3 has one of these items we have $\boldsymbol{e_3} = \langle 0, 16, 13, 0\rangle$ and $L(\boldsymbol{e_3}) = \langle 16, 29, 29, 29\rangle$. If a_3 has two of these items we have $\boldsymbol{e_4} = \langle 0, 18, 12, 0\rangle$ and $L(\boldsymbol{e_4}) = \langle 18, 30, 30, 30\rangle$. Finally if a_3 has all these items we have $\boldsymbol{e_5} = \langle 0, 20, 11, 0\rangle$ and $L(\boldsymbol{e_5}) = \langle 20, 31, 31, 31\rangle$. All $\boldsymbol{e_3}$ $\boldsymbol{e_4}$ and $\boldsymbol{e_5}$ are strongly Lorenz dominated by $\boldsymbol{e}$. As we know that minimizing fair OWA of a vector is consistent with the Lorenz dominance (see Theorem 1), it means that if a solution strongly Lorenz dominates another, then its fair OWA value will be strictly lower (in a minimization problem such as ours) for any non-creasing weight. We can then conclude that the squared allocation is indeed the only min OWA envy one and it is not EF1.

However, a significant number of experiments actually suggest that for almost any instance, some EF1 allocation is also min-OWA, either for the weight vector $\langle 1, 0, \ldots 0\rangle$, or for the weight vector $\langle 1, 1, \ldots 1\rangle$. Moreover, we have a positive result in the restricted domain where agents have binary utilities.

Proposition 4. $\forall I \in \mathcal{I}^b : EFX_0(I) \cap \forall\text{-}min\ OWA(I) \neq \emptyset$

Proof. First note that if the instance is EF then the min OWA envy allocation will be EF and thus EF1 and the proof concludes. Hence we will consider instances that are not EF. As we consider binary utilities, we know thanks to [8] that an EFX_0 allocation always exists. We can easily notice that any such allocation is such that the envy of each agent is at most 1. Hence, as with the weight vector $\langle 1, 0, \ldots 0\rangle$ the OWA envy value of an EFX_0 allocation is 1 (as we supposed no EF allocation exists), it is the minimum OWA envy value possible. It can thus be returned by minimizing the OWA envy value.

5 Experimental results

We drew some experiments to compare the performances of the allocations obtained by min OWA envy with the Maximization of Nash Welfare. More precisely we implemented the linearization described in [7] that returns an allocation approximating MNW but closely enough to keep interesting properties such as EF1 and Pareto Optimality. As we have seen through this paper the range of possibilities offered by the fact that OWA is parameterized is interesting. We will see how three different weights $\boldsymbol{\alpha_1} = \langle 1, 0 \dots 0 \rangle$, $\boldsymbol{\alpha_2} = \langle \frac{1}{2}, \frac{1}{4}, \dots \frac{1}{2^n} \rangle$ and $\boldsymbol{\alpha_3} = \langle 1, 1, \dots 1 \rangle$ compare to each other. $\boldsymbol{\alpha_1}$ and $\boldsymbol{\alpha_3}$ correspond to respectively minimize the max envy and the sum of the envies. $\boldsymbol{\alpha_2}$ is somewhere in the middle of those two extrema with a strictly decreasing weight vector.

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 Mixed Integer Programs[3]. We have tested our methods on two types of instances: Spliddit instances [13] and synthetic instances under uniformly distributed commensurable preferences (that is, for each agent a_i and object o_j, utilities are drawn i.i.d. following the uniform distribution on some interval $[x, y]$ and such that the utilities of each agent sums to $5m$).

We evaluate the performances of the OWA envy minimization outcome for both types of instances through the following criterion: EF, EFX_0, EFX, EF1 and Pareto dominance. Tables 1 and 2 present the percentage of min OWA envy outcomes that satisfy each criterion. We also study how the vector of weights of the OWA influences the characteristics of the outcomes. The computation time (in seconds) of each approach is also mentioned. We recall the strong connections between the 4 first fairness notions as EF $\implies$ EFX_0 $\implies$ EFX $\implies$ EF1. As it can be checked in Tables 1 and 2, the percentage of EF allocations should always be lower than or equal to the number of EFX_0 ones which should be lower than or equal to the number of EFX allocations and so on.

5.1 Spliddit instances

Our first set of experiments has been performed on real-world data from the fair division website Spliddit [13]. 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 MIPs minimizing the OWA envy with the three different weights' vectors described above with a timeout of 1 minute (after this duration the best current solution, if it exists, is returned) we were able to solve all the instances to optimal. The results of these experiments are presented in Table 1. The first three columns respectively correspond to the results of minimization of the OWA envy with respectively $\boldsymbol{\alpha_1}$, $\boldsymbol{\alpha_2}$ and $\boldsymbol{\alpha_3}$, while the fourth column presents the results of the optimization of MNW.

Minimizing the OWA envy provably returns an EF allocation if there exists one. Hence, among the Spliddit instances 65.4 % are envy-free. Note that only

[3] The code is available at `https://gitlab.com/MrPyrom/balancing-envy`

Table 1. Performances for minimizing the OWA envy (with weights $\boldsymbol{\alpha_1}$, $\boldsymbol{\alpha_2}$, $\boldsymbol{\alpha_3}$) or maximizing the Nash Welfare on Spliddit instances

	$\boldsymbol{\alpha_1}$	$\boldsymbol{\alpha_2}$	$\boldsymbol{\alpha_3}$	MNW
%EF	65.4	65.4	65.4	57.2
%EFX$_0$	90.0	93.0	92.7	90.9
%EFX	98.5	99.4	99.0	94.9
%EF1	99.4	99.8	99.3	100
%Pareto	77.1	78.7	79.2	100
%EF+PO	45.7	45.6	46.0	57.2
time(s)	$3.5*10^{-3}$	$5.7*10^{-7}$	$6.9*10^{-7}$	$1.1*10^{-6}$

57.2 % of the allocations returned by MNW are EF which means that for around 8.2 % of the Spliddit instances, an EF allocation exists but MNW failed to return it. Moreover, without any surprise as Pareto optimality (PO) of the MNW allocations is guaranteed, minimizing OWA envy returns fewer PO allocations than MNW. However, around slightly less than 80% of the min OWA envy allocations are PO. It is also guaranteed that MNW returns an EF1 solution. However, we can observe, for every weight, that more than 99% of the allocations returned by min OWA envy are EF1. This balances the negative result in Proposition 3. Moreover, it can be very interestingly observed that the percentage of EFX$_0$ is greater for $\boldsymbol{\alpha_2}$ and $\boldsymbol{\alpha_3}$ than for MNW. The same holds for the percentage of EFX but for the 3 weights' vectors and by a more noticeable margin of around 5%. However, MNW performs slightly better than min OWA when we consider EF alongside with PO. Finally, we can see that all the optimization programs run very quickly in average with a slightly longer time for $\boldsymbol{\alpha_1}$.

5.2 Synthetic instances

For each couple ($|\mathcal{N}|$,$|\mathcal{O}|$) from $(3, 4)$ to $(10, 12)$, we generated 100 synthetic add-MARA instances with uniformly distributed preferences. We then ran the four optimization methods described above on the generated instances. We considered such couples of values in order to produce settings where few EF allocations exist as suggested in [10]. Although it is interesting to consider EF instances to compare with MNW, minimizing OWA envy is even more relevant when no EF allocation exists. Due to lack of space, Table 2 presents the results for only 4 couples (n, m) but similar trends can be observed for the other couples of values. As witnessed for the Spliddit instances, MNW often fails to return an EF allocation even when there exists one. As shown in Table 2, the number of EF allocations missed by MNW can be quite important as shown by the gap between the percentage of EF allocations returned by min OWA envy and the percentage for MNW. This is exemplified in Table 2 for 2 agents and 5 agents where the gap is respectively of 16% and 31%. Even more significantly, it turns out that min OWA outperforms MNW when we consider EF together with PO. Once again and in an even stronger way than for the Spliddit instances, these results heavily balance the result of Proposition 3: in practice the allocations returned by the min OWA envy were always EF1. Concerning EFX_0 and EFX we also obtained very positive results. Indeed, min OWA envy returns around 10% more EFX_0 and EFX instances than MNW. Note that we confirm Proposition 1 as we have 100% of EFX allocations when $n = 2$. Note that we did not adjust the allocation returned by the min OWA optimization to break ties as discussed in the proof of Proposition 2. Thus, we get 97% of EFX_0 but this percentage could be even higher. However, these positive results about EF, EFX_0 and EFX come with a price on efficiency as we can see that PO is not guaranteed and the percentage gets lower as the number of agents increases but is still above 60% for $\boldsymbol{\alpha_2}$ and $\boldsymbol{\alpha_3}$. This highlights the inherent compromise and tension between efficiency and fairness. Besides, as it was the case for the Spliddit instances we can see that the computation is overall quite fast. We can notice that the MNW computation never surpasses 0.02 seconds whereas for 10 agents, min OWA envy optimization is slightly faster than a second for $\boldsymbol{\alpha_1}$ and $\boldsymbol{\alpha_3}$ and around 2 seconds for $\boldsymbol{\alpha_2}$. Finally, we can see that the three different weights considered here lead to quite similar performances. We can globally notice more encouraging results for $\boldsymbol{\alpha_3}$ except for EFX. However, keep in mind that the advantage of using a parameterized function is its rich expressiveness so we could see our method as a combination of the results of the 3 weights.

Table 2. Performances for minimizing the OWA envy (with weights $\boldsymbol{\alpha_1}$, $\boldsymbol{\alpha_2}$, $\boldsymbol{\alpha_3}$) or maximizing the Nash Welfare on synthetics instances (as a function of the number of agents and objects (n, m) ($\epsilon \leq 10^{-3}$)).

	(2,3)				(5,7)				(8,10)				(10,12)			
	$\boldsymbol{\alpha_1}$	$\boldsymbol{\alpha_2}$	$\boldsymbol{\alpha_3}$	MNW	$\boldsymbol{\alpha_1}$	$\boldsymbol{\alpha_2}$	$\boldsymbol{\alpha_3}$	MNW	$\boldsymbol{\alpha_1}$	$\boldsymbol{\alpha_2}$	$\boldsymbol{\alpha_3}$	MNW	$\boldsymbol{\alpha_1}$	$\boldsymbol{\alpha_2}$	$\boldsymbol{\alpha_3}$	MNW
%EF	74	74	74	58	48	48	48	17	10	10	10	1	1	1	1	0
%EFX_0	97	97	97	88	96	96	96	88	88	86	88	78	72	82	83	80
%EFX	100	100	100	92	97	97	98	91	98	96	93	85	87	95	92	84
%EF1	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100
%Pareto	100	100	100	100	73	76	72	100	64	66	67	100	51	62	64	100
%EF+PO	74	74	74	58	33	34	32	17	5	6	5	1	1	1	1	0
time(s)	ϵ	ϵ	ϵ	ϵ	0.02	0.02	0.02	0.01	0.1	0.4	0.1	0.04	0.5	2.5	0.7	0.07

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. This generalizes several approaches using various definitions of degree of envies, which can be captured by adequate weight vector. In practice, we put a special focus on the egalitarian variant (minimizing the highest envy), the utilitarian variant (minimizing the sum of envies), and the compromise consisting of using the fair vector of decreasing weights. After implementing a MIP to compute min OWA allocations, we unveil several connections between the min OWA allocation and other famous fairness measures. In particular, we compare our approach with the alternative relaxations consisting of seeking "envy-freeness up to some/any good". Some of our conclusions show that these approaches correspond to very different perspectives: we show in particular that no algorithm minimizing a fair OWA can ever guarantee to return an EF1 (and thus nor EFX) allocation. This is however balanced by the fact that it never occured in our experiments. Indeed, even in the very few cases for which the min OWA allocation was not EF1 we easily found a weight for which it was the case. This raises the question of choosing the appropriate weight vector for example by elicitating it. We left that question open for now. Indeed, we also ran some experiments to test the performances of our method and compared it with other allocation protocols. The results are extremely encouraging. Our min OWA approaches do very well (in particular regarding the likelihood to return an EFX allocation, which may be somewhat paradoxical given our previous remarks) in terms of fairness, both on real Spliddit instances and randomly generated ones. In comparison, Nash social welfare –despite its guarantee to return an EF1 allocation– is dominated on that respect, as well as on the likelihood to return an EF and Pareto optimal allocation.

References

1. Aleksandrov, M., Ge, C., Walsh, T.: Fair division minimizing inequality. In: Oliveira, P.M., Novais, P., Reis, L.P. (eds.) Progress in Artificial Intelligence, 19th EPIA Conference on Artificial Intelligence, EPIA 2019, Vila Real, Portugal. LNCS, vol. 11805, pp. 593–605. Springer (2019)
2. Amanatidis, G., Birmpas, G., Filos-Ratsikas, A., Hollender, A., Voudouris, A.A.: Maximum nash welfare and other stories about EFX. In: Bessiere, C. (ed.) Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020. pp. 24–30. ijcai.org, Yokohama, Japan (2020)
3. Amanatidis, G., Birmpas, G., Markakis, V.: Comparing approximate relaxations of envy-freeness. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018. pp. 42–48. Stockholm, Sweden (2018)
4. Bansal, N., Sviridenko, M.: The Santa Claus problem. In: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing. pp. 31–40. STOC '06, ACM, New York, NY, USA (2006)
5. Bouveret, S., Lemaître, M.: Characterizing conflicts in fair division of indivisible goods using a scale of criteria. Autonomous Agents and Multi-Agent Systems **30**(2), 259–290 (2016)
6. Budish, E.: The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. Journal of Political Economy **119**(6), 1061–1103 (2011)
7. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of Maximum Nash Welfare. In: Proceedings of the 2016 ACM Conference on Economics and Computation. pp. 305–322. EC '16, ACM, New York, NY, USA (2016)

8. Chaudhury, B.R., Garg, J., Mehlhorn, K.: EFX exists for three agents. In: Proceedings of the 21st ACM Conference on Economics and Computation. p. 1–19. EC '20, Association for Computing Machinery, New York, NY, USA (2020)
9. Chevaleyre, Y., Endriss, U., Maudet, N.: Distributed fair allocation of indivisible goods. Artif. Intell. **242**, 1–22 (2017)
10. Dickerson, J.P., Goldman, J., Karp, J., Procaccia, A.D., Sandholm, T.: The computational rise and fall of fairness. In: Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI-14). pp. 1405–1411. AAAI Press, Québec City, Québec, Canada (Jul 2014)
11. Endriss, U.: Reduction of economic inequality in combinatorial domains. In: Gini, M.L., Shehory, O., Ito, T., Jonker, C.M. (eds.) International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '13, Saint Paul, MN, USA, May 6-10, 2013. pp. 175–182. IFAAMAS (2013)
12. Foley, D.K.: Resource allocation and the public sector. Yale Economic Essays **7**(1), 45–98 (1967)
13. Goldman, J., Procaccia, A.D.: Spliddit: Unleashing fair division algorithms. SIGecom Exch. **13**(2), 41–46 (Jan 2015)
14. Heinen, T., Nguyen, N.T., Rothe, J.: Fairness and rank-weighted utilitarianism in resource allocation. In: Walsh, T. (ed.) Algorithmic Decision Theory. pp. 521–536. Springer International Publishing, Cham (2015)
15. Kyropoulou, M., Suksompong, W., Voudouris, A.A.: Almost envy-freeness in group resource allocation. Theoretical Computer Science **841**, 110 – 123 (2020)
16. Lesca, J., Minoux, M., Perny, P.: The fair OWA one-to-one assignment problem: NP-Hardness and polynomial time special cases. Algorithmica **81**(1), 98–123 (2019)
17. Lipton, R., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of divisible goods. In: Proceedings of the 5th ACM Conference on Electronic Commerce (EC-04). pp. 125–131. ACM, New York, NY (May 2004)
18. Nguyen, T.T., Rothe, J.: How to decrease the degree of envy in allocations of indivisible goods. In: Perny, P., Pirlot, M., Tsoukiàs, A. (eds.) Algorithmic Decision Theory. pp. 271–284. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
19. Ogryczak, W., Śliwiński, T.: On solving linear programs with the ordered weighted averaging objective. European Journal of Operational Research **148**, 80–91 (2003)
20. Perny, P., Spanjaard, O.: An axiomatic approach to robustness in search problems with multiple scenarios. In: Meek, C., Kjærulff, U. (eds.) UAI '03, Proceedings of the 19th Conference in Uncertainty in Artificial Intelligence, Acapulco, Mexico, August 7-10 2003. pp. 469–476. Morgan Kaufmann (2003)
21. Plaut, B., Roughgarden, T.: Almost envy-freeness with general valuations. In: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms. p. 2584–2603. SODA '18, USA (2018)
22. Rawls, J.: A Theory of Justice. Harvard University Press, Cambridge, Mass. (1971)
23. Schneckenburger, S., Dorn, B., Endriss, U.: The Atkinson inequality index in multiagent resource allocation. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil. pp. 272–280. ACM (2017)
24. Weymark, J.A.: Generalized Gini inequality indices. Mathematical Social Sciences **1**(4), 409 – 430 (1981)
25. Yager, R.R.: On ordered weighted averaging aggregation operators in multicriteria decision making. IEEE Transactions on Systems, Man, and Cybernetics **18**, 183–190 (1988)