

Characterizing Conflicts in Fair Division of Indivisible Goods Using a Scale of Criteria

Sylvain Bouveret
LIG – Grenoble INP, France
sylvain.bouveret@imag.fr

Michel Lemaître
Formerly Onera Toulouse, France
michel.lemaitre.31@gmail.com

ABSTRACT

We investigate five different fairness criteria in a simple model of fair resource allocation of indivisible goods based on additive preferences. We show how these criteria are connected to each other, forming an ordered scale that can be used to characterize how conflicting the agents' preferences are: the less conflicting the preferences are, the more demanding criterion this instance will be able to satisfy, and the more satisfactory the allocation will be. We analyze the computational properties of the five criteria, give some experimental results about them, and further investigate a slightly richer model with k -additive preferences.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent Systems*; J.4 [Computer Applications]: Social and behavioral sciences—*Economics*

General Terms

Economics, Theory

Keywords

Computational social choice, resource allocation, fair division, indivisible goods, preferences.

1. INTRODUCTION

The problem of fairly allocating some resources to a set of economically motivated agents is an important and frequent problem. Fair division of indivisible goods in particular, on which we focus in this paper, arises in a wide range of real-world applications, including auctions, divorce settlements, airport traffic management, spatial resource allocation [13], fair scheduling, allocation of tasks to workers, articles to reviewers, courses to students [18].

More precisely, we study here a simple model of fair division of indivisible goods based on the following assumptions. (i) A set of indivisible goods which will be called objects must be distributed among a set of agents. (ii) Agents have numerical additive preferences over the objects (except in Section 7 where we consider a more general model). (iii)

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

The allocation process is centralized, that is, it is decided by a neutral arbitrator or computation, taking into account only agents' preferences, in a single step. (iv) No monetary transfer is possible between agents. Even if this model seems restrictive (especially assumption (ii)), it has been largely investigated and offers a natural compromise between simplicity and expressiveness [3; 14; 2; 1; 9; 15, among others].

Defining fairness and evaluating it is a critical issue in fair allocation mechanisms. Two main options are available. The first one consists in defining a collective utility function (CUF) aggregating individual agents' utilities. The outcome of a well chosen CUF, when applied to individual utilities, reflects the fairness (and possibly other desirable criteria) of a given allocation. The arbitrator just looks for an allocation maximizing this CUF. This is the approach chosen *e.g.* by [2, 1] with the egalitarian (min) CUF for the "Santa-Claus" problem. The other option defines fairness as a Boolean (logical) criterion to satisfy. This is the approach followed by Lipton *et al.* [14] among others for envy-freeness.

In this article, we mainly investigate the second option. While most papers in fair division focus on a specific criterion, here we consider five of them and investigate their connection to each other. Four of these criteria are classical or already known, namely: max-min fair-share (MFS), proportional fair-share (PFS), envy-freeness (EF) and competitive equilibrium from equal incomes (CEEI), and we introduce an original one: min-max fair-share (mFS). All these criteria have a natural interpretation and a very appealing quality: they do not need a common scale of agents' utilities.¹ We show in this paper that these criteria actually form a linear scale that can be used to characterize formally (i) the level of fairness of a given allocation, and (ii) to which extent it will be possible, for a given resource allocation instance, to find a satisfactory (fair) allocation: the less conflicting the agents' preferences are about the objects, the more demanding criterion the central authority will be able to satisfy, and the fairer the resulting allocation will be.

This article is structured as follows. Section 2 describes the model: fair division of indivisible objects under numerical additive preferences. The scale of five criteria characterizing the fairness of an allocation, as well as associated computational complexity results are exposed in Section 3. We go back to the CUF approach in Section 4 to connect the important egalitarian CUF to the scale of criteria. Section 5 is devoted to interesting restricted cases. Some experimental results on the scale of criteria are presented in Section

¹Whereas most CUF – except Nash – only make sense if the utilities are expressed on a common scale or normalized.

6. Extending the model to k -additive preferences, Section 7 presents a quite different perspective.²

2. MODEL

Let $\mathcal{A} = \{1, \dots, n\}$ be a set of *agents* and $\mathcal{O} = \{1, \dots, m\}$ be a set of indivisible *objects*. An *allocation* of the objects to the agents is a vector $\vec{\pi} = \langle \pi_1, \dots, \pi_n \rangle$, where $\pi_i \subseteq \mathcal{O}$ is the *bundle* of objects allocated to agent i , called agent i 's *share*. An allocation $\vec{\pi}$ is *admissible* if and only if it satisfies the two following conditions: (i) $i \neq j \Rightarrow \pi_i \cap \pi_j = \emptyset$ (each object is allocated to at most one agent) and (ii) $\cup_{i \in \mathcal{A}} \pi_i = \mathcal{O}$ (all the objects are allocated). \mathcal{F} is the set of admissible allocations. All the allocations considered in forthcoming definitions and propositions are implicitly admissible.

We make two usual assumptions concerning the way agents express their preferences. First, we consider that they are expressed numerically by a *utility function* $u_i : 2^{\mathcal{O}} \rightarrow \mathbb{R}^+$ specifying, for each agent i , the satisfaction $u_i(\pi)$ she enjoys if she receives bundle π : this is the utilitarian model [16]. Second, we consider (except in Section 7) that the agents' preferences are *additive*, which means that the utility function of an agent is defined as follows: $u_i(\pi) \stackrel{\text{def}}{=} \sum_{l \in \pi} w(i, l)$, where $w(i, l)$ is the *weight* given by agent i to object l .

Adapting the terminology from Chevaleyre *et al.* [7], we will hence define an *additive MultiAgent Resource Allocation* instance (add-MARA instance for short) as a triple $\langle \mathcal{A}, \mathcal{O}, w \rangle$, where \mathcal{A} is a set of agents, \mathcal{O} is a set of objects, and $w : \mathcal{A} \times \mathcal{O} \rightarrow \mathbb{R}^+$ is the weight function.

In the following, indices i and j always refer to agents, and l to objects. In examples, we use a matrix representation W for the weight function w : the element at row i and column l represents the weight $w(i, l)$. Finally, we write \mathcal{I} the set of all add-MARA-instances.

The basic notions of computational complexity [19] are supposed to be well-known by the reader: P and NP refer to the two usual complexity classes; Σ_2^P is the class of problems that can be solved in non-deterministic polynomial time by a Turing machine augmented by an NP oracle.

3. FIVE FAIRNESS CRITERIA

This section introduces five fairness criteria, among which the two most prominent ones are proportionality [22] and envy-freeness [10]. We write $\vec{\pi} \models \mathcal{C}$ if the allocation $\vec{\pi}$ satisfies criterion \mathcal{C} ; $\mathcal{I}_{\mathcal{C}}$ denotes the set of add-MARA instances admitting at least one allocation satisfying criterion \mathcal{C} . Beyond fairness criteria, we will also deal with Pareto-efficiency: an allocation $\vec{\pi}$ is Pareto-efficient if no allocation $\vec{\pi}'$ is such that $u_i(\pi'_i) \geq u_i(\pi_i)$ for all i , with at least one strict inequality.

3.1 Max-min fair-share

An important fairness criterion in resource allocation problems is proportional fair-share (discussed in details in Section 3.2). This criterion, coined by Steinhaus [22] in the context of continuous fair division (cake-cutting) problems, states that each agent should get from the allocation at least $1/n$ of the total utility she would have received if she were alone. However, when one deals with indivisible objects, it is often too demanding: consider for example a problem with

²Due to lack of space, most proofs are omitted. Complete proofs can be found at <http://recherche.noiraudes.net/resources/papers/AAMAS14-full.pdf>.

one object and two agents, where obviously no allocation can give her fair share to each agent. That is why it has been recently adapted to this context by Budish [6], which defines the *max-min fair share*, whose original definition is purely ordinal, but which can be defined in our setting as follows:

DEFINITION 1. Let $(\mathcal{A}, \mathcal{O}, w)$ be an add-MARA instance. The max-min fair share of agent i for this instance is

$$u_i^{\text{MFS}} \stackrel{\text{def}}{=} \max_{\vec{\pi} \in \mathcal{F}} \min_{j \in \mathcal{A}} u_j(\pi_j)$$

We say that the allocation $\vec{\pi}$ satisfies the max-min fair-share criterion, written $\vec{\pi} \models \text{MFS}$, if $u_i^{\text{MFS}} \leq u_i(\pi_i)$ for all i (each agent obtains at least her max-min fair share in $\vec{\pi}$).

EXAMPLE 1. Let us consider the 2 agents / 4 objects instance defined by the following weight matrix:

$$W = \begin{pmatrix} *7 & 2 & 6 & *10 \\ 4 & *7 & *7 & 7 \end{pmatrix}$$

$u_1^{\text{MFS}} = 12$ (with share $\{2, 4\}$); $u_2^{\text{MFS}} = 11$ (with share $\{1, 2\}$). The starred allocation $\langle \{1, 4\}, \{2, 3\} \rangle$ satisfies MFS.

The max-min fair-share of an agent is the maximal utility that she can hope to get from an allocation if all the other agents have the same preferences as her, when she always receive the worst share (it is the best of the worst shares).

The max-min fair-share is the minimal amount of utility that an agent could feel to be entitled to, based on the following argument: if all the other agents have the same preferences as me, there is at least one allocation that gives me this utility, and makes every other agent better off; so there is no reason to give me less. It is also the maximum utility an agent can get for sure in the allocation game "I cut, I choose last": the agent proposes an allocation (that we will refer to as a *max-min cut*) and leaves all the other ones choose one share before taking the remaining one.

The max-min fair-share level is loosely connected to a result from [12], recently refined by [15], which establishes a worst case guarantee on the utility an agent can have. However, this guarantee only depends on the maximum weight of an agent, and so is not very informed, often being just 0.

Beyond its appealing formulation, max-min fair-share has a computational drawback: the computation of the max-min fair-share u_i^{MFS} itself for a given agent is complex. More precisely, the following decision problem is NP-complete:

Problem 1 [MFS-COMP]	
Input:	An add-MARA instance $\langle \mathcal{A}, \mathcal{O}, w \rangle$, an agent i , an integer K .
Question:	Do we have $u_i^{\text{MFS}} \geq K$?

PROPOSITION 1. [MFS-COMP] is NP-complete, $\forall n \geq 2$.

PROOF. Membership to NP is obvious. NP-hardness can be proved by reduction from the [PARTITION] problem: given a set $\mathcal{X} = \{s_1, \dots, s_n\}$ of integers whose sum is $2L$, is it possible to find a partition $(\mathcal{X}_1, \mathcal{X}_2)$ of \mathcal{X} such that $\sum_{s_i \in \mathcal{X}_1} s_i = \sum_{s_i \in \mathcal{X}_2} s_i$?

From an instance of [PARTITION], we create an instance of [MFS-COMP] with 2 agents and n objects. The agents' preferences are identical and defined as $w(1, l) = w(2, l) = s_l$. Integer K is defined as L , completing the reduction.³ \square

³We use here a very similar idea to the one used by [14, p4].

Let us now focus on the problem [MFS-EXIST] of determining, for a given add-MARA instance, if there is an allocation satisfying the max-min fair-share criterion. Strong evidences led us to think that every add-MARA instance had at least one such allocation: it is true in many restricted cases (see Section 5), and no counterexample was found in thousands of randomly generated instances (see Section 6). However, surprisingly, Procaccia and Wang [20] have recently proved that there actually exists add-MARA instances for which there is no allocation satisfying max-min fair-share. Put in other words, we have $\mathcal{I}_{\text{MFS}} \subsetneq \mathcal{I}$.

The complexity of [MFS-EXIST] is still open. We only know that this problem belongs to Σ_2^P , because it can be solved by the following non-deterministic polynomial algorithm: (i) guess an allocation $\vec{\pi}$; (ii) for all $i \in A$, compute the max-min fair share u_i^{MFS} of agent i ; (iii) for all $i \in A$, check that $u_i(\pi_j) \geq u_i^{\text{MFS}}$.

3.2 Proportional fair-share

The aforementioned concept of proportional fair-share was originally defined not on the utilities but on the resources themselves. A lot of authors have since given a natural utilitarian interpretation of this notion, like the one that follows:

DEFINITION 2. Let $\langle A, \mathcal{O}, w \rangle$ be an add-MARA instance. The proportional fair-share of agent i for this instance is

$$u_i^{\text{PFS}} \stackrel{\text{def}}{=} \frac{1}{n} u_i(\mathcal{O}) = \frac{1}{n} \sum_{l \in \mathcal{O}} w(i, l).$$

Allocation $\vec{\pi}$ satisfies the proportional fair-share criterion, written $\vec{\pi} \models \text{PFS}$, if $u_i^{\text{PFS}} \leq u_i(\pi_i)$ for all i (that is, each agent obtains at least her proportional fair-share in $\vec{\pi}$).

The proportional fair-share of an agent represents the maximal utility she would receive from a virtual perfectly equitable allocation if all the agents had exactly the same preferences as her (for all $i, j, l : w(j, l) = w(i, l)$). Moreover, in the virtual allocation obtained by dividing each object into n parts, each one allocated to a different agent, each single agent would enjoy exactly her proportional fair-share.

This is obviously stronger than max-min fair-share:

PROPOSITION 2. Let $\langle A, \mathcal{O}, w \rangle$ be an add-MARA instance. We have $u_i^{\text{MFS}} \leq u_i^{\text{PFS}}$, for all $i \in A$. So, for all $\vec{\pi}$, we have $\vec{\pi} \models \text{PFS} \implies \vec{\pi} \models \text{MFS}$, and thus $\mathcal{I}_{\text{PFS}} \subset \mathcal{I}_{\text{MFS}}$.

PROOF. Let $\vec{\pi}$ be an allocation and i an agent. We have $\sum_{j \in A} u_i(\pi_j) = u_i(\mathcal{O})$, and

$$\min_{j \in A} u_i(\pi_j) \leq \frac{1}{n} \sum_{j \in A} u_i(\pi_j) = \frac{1}{n} u_i(\mathcal{O}) = u_i^{\text{PFS}}$$

Taking the max over the set of allocations on both sides of the latter inequality gives the result: $u_i^{\text{MFS}} \leq u_i^{\text{PFS}}$. \square

The inclusion in Proposition 2 is strict: in an instance with two agents and one object, every allocation satisfies max-min fair-share, but none satisfies proportional fair-share.

Contrary to max-min fair-share, computing the proportional fair-share for a given agent is easy. However, the problem [PFS-EXIST] of determining whether a given add-MARA instance has an allocation satisfying proportional fair-share is computationally hard (proof similar to Prop. 1):

PROPOSITION 3. [PFS-EXIST] is NP-complete, $\forall n \geq 2$.

3.3 Min-max fair-share

The min-max fair-share criterion that we now introduce is, to the best of our knowledge, original. It can be seen as the symmetrical version or the max-min fair-share criterion.

DEFINITION 3. Let $\langle A, \mathcal{O}, w \rangle$ be an add-MARA instance. The min-max fair-share of agent i for this instance is

$$u_i^{\text{mFS}} \stackrel{\text{def}}{=} \min_{\vec{\pi} \in \mathcal{F}} \max_{j \in A} u_i(\pi_j)$$

Allocation $\vec{\pi}$ satisfies the min-max fair-share criterion, written $\vec{\pi} \models \text{mFS}$, if $u_i^{\text{mFS}} \leq u_i(\pi_i)$ for all i (that is, each agent obtains at least her min-max fair share in $\vec{\pi}$).

The min-max fair-share of an agent is the minimal utility that she can hope to get from an allocation if all the other agents have the same preferences as her, when she always receive the best share (it is the worst of the best shares). It is also the minimal utility that an agent can get for sure in the allocation game ‘‘Someone cuts, I choose first’’. The following result is the equivalent of Proposition 2 and is proved in a similar way:

PROPOSITION 4. Let $\langle A, \mathcal{O}, w \rangle$ be an add-MARA instance. We have $u_i^{\text{PFS}} \leq u_i^{\text{mFS}}$, for all $i \in A$. So, for all $\vec{\pi}$, we have $\vec{\pi} \models \text{mFS} \implies \vec{\pi} \models \text{PFS}$ and thus $\mathcal{I}_{\text{mFS}} \subset \mathcal{I}_{\text{PFS}}$.

This inclusion is strict, as the following example shows.

EXAMPLE 2. Let us consider the 3 agents / 3 objects instance defined by the following weight matrix:

$$W = \begin{pmatrix} 2 & 2 & *2 \\ 3 & *2 & 1 \\ *3 & 2 & 1 \end{pmatrix}$$

Obviously $u_i^{\text{PFS}} = 2$ for each agent. Thus the starred allocation gives to each agent her proportional fair-share. However, no allocation gives to each agent her min-max fair-share (which is 2 for agent 1 and 3 for the other ones).

Exactly like the max-min fair-share, and for similar reasons, the computation of the min-max fair-share for a given agent is hard. More precisely, with [MFS-COMP] being the equivalent for min-max fair-share of decision Problem 1, the following proposition holds.

PROPOSITION 5. [MFS-COMP] is coNP-complete, $\forall n \geq 2$.

The decision problem is coNP-complete because min-max fair-share is defined as a minimization, and that we want to know, as for the max-min fair-share, whether the min-max fair-share of a given agent is greater than a given threshold. The proof is very similar to the one of Proposition 1. The decision problem of determining whether there exists an allocation satisfying min-max fair-share is very likely to be hard, but its precise complexity remains unknown⁴.

3.4 Envy-freeness

Among all fairness criteria, envy-freeness [10] is probably the most prominent one.

DEFINITION 4. Let $\langle A, \mathcal{O}, w \rangle$ be an add-MARA instance. The allocation $\vec{\pi}$ is envy-free, written $\vec{\pi} \models \text{EF}$, when for all $i, j : u_i(\pi_i) \geq u_i(\pi_j)$ (no agent strictly prefers the share of another agent to her own share).

⁴All that we can say for sure is that this problem is in Σ_2^P .

A known fact (cited at least in some working papers) in that envy-freeness implies proportionality for additive preferences. The following proposition is actually a bit stronger:

PROPOSITION 6. *Any envy-free allocation gives to each agent her min-max fair-share. In other words, for all $\vec{\pi} : \vec{\pi} \models \text{EF} \implies \vec{\pi} \models \text{mFS}$, so $\mathcal{I}_{|\text{EF}} \subset \mathcal{I}_{|\text{mFS}}$.*

PROOF. Let $\vec{\pi}$ be an envy-free allocation. Then for all $i, j : u_i(\pi_i) \geq \max_{j \in \mathcal{A}} u_i(\pi_j)$ by definition. Since $\vec{\pi} \in \mathcal{F}$, $u_i(\pi_i) \geq \min_{\vec{\pi} \in \mathcal{F}} \max_{j \in \mathcal{A}} u_i(\pi_j) = u_i^{\text{mFS}}$. \square

The inclusion introduced in this proposition is again strict:

EXAMPLE 3. *Let us consider the 3 agents / 4 objects instance defined by the following weight matrix:*

$$W = \begin{pmatrix} *10 & 6 & 6 & 1 \\ 10 & *6 & *6 & 1 \\ 1 & 6 & 6 & *10 \end{pmatrix}$$

We have $u_i^{\text{mFS}} = 10$ for each agent, thus the starred allocation gives the min-max fair-share to every agent. Now suppose that there exists an envy-free allocation $\vec{\pi}$. This $\vec{\pi}$ should give the same utility to agent 1 and 2 since they have the same preferences (otherwise they would be envious): either $\vec{\pi}$ gives nothing to them, or it gives 6 to each of them. In both cases they envy agent 3. So there is no envy-free allocation for this instance.

3.5 CEEI

The last introduced criterion is a classical notion in microeconomics [17, for example]. It has, to the best of our knowledge, almost never been considered in computer science, with the notable exception of [18]. This criterion is based on the following argument: the sharing process should be considered as a search for an equilibrium between the supply (the set of objects, each one having a public price) and the demand (the agents' desires, each agent having the same budget for buying the objects). A competitive equilibrium is reached when the supply matches the demand. The fairness argument is very straightforward: prices and budgets are the same for everyone. Several variants of this notion exist; the following definition is adapted from Budish [6].

DEFINITION 5. *Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance, $\vec{\pi}$ an allocation, and $\vec{p} \in [0, 1]^m$ a price vector. A pair $(\vec{\pi}, \vec{p})$ is said to form a competitive equilibrium from equal incomes (CEEI), if for each agent i ,*

$$\pi_i \in \operatorname{argmax}_{\pi \subseteq \mathcal{O}} \{u_i(\pi) : \sum_{l \in \pi} p_l \leq 1\}.$$

In other words, π_i is one of the maximal shares that i can buy with a budget of 1, given that the price of object l is p_l . We say that the allocation $\vec{\pi}$ satisfies the CEEI criterion, written $\vec{\pi} \models \text{CEEI}$, if there exists a price vector \vec{p} such that $(\vec{\pi}, \vec{p})$ forms a CEEI.

EXAMPLE 4. *Let us consider the 2 agents / 4 objects instance defined by the following weight matrix:*

$$W = \begin{pmatrix} *7 & 2 & 6 & *10 \\ 7 & *6 & *8 & 4 \end{pmatrix}$$

The starred allocation forms a CEEI, together with price vector $\langle 0.8, 0.2, 0.8, 0.2 \rangle$.

The following proposition holds for a lot of continuous resource allocation instances (divisible goods, presence of monetary compensations...). It also holds in our discrete model:

PROPOSITION 7. *Every CEEI allocation is envy-free : for all $\vec{\pi} : \vec{\pi} \models \text{CEEI} \implies \vec{\pi} \models \text{EF}$. Therefore $\mathcal{I}_{|\text{CEEI}} \subset \mathcal{I}_{|\text{EF}}$.*

PROOF. Let $\vec{\pi}$ be a CEEI allocation, and suppose that $u_i(\pi_j) > u_i(\pi_i)$ (agent i envies j). Budgets and prices being the same for everyone, π_i is not the best share that agent i can buy, contradicting the definition of the CEEI. \square

The CEEI also has the following interesting property:

PROPOSITION 8. *When the agents' preferences are strict (i.e distinct shares have distinct utilities), every CEEI allocation is Pareto-efficient.*

PROOF. Let $(\vec{\pi}, \vec{p})$ be a CEEI, and $p(\pi) \stackrel{\text{def}}{=} \sum_{l \in \pi} p_l$ for any π . Suppose $\vec{\pi}$ dominated by $\vec{\pi}'$. Then $u_i(\pi_i) \leq u_i(\pi'_i)$ for all i , with at least one strict inequality. Since $\vec{\pi}$ is optimal under budget \vec{p} , we have $u_i(\pi_i) < u_i(\pi'_i) \implies p(\pi_i) < p(\pi'_i)$. Preferences are strict, so $u_i(\pi_i) = u_i(\pi'_i) \implies \pi_i = \pi'_i \implies p(\pi_i) = p(\pi'_i)$. Therefore $\sum_{i \in \mathcal{A}} p(\pi_i) < \sum_{i \in \mathcal{A}} p(\pi'_i)$, which is impossible. Thus $\vec{\pi}$ is Pareto-efficient. \square

As a consequence of Propositions 7 and 8, when preferences are strict, a necessary condition for the existence of a CEEI is the existence of an envy-free Pareto-efficient allocation (which is known to be Σ_2^P -complete [9]). With this necessary condition, we can prove that the inclusion in Proposition 7 is strict, as the following example shows:

EXAMPLE 5. *Consider the 3 agents / 5 objects instance defined by the following weight matrix:*


$$W = \begin{pmatrix} 2 & 12 & 7* & \dagger 15 & * \dagger 11 \\ * \dagger 12 & 15 & \dagger 11 & *7 & 2 \\ 15 & * \dagger 20 & 9 & 2 & 1 \end{pmatrix}$$

It can be proved that the starred allocation is the only envy-free allocation. However, it is not Pareto-efficient, as it is dominated by the one marked with \dagger . Hence there is no Pareto-efficient envy-free allocation. The preferences being strict, Proposition 8 implies that there is no CEEI allocation.

Three open questions remain: determining whether the necessary condition of Propositions 7 and 8 is also sufficient, and finding the precise complexity of determining whether a given allocation is CEEI, and determining whether such an allocation exists for a given add-MARA instance.

3.6 A scale of criteria

Putting Propositions 2, 4, 6 and 7 together leads to the following implication sequence, for all $\vec{\pi} : (\vec{\pi} \models \text{CEEI}) \implies (\vec{\pi} \models \text{EF}) \implies (\vec{\pi} \models \text{mFS}) \implies (\vec{\pi} \models \text{PFS}) \implies (\vec{\pi} \models \text{MFS})$. In other words, these criteria can be ranked from the least to the most demanding as follows:



As the propositions also show, these results can also be interpreted the other way around, in terms of add-MARA instances: $\mathcal{I}_{|\text{CEEI}} \subset \mathcal{I}_{|\text{EF}} \subset \mathcal{I}_{|\text{mFS}} \subset \mathcal{I}_{|\text{PFS}} \subset \mathcal{I}_{|\text{MFS}} \subset \mathcal{I}$, all these inclusions being strict. The five criteria can thus be used to characterize the level of conflict one can expect

from a given add-MARA instance. In an instance for which it is proved to exist a CEEI, the level of conflict is very low, and thus it is possible to find an allocation which is satisfactory for everyone. However, an instance for which the best we can find is an allocation satisfying MFS is very prone to conflicts, and in that case, the benevolent arbitrator will have no choice but to leave some agents unsatisfied.

Beyond their differences, these criteria all have a common appealing feature: they do not rely on interpersonal comparison of utilities.⁵ It leads to the following (easy) result:

PROPOSITION 9. *The MFS, PFS, mFS, EF and CEEI criteria are preserved by any linear dilatation of individual utility scales.*

In other words, if $\langle \mathcal{A}, \mathcal{O}, w \rangle$ is an add-MARA instance and $\vec{\pi}$ an allocation satisfying criterion \mathcal{C} , then $\vec{\pi}$ also satisfies \mathcal{C} for any instance $\langle \mathcal{A}, \mathcal{O}, w_K \rangle$, where $K : \mathcal{A} \rightarrow \mathbb{R}^+$ and w_K is defined as follows: $w_K(i, l) = K(i) \times w(i, l)$.

Finally, max-min fair-share, proportional fair-share and min-max fair share have an interesting feature: if we can find an allocation satisfying one of these criteria \mathcal{C} , then we can find one that satisfies both \mathcal{C} and Pareto-efficiency. This is not the case for the envy-free criterion: as Example 5 shows, one can find instances having envy-free allocations, none of them being Pareto-efficient.

4. THE EGALITARIAN APPROACH

As pointed out at the beginning of the paper, an orthogonal approach for ensuring fairness in resource allocation problems is to choose a CUF and find an allocation that maximizes it. The most prominent one is the egalitarian CUF, which is defined in our context as the function $g_e : \vec{\pi} \mapsto \min_{i \in \mathcal{A}} u_i(\pi_i)$. This CUF is the formal translation of Rawlsian egalitarianism [21], which recommends to maximize the utility of the least well-off agent. Any allocation maximizing the egalitarian CUF will be called *min-optimal*.

Since egalitarianism and the approach based on the five aforementioned criteria are two different ways of defining fairness, a natural question is to investigate the links between them. Interestingly, it turns out that the compatibility between the two approaches depends on the criterion considered. Envy-freeness *e.g.* can be somewhat antagonistic with egalitarianism,⁶ as it has been pointed out by Brams and King [5], but as we will see egalitarianism is more compatible with proportional fair-share and max-min fair-share.

As egalitarianism requires interpersonal comparisons of utilities, we assume here normalized weights, namely: there is a constant K such that for all i , $\sum_{l \in \mathcal{O}} w(i, l) = K$.

PROPOSITION 10. *If there is an allocation satisfying the proportional fair-share criterion, then any min-optimal allocation satisfies it.*

PROOF. For all i , $u_i^{\text{PFS}} = K/n$. If there is $\vec{\pi}$ such that $\vec{\pi} \models \text{PFS}$, then $K/n \leq \min_{i \in \mathcal{A}} u_i(\pi_i)$. Let $\vec{\pi}^*$ be a min-optimal allocation. By definition $\min_i u_i(\pi_i) \leq \min_i u_i(\pi_i^*)$, hence $K/n \leq \min_{i \in \mathcal{A}} u_i(\pi_i^*)$ and $K/n \leq u_i(\pi_i^*)$, for all i . \square

This proposition also gives a practical way to find an allocation satisfying proportional fair-share if there is one, by normalizing weights and finding a min-optimal allocation.

⁵ Actually, four of them are even purely ordinal (PFS is not).

⁶ In the sense that an envy-free allocation can be far away from being min-optimal, and vice-versa.

Things are less clear for max-min fair-share. On the one hand, the latter result does not hold for max-min fair-share, as the following example shows.

EXAMPLE 6. *Consider the following instance:*

$$\begin{pmatrix} 58 & \dagger 15 & \dagger * 19 & 8 \\ \dagger 63 & * 5 & 25 & * 7 \\ 37 & 10 & * 27 & \dagger 26 \end{pmatrix} \begin{array}{l} \rightarrow * 19 / \dagger 34 \\ \rightarrow * 12 / \dagger 63 \\ \rightarrow * 27 / \dagger 26 \end{array}$$

The max-min fair-share of each agent (on the right) and the corresponding shares are starred. A min-optimal allocation and the corresponding utilities are marked with \dagger . The third agent does not get her max-min fair-share (expecting at least 27 but getting only 26).

On the other hand however, such a counter-example is very rare in practice: for example, using a uniform generation process similar to the impartial culture in vote theory, about 1 instance over 3500 is a counter-example similar to Example 6. This shows that the max-min fair-share has a good correlation with the egalitarian approach.

5. RESTRICTED CASES

We examine here some restrictions of the resource allocation problems, concerning the agents' preferences and the number of agents and objects. The main result here, is that for all these restrictions (even if some of them are very general), it is always possible to find an allocation satisfying max-min fair-share.

5.1 Preferences

0–1 preferences We first consider the case where the weights are binary, which corresponds to the MARA version of approval voting. Interestingly, we can prove that an MFS allocation can always be found, using a decentralized protocol where each agent takes in turn (according to a predefined sequence) one of its preferred (approved, here) objects among the remaining ones. Such a picking protocol is known as *product of sincere choices* [5] or *elicitation-free sequential protocol* [4]. Using this protocol with an alternating sequence of agents always yields an allocation satisfying max-min fair-share (if every agent acts sincerely):

PROPOSITION 11. *Any add-MARA instance with weights restricted to 0, 1 belongs to \mathcal{I}_{MFS} .*

Identical preferences When the agents have identical preferences, our scale of criteria has only two levels:

PROPOSITION 12. *For any add-MARA instance for which $w(j, l) = w(i, l)$ for all i, j, l :*

- (i) *any min-optimal allocation $\vec{\pi}$ satisfies MFS;*
- (ii) *if preferences are strict, no allocation satisfies PFS (and thus none satisfies the more demanding criteria);*
- (iii) *for all $\vec{\pi}$, the five following propositions are equivalent: (a) all the agents get the same utility in $\vec{\pi}$; (b) $\vec{\pi} \models \text{CEEI}$; (c) $\vec{\pi} \models \text{EF}$; (d) $\vec{\pi} \models \text{mFS}$; (e) $\vec{\pi} \models \text{PFS}$.*

PROOF. Point (i) is a direct implication of the definition of MFS. (ii) If preferences are strict, for any $\vec{\pi}$, the n numbers $u_i(\pi_i)$ are different. One of them at least is strictly smaller than their mean. (iii) Let $\vec{\pi}$ be an allocation giving the same utility to every agent. One can check that the price vector defined as $p_l = nw(i, l)/u_i(\mathcal{O})$ forms a CEEI with $\vec{\pi}$. So (a) \Rightarrow (b). Implications (b) \Rightarrow (c) \Rightarrow (d) \Rightarrow (e) follow from Section 3. (e) \Rightarrow (a) is easily proved. \square

Same-order preferences and scoring functions Intuitively, the more similar the agents preferences are, the more likely they are in conflict, and the harder it will be to satisfy the aforementioned fairness criteria. This notion of similarity is well captured by the concept of same-order preferences (SOP for short). Formally, an add-MARA instance satisfies SOP if for all $i, l, l' : l < l' \Rightarrow w(i, l) \geq w(i, l')$. In other words, all the agents agree on the same ranking of objects (1 is one of the best, m is one of the worst), but can give different weights.⁷ For any weight function w , we will write w^\uparrow the function $i, l \mapsto w(i, \sigma_i(l))$, where σ_i is a permutation of $\llbracket 1, m \rrbracket$ such that $l < l' \Rightarrow w(i, \sigma_i(l)) \geq w(i, \sigma_i(l'))$. Obviously, w^\uparrow is a ‘‘SOP’’ version of w . It turns out that if we can find a MFS allocation for a given SOP add-MARA instance, then we can find one for every permutation derived from it:

PROPOSITION 13. *Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. Then we have $\langle \mathcal{A}, \mathcal{O}, w^\uparrow \rangle \in \mathcal{I}_{\text{MFS}} \Rightarrow \langle \mathcal{A}, \mathcal{O}, w \rangle \in \mathcal{I}_{\text{MFS}}$.*

PROOF. We will here once again use the aforementioned idea of sequence of sincere choices. Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance, and let $\vec{\pi}^\uparrow$ be an allocation satisfying MFS for the SOP instance $\langle \mathcal{A}, \mathcal{O}, w^\uparrow \rangle$. Let $S = S_1, S_2, \dots, S_m$ be the sequence of agents defined as follows: S_i is the agent who receives object l in $\vec{\pi}^\uparrow$. Because agents have the same preference order, $\vec{\pi}^\uparrow$ is (one of) the allocation(s) obtained by the sequence of sincere choices S_1, S_2, \dots, S_m .

The key is to notice that the allocation $\vec{\pi}$ obtained by the same picking protocol (with the same sequence) used with the original instance $\langle \mathcal{A}, \mathcal{O}, w \rangle$, will make every agent at least as well-off as in $\vec{\pi}^\uparrow$. To see it, notice that before step p , exactly $p - 1$ objects have been chosen, so the worst object that agent S_p could have at step p is the object p obtained in $\vec{\pi}^\uparrow$. Consequently, for each agent i and each object of π_i^\uparrow , there is an object in π_i which is weakly better for i : the utility of i weakly increases from $\vec{\pi}^\uparrow$ to $\vec{\pi}$.

Since the MFS of an agent only depends on the set of weights (not on their ordering), it is the same for the SOP instance and the original one. Since $\vec{\pi}^\uparrow \models \text{MFS}$, and $\vec{\pi}$ makes everyone at least as well-off, we conclude $\vec{\pi} \models \text{MFS}$. \square

Because every add-MARA instance can be considered as a derivation (by permutations of weights) of a SOP one, this proposition shows that SOP instances are the most difficult as far as the MFS criterion is concerned.⁸

Interestingly, we can use this result to show that \mathcal{I}_{MFS} contains another huge family of MARA instances, namely, the ones where all the agents have the same multiset of weights: $\{\{w(i, l) \mid l \in \mathcal{O}\} \mid i \in \mathcal{A}\} = \{\{w(j, l) \mid l \in \mathcal{O}\} \mid j \in \mathcal{A}\}$ for all i, j . Equivalently we could say that agents use the same *scoring function*. A scoring function is a weakly decreasing function $g : \llbracket 1, m \rrbracket \rightarrow \mathbb{R}^+$. It can be used to convert a purely ordinal expression of preferences into to a numerical one, as follows. Consider that each agent ranks strictly the objects from 1 (the most preferred) to m (the least preferred). If $r(i, l)$ is the rank given to object l by agent i , then the weight $w(i, l)$ is defined as $g(r(i, l))$. Using a scoring function to ‘‘cardinalize’’ ordinal preferences is standard in social choice, especially in voting theory, where it is at the basis of well-known scoring procedures (plurality, veto, Borda among others).

⁷This property is sometimes known as full-correlation [4].

⁸This also seems to be true for more demanding criteria as our experiments show in Section 6.

PROPOSITION 14. *Any add-MARA instance in which preferences are defined by the same scoring function is in \mathcal{I}_{MFS} , and any min-optimal allocation satisfies MFS in this case.*

PROOF. By Proposition 13 it is enough to consider SOP instances, which are in this case instances with identical preferences. Then use Proposition 12. Remark: the other conclusions of Proposition 12 are also satisfied. \square

5.2 Number of agents and objects

Two agents The 2-agents case is interesting because the famous cut-and-choose game gives the max-min fair-share to both agents.

PROPOSITION 15. *Any 2-agents add-MARA instance belongs to \mathcal{I}_{MFS} .*

PROOF. Agent 1 cuts. Agent 2 chooses first, getting her mFS, therefore her MFS too. \square

Restricted number of objects When $m \leq n$, the scale of criteria somewhat collapses. The case $m = n$ brings to light matchings (allocations giving one object to each agent) and min-max fair-share.

PROPOSITION 16. *Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. If $m < n$, then every allocation satisfies MFS, but none satisfies the other criteria. If $m = n$, then (i) any matching $\vec{\pi}$ satisfies MFS; (ii) if $\vec{\pi}$ satisfies mFS, then $\vec{\pi}$ is a matching, is Pareto-efficient, envy-free, and satisfies CEEI.*

We can go a little bit further for max-min fair-share by proving that any add-MARA instance with up to three more objects than agents belongs to \mathcal{I}_{MFS} . We first introduce some useful preliminary lemmas showing how u^{MFS} evolves if we add some agents and objects to an add-MARA instance. Let us first define the extension of a MARA instance:

DEFINITION 6. *Let $I = \langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. A (p, q) -extension of I is an add-MARA instance $I_{+p, +q} = \langle \mathcal{A}', \mathcal{O}', w' \rangle$ such that $\mathcal{A}' = \mathcal{A} \cup \{n + 1, \dots, n + p\}$, $\mathcal{O}' = \mathcal{O} \cup \{m + 1, \dots, m + q\}$, and $w'(i, l) = w(i, l)$ for all $(i, l) \in \mathcal{A} \times \mathcal{O}$.*

PROPOSITION 17. *Any add-MARA instance with n agents and $(n + 1)$ objects belongs to \mathcal{I}_{MFS} .*

PROOF. Following Proposition 13, we can restrict to instances satisfying SOP. Since objects n and $n + 1$ are the worst ones, it can be seen that all the shares from allocation $\langle (1)(2) \cdots (n - 1)(n, n + 1) \rangle$ give to each agent her MFS. \square

The cases $m = n + 2$ and $m = n + 3$ are a bit trickier and require additional lemmas.

LEMMA 1. *Let $I = \langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. Then for all $i \in \mathcal{A}$, u_i^{MFS} does not increase from I to any $(1, 1)$ -extension of I .*

PROOF. Let $u_i^{\text{MFS}}(I)$ denote the MFS of agent i in instance I . Start from an allocation $\vec{\pi}'$ of $I_{+1, +1}$ such that $\min_{j=1}^{n+1}(u_i(\pi'_j)) = u_i^{\text{MFS}}(I_{+1, +1})$. Removing from $\vec{\pi}'$ the share containing object $m + 1$ yields a valid (possibly incomplete) allocation $\vec{\pi}$ for I . Hence, $u_i^{\text{MFS}}(I) \geq \min_{j=1}^n(u_i(\pi_j)) = \min_{j=1}^n(u_i(\pi'_j)) \geq \min_{j=1}^{n+1}(u_i(\pi'_j)) = u_i^{\text{MFS}}(I_{+1, +1})$. \square

The next lemma gives an upper bound of u^{MFS} .

LEMMA 2. Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. For all i , we have $u_i^{\text{MFS}} \leq \lfloor \frac{m}{n} \rfloor \max_{l=1}^m w(i, l)$.

The proof is not difficult. We can now use these two lemmas to show the following more general result:

LEMMA 3. Let $I = \langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. If $I \in \mathcal{I}_{\text{MFS}}$ and $n \leq m \leq 2n$ then any (p, p) -extension of I is in \mathcal{I}_{MFS} .

PROOF SKETCH. First prove the lemma for $p = 1$, which can be done, considering (after Proposition 13) only SOP instances, together with Lemmas 1 and 2. An induction argument gives the result for general p . \square

Proposition 15 (showing that any instance with 2 agents and 4 objects belongs to \mathcal{I}_{MFS}) and Lemma 3 directly imply:

PROPOSITION 18. Any add-MARA instance with n agents and $(n + 2)$ objects belongs to \mathcal{I}_{MFS} .

The case with n and $n+3$ objects can also be proved using Lemma 3, but for that we need to prove the base case with 3 agents and 6 objects.

LEMMA 4. Any add-MARA instance with 3 agents and 6 objects belongs to \mathcal{I}_{MFS} .

PROOF. Once again, we consider a SOP instance I with 3 agents and 6 objects and investigate two cases.

(i) Suppose that there is an agent (say 3) s. t. $w(3, 1) \geq u_3^{\text{MFS}}$. Let I' be the restriction of I where object 1 and agent 3 have been removed. By Proposition 15, there is a $\vec{\pi}'$ such that $\vec{\pi}' \models \text{MFS}$ in I' . By Lemma 1 together with the fact that $w(3, 1) \geq u_3^{\text{MFS}}$, $\langle \pi'_1, \pi'_2, \{1\} \rangle$ satisfies MFS in I .

(ii) Otherwise, $w(i, l) < u_i^{\text{MFS}}$ for all i, l . Hence for all i every max-min cut has 3 shares of 2 objects each, which is $\langle (1, 6)(2, 5)(3, 4) \rangle$, the same for every agent. \square

PROPOSITION 19. Any add-MARA instance with n agents and $(n + 3)$ objects belongs to \mathcal{I}_{MFS} .

6. EXPERIMENTS

For each combination of n agents and m objects we considered 1000 randomly generated instances, and their SOP versions, with weights uniformly drawn from $[0, 1]$. Table 6 shows for each criterion \mathcal{C} the number of instances (out of 1000) belonging to $\mathcal{I}_{\mathcal{C}}$. Results about MFS are not shown in the table, as all generated instances were in \mathcal{I}_{MFS} . As exactly characterizing a CEEI allocation is computationally too difficult in general [18, Section 3], we used envy-freeness and Pareto-efficiency (EFP) as a proxy for this criterion.⁹

Several facts can be noticed, which confirm our theoretical results. (1) Main result: the scale of criteria is really significant. The numbers weakly decrease from left to right, and often strictly decrease, showing that the scale is not trivial. (2) SOP instances are more conflicting than non SOP ones, in accordance with Proposition 13. (3) For a fixed number of agents, instances are less conflict-prone as the number of objects increases: intuitively, we get closer to the continuous case. (4) As said before, all generated instances are in

⁹As we saw in Section 3.5, EFP implies CEEI when preferences are strict, but we believe that they are not equivalent in the context of this discrete model.

n, m	Non SOP instances				SOP instances			
	PFS	mFS	EF	EFP	PFS	mFS	EF	EFP
3, 3	618	231	231	231	0	0	0	0
3, 4	821	563	318	318	340	2	2	2
3, 5	829	730	530	477	652	237	218	218
3, 6	991	967	933	890	775	500	374	374
3, 7	1000	999	997	989	942	780	615	611
3, 8	1000	999	997	995	990	958	869	831
3, 9	1000	1000	1000	1000	1000	995	983	965
3, 10	1000	1000	1000	1000	1000	1000	1000	990
3, 11	1000	1000	1000	1000	1000	1000	1000	999
4, 4	746	86	86	86	0	0	0	0
4, 5	945	511	130	130	159	0	0	0
4, 6	927	744	217	192	563	2	1	1
4, 7	920	843	530	434	809	131	86	86
4, 8	998	998	978	923	868	500	241	240
4, 9	1000	1000	998	984	972	751	442	433
4, 10	1000	1000	1000	999	1000	952	752	706
4, 11	1000	1000	1000	1000	1000	999	962	912
5, 5	839	43	43	43	0	0	0	0
5, 6	991	376	38	38	62	0	0	0
5, 7	989	726	73	61	430	0	0	0
5, 8	970	835	178	130	764	0	0	0
5, 9	964	903	561	387	896	70	29	29
5, 10	1000	997	985	953	941	449	142	138
5, 11	1000	1000	1000	998	987	732	302	286

Table 1: Experimental results.

\mathcal{I}_{MFS} , showing that it is quite unlikely to find an instance not in \mathcal{I}_{MFS} even if such instances exist [20].

Other experiments have been conducted with different distributions of weights (not reported here for lack of space). Notably, with a Gaussian distribution with small variance, instances where m is close to a multiple of n are less conflict-prone than others, which is not very surprising.

7. BEYOND ADDITIVE PREFERENCES

Even if, as we have seen earlier, it is almost always possible, for a given add-MARA instance, to find an MFS allocation, things are surprisingly different for more general non-additive preferences. The most natural way of relaxing preference additivity while keeping some conciseness is to allow limited synergies (complementarities or substitutabilities) between objects, which is the exact idea behind k -additive functions [11, 8].

Formally, we consider in this section k -additive multi-agent resource allocation instances (k -add-MARA instances for short), defined as triples $\langle \mathcal{A}, \mathcal{O}, w \rangle$, where w is now a mapping from $\mathcal{A} \times 2^{\mathcal{O}}$ to \mathbb{R} such that $w(i, \pi) = 0$ for all agent i and subset π such that $|\pi| > k$. In other words, w gives a weight for all agent and all subset of less than k objects. The utility function is, as before, defined additively: $u_i(\pi) = \sum_{\pi' \subseteq \pi} w(i, \pi')$. Obviously, 1-additive functions are the additive functions, so the 1-add-MARA instances are exactly the add-MARA instances considered earlier.

As soon as we switch from 1-additive to 2-additive functions, finding an instance for which no MFS allocation exists is not challenging anymore:

EXAMPLE 7. Let us consider the 2 agents / 4 objects instance defined by the following weight functions:

- $w(1, \{1, 2\}) = w(1, \{3, 4\}) = 1$
- $w(2, \{1, 3\}) = w(2, \{2, 4\}) = 1$
- $w(i, \pi) = 0$ for every other share π .

It is not hard to see that $u_i^{\text{MFS}} = 1$ for both agents, and no allocation giving at least 1 to both agents exist.

Actually, the problem of determining whether there exists

an allocation satisfying max-min fair share (further referred to as $[k\text{-ADD-MFS-EXIST}]$) is even hard:

PROPOSITION 20. $[k\text{-ADD-MFS-EXIST}]$ is NP-hard, for $k \geq 2$ and $n \geq 3$.

PROOF SKETCH. NP-hardness can be proved by reduction from the [PARTITION] problem. Let $\{s_1, \dots, s_n\}$ be an instance of this problem. From this instance, we create a 3-agents / $n + 4$ objects k -add-MARA instance, where the agents' preferences are defined as follows:

- for all i , $w(i, \{l\}) = s_l$ and $w(i, \{l, n + m\}) = -3L$ for all $l \in \llbracket 1, n \rrbracket$ and $m \in \llbracket 1, 4 \rrbracket$;
- $w(1, \{n + 1, n + 2\}) = w(1, \{n + 3, n + 4\}) = L$
- $w(2, \{n + 1, n + 3\}) = w(2, \{n + 2, n + 4\}) = L$
- $w(3, \{n + 1, n + 4\}) = w(3, \{n + 2, n + 3\}) = L$
- $w(i, \pi) = 0$ for every other share π .

The key of the proof is based on the idea of Example 7: obviously each agent can individually enjoy L with 2 different shares of $\{n + 1, \dots, n + 4\}$, but no allocation of these 4 objects can give L to two different agents. The only way to give at least L to all the agents is to partition the first n objects into 2 shares, and give the rest to the third agent. \square

It can be noticed that Proposition 20 only gives a NP-hardness result, as it is not known yet whether $[k\text{-ADD-MFS-EXIST}]$ belongs to NP. We only know for now that this problem belongs to Σ_2^P , because it can be solved by the same algorithm as in the additive case (see Section 3.1).

8. CONCLUSION AND FUTURE WORK

In this paper we have introduced five fairness criteria for resource allocation, two of which being classical, two of which being less well-known, and one being original. We have shown how these criteria form, in the context of multiagent resource allocation with additive preferences, an ordered scale that can be used as a basis not only for finding satisfactory (fair) allocations, but also for measuring to which extent it is possible to find some. We have also run some experiments that give some insights on how instances divide up on this scale of properties, and finally we have shown that the extension of these criteria to more general preferences is likely to have quite different properties.

This work raises many interesting questions, beyond the several open (complexity) problems presented in the paper. Among others, the question of efficiently computing allocations satisfying some criteria is crucial and not trivial, especially for CEEI (where no efficient complete algorithm is known so far [18]). From a more theoretical point of view, the question of extending the results to non-additive problems is worth being further investigated. Lastly, since four of the five criteria introduced are purely ordinal (PFS is not), it would be interesting to analyze to which extent our results carry over to an ordinal setting with separable preferences: unlike numerical additivity, ordinal separability leaves many pairs of allocations incomparable. Hence, even if the criteria themselves can be directly expressed ordinally, the way they must be adapted to deal with incomparable pairs is not so clear and deserves further investigation.

Acknowledgements We thank Ariel Procaccia and members of COST Action IC-1205 for helpful discussions.

9. REFERENCES

- [1] A. Asadpour and A. Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. In *Proc. of STOC-2007*, 2007.
- [2] N. Bansal and M. Sviridenko. The Santa Claus problem. In *Proc of STOC'06*, 2006.
- [3] C. Beviá. Fair allocation in a general model with indivisible goods. *Review of Economic Design*, 3:195–213, 1998.
- [4] S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proc. of IJCAI'11*, 2011.
- [5] S. J. Brams and D. King. Efficient fair division – help the worst off or avoid envy? *Rationality and Society*, 17(4), 2005.
- [6] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6), 2011.
- [7] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.
- [8] Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Multiagent resource allocation with k -additive utility functions. In *Proc. DIMACS-LAMSADE Workshop on Computer Science and Decision Theory*, 2004.
- [9] B. de Keijzer, S. Bouveret, T. Klos, and Y. Zhang. On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In *Proc. of ADT'09*. Springer Verlag, 2009.
- [10] D. Foley. Resource allocation and the public sector. *Yale Econ. Essays*, 7(1), 1967.
- [11] M. Grabisch. k -order additive discrete fuzzy measure and their representation. *Fuzzy Sets and Systems*, 92, 1997.
- [12] T. P. Hill. Partitioning general probability measures. *The Annals of Probability*, 15(2):804–813, 1987.
- [13] M. Lemaître, G. Verfaillie, and N. Bataille. Exploiting a common property resource under a fairness constraint: a case study. In *Proc. of IJCAI'99*, 1999.
- [14] R. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of EC'04*, 2004.
- [15] E. Markakis and C.-A. Psomas. On worst-case allocations in the presence of indivisible goods. In *Proc. of WINE 2011*, Singapore, December 2011.
- [16] J. S. Mill. *Utilitarianism*. U. of Chicago Press, 1906.
- [17] H. Moulin. *Cooperative Microeconomics, A Game-Theoretic Introduction*. Prentice Hall, 1995.
- [18] A. Othman, T. Sandholm, and E. Budish. Finding approximate competitive equilibria: efficient and fair course allocation. In *Proceedings of AAMAS'10*, 2010.
- [19] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [20] A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. Manuscript, February 2014.
- [21] J. Rawls. *A Theory of Justice*. Harvard University Press, 1971.
- [22] H. Steinhaus. The problem of fair division. *Econometrica*, 16(1), 1948.