

Minimizing and balancing envy among agents using Ordered Weighted Average

Parham Shams¹ Aurélie Beynier¹ Sylvain Bouveret² Nicolas Maudet¹

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

`{aurelie.beynier,nicolas.maudet,parham.shams}@lip6.fr`

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

`sylvain.bouveret@imag.fr`

Keywords : *Computational Social Choice, resource allocation, fair division, 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 [4]. 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 [5] and envy-freeness up to any good [2]. 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.* [5] or Endriss *et al.* [3] for instance. Several approximation algorithms dedicated to minimize these measures were subsequently designed, see e.g. [6].

2 Problem definition

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 $u_i(\pi) \stackrel{\text{def}}{=} \sum_{o_k \in \pi} w(i, o_k)$, 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 [5, for instance] and is considered a good compromise between expressivity and conciseness.

3 Contribution

We first 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. In order to do so, we translate the OWA minimization problem into the following mixed integer program where the z_i^j variables encode whether agent a_i possesses object o_j or not, the e_i variables represent the envy of agent a_i while the b_i^k and r_k variables help us linearize the OWA :

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

Then, we relate this criterion to other fairness notions (such as max-min fair share, envy-freeness up to one good and envy-freeness up to any good) in the general and the 2 agents settings. Moreover, we study properties of the allocations obtained by minimizing the OWA of the envy vector and why commensurability between agents is of paramount importance. Finally, we present some experimental results investigating the fairness of min OWA envy solutions.

References

- [1] 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 2018, July 13-19, 2018, Stockholm, Sweden.*, pages 42–48, 2018.
- [2] 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*, pages 305–322, New York, NY, USA, 2016. ACM.
- [3] Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. Allocating goods on a graph to eliminate envy. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI-07)*, Vancouver, British Columbia, Canada, July 2007.
- [4] Duncan K. Foley. Resource allocation and the public sector. *Yale Economic Essays*, 7(1) :45–98, 1967.
- [5] 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)*, pages 125–131, New York, NY, May 2004. ACM.
- [6] Trung Thanh Nguyen and Jörg Rothe. How to decrease the degree of envy in allocations of indivisible goods. In Patrice Perny, Marc Pirlot, and Alexis Tsoukiàs, editors, *Algorithmic Decision Theory*, pages 271–284, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

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