Characterizing Conflicts in Fair Division of Indivisible Goods Using a Scale of Criteria[1]

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

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

May 2, 2014

[1]This article is the long version of an article published in the proceedings of AAMAS'14, with the same title and authors. This long version includes detailed proofs and experiments that do not appear in the AAMAS version.

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: for a given instance of a resource allocation problem, the less conflicting the agents' preferences are, the more demanding criterion this instance is able to satisfy, and the more satisfactory the allocation can 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.

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

1 Introduction

In the context of economically motivated agents, the fair allocation of resources is an important and frequent problem. A subcase of this problem, subject of this article, namely the allocation of indivisible goods (or objects) to a set of agents, arises in a wide range of real world applications, including auctions, divorce setlements, wave frequency allocation, airport traffic management, spatial resource allocation [21], fair scheduling, allocation of tasks to workers, articles to reviewers, courses to students [26].

More precisely, we study here a simple model of fair division of indivisible goods based on the following assumptions.

- A set of indivisible goods which will be called objects must be distributed among a set of agents.
- Agents have numerical additive preferences over the objects (except in Section 7 where we consider a more general model).
- 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.
- No monetary transfer is possible between agents.

Even if this model seems restrictive (especially the second assumption), it has been largely investigated (see for example [3, 16, 11, 19, 9, 22, 6, 2, 1, 7, 5, 15, 23]) because it offers a natural compromise between simplicity and expressiveness.

Our contribution consists mainly in providing a logical scale characterizing the degree of conflict inherent to each problem instance.

An important point in this context is how agents express their preferences. In a centralized allocation process, the agents have first to communicate and hence explicitly describe their preferences over the objects. Two main approaches are appropriate for that. The first rests on a purely ordinal expression of preferences, such as a weak (partial or total) order. The second one exploits a numerical expression of preferences taking the form of utility functions. This article, for convenience, rests on the second approach, but many of the results presented here could be transposed in the first, purely ordinal one.

Another crucial point in fair allocation mechanisms is the following : how to define fairness and how can it be evaluated ? Once again two main options are available. The first one consists in defining a collective utility function (CUF) aggregating individual agents' utilities. If the CUF is well chosen, its outcome when applied to individual utilities reflects the fairness (and possibly other desirable properties) of a given allocation. The arbitrator just looks for an allocation maximizing this CUF. The second option consists in defining, by mean of a boolean (logical) criterion, what is considered as fair. This is the approach followed by Lipton *et al.* [22] among others for envy-freeness. This article explores mainly this logical 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]

Our contribution. Some instances of fair sharing problems are more conflicting than others. When objects are numerous and agents prefer somewhat different objects, a well-balanced allocation, satisfying all participants, is likely to be found. On the opposite, when agents have similar preferences (they want more or less the same objects with the same intensity), or when there are only a few objects to distribute, the sharing out will be for sure conflicting.

Our main and original contribution is the following. Starting from the simple model of fair division sketched above, we show that the five criteria cited above form a linear scale of increasing requirements, that can be used to characterize formally the level of fairness of a given allocation. The more demanding criterion is satisfied by an allocation, the more this allocation is fair, harmonious and not conflicting.

This scale of properties can be used to characterize not only an allocation, as said before, but also a resource allocation problem instance: the degree of non conflictness of an instance is measured by the most demanding criterion an allocation from this instance can satisfy.

This article is structured as follows. Section 2 describes the model : fair division of indivisible objects under numerical additive preferences. The scale of five properties characterizing the fairness of an allocation, as well as associated computational complexity results are exposed in Section 3. We go back to the collective utility function approach in Section 4 to connect the important egalitarian CUF to the scale of criteria. Section 5 is devoted to a bunch of interesting restricted cases. Some experimental results on the scale of criteria are presented in Section 6. Extending the model to k-additive preferences, Section 7 presents a quite different perspective.

2 Model

Let $\mathcal{A} = \{1, \ldots, n\}$ be a set of *agents* and $\mathcal{O} = \{1, \ldots, m\}$ be a set of indivisible *objects*. An *allocation* of the objects to the agents is a vector $\overrightarrow{\pi} = \langle \pi_1, \ldots, \pi_n \rangle$, where $\pi_i \subseteq \mathcal{O}$ is the *bundle* of objects allocated to agent i, called agent i's *share*. An allocation $\overrightarrow{\pi}$ is said to be *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). We will

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

write $\mathcal{F}_{n,m}$ the set of admissible allocations for a given set of n agents and m objects (n and m will be omitted when the context is clear). All the allocations considered in this paper are implicitly admissible.

To find a "good" allocation, it is necessary to know the agents' preferences over the sets of objects they may receive. We make two usual assumptions concerning the way agents express their preferences. First, we consider that they are expressed numerically by a *utility function* $u_i : 2^{\mathcal{O}} \to \mathbb{R}^+$ specifying, for each agent i, the satisfaction $u_i(\pi)$ she enjoys if she receives bundle π: this is the utilitarian model [24]. 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:

$$u_i(\pi) \stackrel{def}{=} \sum_{l \in \pi} w(i,l), \tag{1}$$

where $w(i,l)$ is the *weight* given by agent i to object l. This assumption, as restrictive it may seem to be, is made by a lot of authors [2, 22, *e.g*] and offers a good compromise between preference expressive power and conciseness.

Adapting the terminology from the survey by Chevaleyre *et al.* [13], we 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 $w : \mathcal{A} \times \mathcal{O} \to \mathbb{R}^+$ is a function specifying the weight $w(i,l)$ given by agent i to object l.

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

The basic notions of computational complexity [27] are supposed to be well-known by the reader: P and NP refer to the two usual complexity classes; $\mathsf{\Sigma_2^P}$ 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

Even before any fairness consideration, the most basic desirable criterion for a resource allocation is Pareto-efficiency, of which the definition is recalled here:

Definition 1. *Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. We say that allocation $\overrightarrow{\pi}$* dominates *allocation $\overrightarrow{\pi}'$ if and only if $u_i(\pi_i) \geq u_i(\pi'_i)$ for all i, with at least one strict inequality. A* Pareto-efficient allocation *is an undominated allocation.*

The Pareto-efficiency criterion encodes the idea that the resource to be shared should not be wasted or under-exploited, but tells nothing about fairness. Two approaches are possible to deal with the fairness requirement.

1. If the preferences are numerical, we can use a *collective utility function* (CUF) to aggregate the individual preferences into a collective preference,

and look for an allocation that maximizes this function. If this function is carefully chosen, it can encode some idea of fairness (like the egalitarian criterion min for example, discussed in Section 4).

2. We can chose a fairness criterion and look for an allocation that obeys it, if some exists. The two prominent fairness criteria are envy-freeness [17] and proportional fair-share [31].

In this paper, we adopt the second point of view. We will introduce five fairness criteria (including envy-freeness and proportional fair-share) and show how they form together a scale of criteria of increasing strength. This scale provides an evaluation of the degree of fairness of a given allocation on the one hand, and can give an idea of the degree of "conflictuality" of a given add-MARA instance. For each one of these criteria, we write $\overrightarrow{\pi} \models \mathcal{C}$ if the allocation $\overrightarrow{\pi}$ satisfies criterion $\mathcal{C}$; $\mathcal{I}_{|\mathcal{C}}$ denotes the set of add-MARA instances admitting at least one allocation satisfying criterion $\mathcal{C}$.

3.1 Max-min fair-share

One of the most prominent fairness criteria in resource allocation problems is proportional fair-share, that will be discussed in details in Section 3.2. This criterion, coined by Steinhaus [31] in the context of continuous fair division (cake-cutting) problems, states that each agent should get from the allocation at least the n^{th} 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 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 [12], which defines the *max-min fair share*, whose original definition is purely ordinal, but which can be defined in our (utilitarian) setting as follows:

Definition 2. *Let* $(\mathcal{A}, \mathcal{O}, w)$ *be an add-MARA instance. The* max-min share *of agent* i *for this instance is*

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

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

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

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

We have $u_1^{\text{MFS}} = 12$ *(with share* $\{2,4\}$*) and* $u_2^{\text{MFS}} = 11$ *(with share* $\{1,2\}$*). The allocation* $\overrightarrow{\pi} = \langle\{1,4\},\{2,3\}\rangle$ *marked with stars satisfies max-min fair share, with* $u_1(\pi_1) = 17 > 12$ *and* $u_2(\pi_2) = 14 > 11$.

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 can be considered as 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; hence there is no reason to give me less. It is also the maximum utility that an agent can get for sure in the allocation game "I cut, I choose last": the agent proposes her best allocation (that will be refered 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 [20], recently refined by [23], which establishes a worst case garantee on the utility an agent can have. However, this garantee 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 $\langle A, \mathcal{O}, w \rangle$, an agent i, an integer K.
Question:	Do we have $u_i^{\text{MFS}} \geq K$?

Proposition 1. [MFS-Comp] *is NP-complete, for all $n \geq 2$.*

Proof. Membership to NP is obvious. NP-hardness can be proved by reduction from the partition problem:

Problem 2 [Partition]	
Input:	A set $\mathcal{X} = \{x_1, \dots, x_n\}$ and a mapping $s : \mathcal{X} \to \mathbb{N}$ such that $\sum_{x_i \in \mathcal{X}} s(x_i) = 2L$.
Question:	Is there a partition $(\mathcal{X}_1, \mathcal{X}_2)$ of $\mathcal{X}$ such that $\sum_{x_i \in \mathcal{X}_1} s(x_i) = \sum_{x_i \in \mathcal{X}_2} s(x_i) = L$?

From a given instance of [Partition], we can create an instance of [MFS-Comp] with two agents and n objects $\{1, \dots, n\}$. The agents' preferences are identical and defined as $w(1, l) = w(2, l) = s(x_l)$. Integer K is defined as L, which completes the reduction.[2] □

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),

[2]We use here a very similar idea to the one used by [22, p4].

and no counterexample was found in thousands of randomly generated instances (see Section 6). However, surprisingly, Procaccia and Wang [28] have recently proved (by a very tricky construction) that there actually exists add-MARA instances for which there is no allocation satisfying max-min fair-share. Put in other words, we thus have $\mathcal{I}_{|\mathrm{MFS}} \subsetneq \mathcal{I}$.

3.2 Proportional fair-share

The aforementioned concept of proportional fair-share was originally defined not on the utilities but on the resources themselves [31]. A lot of authors have since given a natural utilitarian interpretation of this notion, like the one that follows:

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

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

We say that the allocation $\overrightarrow{\pi}$ satisfies the criterion of proportional fair-share*, written $\overrightarrow{\pi} \models \mathrm{PFS}$, if $u_i^{\mathrm{PFS}} \leq u_i(\pi_i)$ for all i (that is, each agent obtains at least her proportional fair-share in $\overrightarrow{\pi}$).*

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 criterion is more demanding than max-min fair-share:

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

Proof. Let $\overrightarrow{\pi}$ be an allocation and i an agent. We have $\sum_{j \in \mathcal{A}} u_i(\pi_j) = u_i(\mathcal{O})$. The minimum of a set of numbers being lower than their mean, we have

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

Taking the max over the set of allocations on the two sides of the latter inequality directly leads to the result: $u_i^{\mathrm{MFS}} \leq u_i^{\mathrm{PFS}}$. □

The inclusion in Proposition 2 is strict: we can consider a instance with two agents and one object, for which 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, determining whether a given add-MARA instance has an allocation satisfying proportional fair-share (problem that we shall call [PFS-EXIST]) is computationally hard:

Proposition 3. [PFS-EXIST] *is* NP*-complete, for all* $n \geq 2$.

This proposition can be proved using a similar reduction as the one used in proof of Proposition 1.

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 defined earlier.

Definition 4. *Let* $(\mathcal{A}, \mathcal{O}, w)$ *be an add-MARA instance. The* max-min share *of agent* i *for this instance is*

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

We say that the allocation $\overrightarrow{\pi}$ *satisfies the criterion of* max-min fair-share*, written* $\overrightarrow{\pi} \models \text{mFS}$*, 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* $\overrightarrow{\pi}$*).*

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 can be proved in a similar way:

Proposition 4. *Let* $\langle \mathcal{A}, \mathcal{O}, w \rangle$ *be an add-MARA instance. We have* $u_i^{\text{PFS}} \leq u_i^{\text{mFS}}$*, for all* $i \in \mathcal{A}$*. Hence, for all* $\overrightarrow{\pi}$*, we have* $\overrightarrow{\pi} \models \text{mFS} \implies \overrightarrow{\pi} \models \text{PFS}$ *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. Hence the allocation marked with stars 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, if [mFS-Comp] is the equivalent for min-max fair-share of decision Problem 1, the following proposition holds.

Proposition 5. [mFS-Comp] *is* coNP*-complete, for all* $n \geq 2$.

The decision problem becomes coNP-complete because the 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, and is thus omitted.

Of course, an add-MARA instance may not always have an allocation satisfying min-max fair-share. The decision problem of determining whether there exists one is very likely to be hard, but its precise complexity remains unknown.[3]

3.4 Envy-freeness

The envy-freeness criterion [17] is probably the most prominent one.

Definition 5. *Let* $\langle \mathcal{A}, \mathcal{O}, w \rangle$ *be an add-MARA instance. The allocation* $\overrightarrow{\pi}$ *satisfies the criterion of* envy-freeness *(or simply is envy-free), written* $\overrightarrow{\pi} \models \mathrm{EF}$, *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).*

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* $\overrightarrow{\pi}$: $\overrightarrow{\pi} \models \mathrm{EF} \implies \overrightarrow{\pi} \models \mathrm{mFS}$, *and hence* $\mathcal{I}_{|\mathrm{EF}} \subset \mathcal{I}_{|\mathrm{mFS}}$.

Proof. In every envy-free allocation, each agent obtains a share which is of maximal utility for her in this allocation. Hence, such a share has a greater utility than her min-max fair-share. More formally : let $\overrightarrow{\pi}$ 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 $\overrightarrow{\pi} \in \mathcal{F}$, $u_i(\pi_i) \geq \min_{\overrightarrow{\pi} \in \mathcal{F}} \max_{j \in \mathcal{A}} u_i(\pi_j) = u_i^{\mathrm{mFS}}$. □

Once again, the inclusion introduced in this proposition is strict, as shows the following example.

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

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

[3] All that we can say for sure is that this problem belongs to Σ_2^{P}.

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

3.5 Competitive Equilibrium from Equal Incomes

The last introduced criterion is a classical notion in microeconomics [25, for example]. It has, to the best of our knowledge, almost never been considered in computer science, with the notable exception of the work of Othman *et al.* [26] about course allocation. 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. A lot of variants of this notion exist; the following definition is adapted from Budish [12].

Definition 6. *Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ an add-MARA instance, $\overrightarrow{\pi}$ an allocation, and $\overrightarrow{p} \in [0,1]^m$ a price vector. A pair $(\overrightarrow{\pi}, \overrightarrow{p})$ is said to form a* competitive equilibrium from equal incomes *(CEEI), if for each agent i,*

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

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 $\overrightarrow{\pi}$ satisfies the CEEI criterion, written $\overrightarrow{\pi} \models \mathrm{CEEI}$, if there exists a price vector $\overrightarrow{p}$ such that $(\overrightarrow{\pi}, \overrightarrow{p})$ 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 marked allocation, associated to price vector $(0.8, 0.2, 0.8, 0.2)$ forms a CEEI.

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

Proposition 7. *Every CEEI allocation is envy-free. That is, for all $\overrightarrow{\pi} : \overrightarrow{\pi} \models \mathrm{CEEI} \implies \overrightarrow{\pi} \models \mathrm{EF}$, and hence $\mathcal{I}_{|\mathrm{CEEI}} \subset \mathcal{I}_{|\mathrm{EF}}$.*

Proof. Let $\overrightarrow{\pi}$ be a CEEI allocation, and suppose that $u_i(\pi_j) > u_i(\pi_i)$ (agent i envies j). Since budgets and prices are the same for everyone, π_i is not the maximal utility share which can be bought by agent i, which contradicts the definition of the CEEI. Thus $\overrightarrow{\pi}$ is envy-free. □

The CEEI also has the following interesting property:

Proposition 8. *When the agents' preferences are strict (meaning that distinct shares have distinct utilities), any CEEI allocation is Pareto-efficient.*

Proof. Let $(\overrightarrow{\pi}, \overrightarrow{p})$ be a CEEI. For a share π, we write $p(\pi) \overset{def}{=} \sum_{l \in \pi} p_l$. Suppose that $\overrightarrow{\pi}$ is not Pareto-efficient. Then there is a $\overrightarrow{\pi}'$ such that $u_i(\pi_i) \leq u_i(\pi'_i)$ for all i, with at least one strict inequality. Since $\overrightarrow{\pi}$ is optimal under budget $\overrightarrow{p}$, we have $u_i(\pi_i) < u_i(\pi'_i) \Rightarrow p(\pi_i) < p(\pi'_i)$. But $u_i(\pi_i) = u_i(\pi'_i) \Rightarrow \pi_i = \pi'_i \Rightarrow p(\pi_i) = p(\pi'_i)$ because preferences are strict. Therefore $\sum_{i \in \mathcal{A}} \overrightarrow{p}(\pi_i) < \sum_{i \in \mathcal{A}} \overrightarrow{p}(\pi'_i)$, which is impossible. Thus $\overrightarrow{\pi}$ is Pareto-efficient. □

The following example shows that the strict preference hypothesis, in the previous proposition, is necessary.

Example 5.

$$W = \begin{pmatrix} *2 & 3 & 3 & *2 \\ 2 & 3 & *4 & 1 \\ 0 & *4 & 2 & 4 \end{pmatrix} \begin{matrix} \to 4 \\ \to 4 \\ \to 4 \end{matrix}$$

In this instance, preferences are not strict. The marked allocation, associated to price vector $\langle 0.5, 1, 1, 0.5 \rangle$ *forms a CEEI. However, it is dominated by the allocation* $\langle (1,2), (3), (4) \rangle$ *which provides utilities* $(5, 4, 4)$*. The marked allocation is CEEI but not Pareto-efficient.*

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 [15]). With this necessary condition, we can prove that the inclusion in Proposition 7 is strict, as the following example shows:

Example 6. *Let us consider the 3 agents / 5 objects instance in which preferences are strict, 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 allocation marked with $*$ is the only envy-free allocation. However, this allocation 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 : do we have, under the hypothesis of strict preferences, EF plus Pareto-efficiency implies CEEI ? This result would give, in this discrete model, under the strict preference hypothesis, the equivalence between CEEI and EF plus Pareto-efficiency. A counter-example of this result would be to find an instance with strict preferences, having an EF and Pareto-efficient allocation but not CEEI
- finding the precise complexity of determining whether a given allocation is CEEI
- finding an "efficient" algorithm determining whether a CEEI allocation exists for a given add-MARA instance and giving one if it exists.

3.6 A scale of criteria

Putting Propositions 2, 4, 6 and 7 together leads to the following implication sequence, for any allocation $\overrightarrow{\pi}$: $(\overrightarrow{\pi} \models \text{CEEI}) \Rightarrow (\overrightarrow{\pi} \models \text{EF}) \Rightarrow (\overrightarrow{\pi} \models \text{mFS}) \Rightarrow (\overrightarrow{\pi} \models \text{PFS}) \Rightarrow (\overrightarrow{\pi} \models \text{MFS})$.

In other words, these criteria can be ranked from the least to the more demanding as follows:

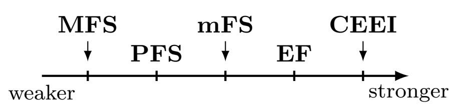

As the propositions also show, these results can also be interpreted the other way around, in terms of add-MARA instances: $\mathcal{I}_{|\text{CEEI}} \subset \mathcal{I}_{|\text{EF}} \subset \mathcal{I}_{|\text{mFS}} \subset \mathcal{I}_{|\text{PFS}} \subset \mathcal{I}_{|\text{MFS}} \subset \mathcal{I}$,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

In can be noticed that all these criteria have a kind of distributed flavor. MFS, PFS and mFS are of similar nature: every agent, only considering her own share, is able to judge whether she is satisfied or not. Envy-freeness requires the additional knowledge of the other shares, but each agent is still able to assert on her own whether she is envious or not. As for the CEEI, once the prices are fixed by the arbitrator, each agent is able to compute her own share (up to some equivalent shares).

Beyond their differences, these five criteria all have a common very appealing feature: they are not based on interpersonal comparison of utilities (actually, four of them are even purely ordinal — PFS is not). This leads to the following (easy) proposition:

Proposition 9. *The max-min fair-share, proportional fair-share, min-max fair-share, envy-freeness and CEEI criteria are preserved by any linear dilatation of individual utility scales.*

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

Finally, MFS, PFS and mFS have an interesting characteristic, which comes from the fact that they are all defined as minimum thresholds to satisfy: if for a given add-MARA instance there is an allocation satisfying one of these three criteria, then either this allocation is Pareto-efficient, or there exists another allocation which both satisfies Pareto-efficiency and this criterion. This is not the case for the envy-free criterion : as Example 6 shows, one can find instances having envy-free allocations, none of them being Pareto-efficient.

4 The egalitarian approach

As pointed out in the beginning of Section 3, an orthogonal approach for ensuring fairness in resource allocation problems, when agents preferences are numerically expressed, is to base the sharing process upon a collective utility function, to be maximized. Probably the most prominent one is the egalitarian CUF, which can be defined as follows in our sharing context :

Definition 7. *Let $\langle \mathcal{A}, \mathcal{O}, w \rangle$ be an add-MARA instance. The egalitarian CUF is the function $g_e : \overrightarrow{\pi} \mapsto \min_{i \in \mathcal{A}} u_i(\pi_i)$. Any allocation maximizing the egalitarian CUF will be called* min-optimal*.*

This CUF is the formal translation of Rawlsian egalitarianism [29], which recommends to maximize the utility of the least well-off agent.

At this point a natural question arises : what are the links between on one hand the egalitarianist approach of fairness (the min CUF) and on the other hand the approach based on criteria (Section 3) ? Actually, some criteria are not fully compatible with egalitarianism. For example, an envy-free allocation can be far away from being min-optimal. This question is raised in [10]. However, as we will see just below, egalitarianism is more compatible with proportional fair-share as well as with max-min fair-share.

As egalitarianism requires inter-agent comparisons of utilities,[4] we assume in this section normalized weights, namely: there is a constant K such that for all i, $\sum_{l \in \mathcal{O}} w(i,l) = K$.

[4]Conversely, the five criteria (Section 3) do not require inter-agent comparisons of utilities, even if our examples use normalized weights.

Proposition 10. *If there is an allocation satisfying the proportional fair-share criterion (with normalized weights), then any min-optimal allocation satisfies this criterion.*

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

This proposition also gives a practical way to find an allocation satisfying proportional fair-share if there is one: normalize weights and compute a min-optimal allocation. If this allocation obeys the criterion, then we get it, otherwise there is no such allocation.

Things are less clear for max-min fair-share. On the one hand, the latter result does not hold for max-min fair-share,[5] as shows the following counter-example with 3 agents and 4 objects ($K = 100$).

Example 7. *Consider the following instance, given by its weight matrix.*

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

The max-min fair-share of each agent (on the right) and the corresponding shares are marked with '$*$'. A min-optimal allocation and the corresponding utilities are marked with '†'[6]*. In this min-optimal allocation, the third agent does not get her max-min fair-share (expecting at least 27 but getting only 26).*

There are for this instance allocations obeying the max-min fair-share criterion, for example $\langle\{2,4\},\{1\},\{3\}\rangle$, but none of them are min-optimal. Moreover, the min-optimal allocation does not provide her proportional fair-share to agent 3 (26 < 100/3). Hence from Proposition 10, we know that this instance admits no PFS *allocation, and from Propositions 6 and 7, it admits neither* mFS, EF *or* CEEI *allocations.*

On the other hand however, such a counter-example has a small chance to appear in practice: for example, using a uniform generation process similar to the impartial culture in vote theory (see Section 6), for 3 agents and 4 objects, only one instance over 3500 is a counter-example similar to Example 7. This shows that max-min fair-share has a good correlation with the egalitarian approach.

---

[5] Actually a similar result holds if weights are normalized such that u_i^{MFS} is equal for all agents (and not u_i^{PFS}).

[6] This min-optimal allocation is also leximin-optimal. The leximin ordering [30] is a refinement of the min ordering for which a lexicographic comparison of sorted vectors of weigths is used, instead of comparing their min values.

5 Restricted cases

In this section we examine the behavior of our criteria – and especially the max-min fair-share one – in some restricted cases, giving to these criteria an additional insight. These restrictions concern 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 the max-min fair-share criterion.

5.1 Restricted preferences

5.1.1 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 allocation satisfying max-min fair-share 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* [10] or *elicitation-free sequential protocol* [8]. 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 ("Approval sharing")**.** *Any add-MARA instance with weights restricted to* $0, 1$ *belongs to* $\mathcal{I}_{|\mathrm{MFS}}$.

Proof. In an instance with n agents and m objects, the max-min fair-share of agent i is $\lfloor \frac{s_i}{n} \rfloor$, with $s_i = \sum_{l=1}^{m} w(i,l)$. The following very simple algorithm (a picking protocol) gives any agent her max-min fair-share :

```
while ( true )
   for i = 1 to n
      Allocate to agent i an object l not allocated yet
         such that w(i,l) = 1 if any,
      otherwise allocate to i any remaining object of weight 0.
      If all objects are allocated, exit.
```

There are exactly $\lfloor \frac{m}{n} \rfloor$ complete passages in the for loop. During each complete passage, n objects are allocated, one to each agent. During each of the first $\lfloor \frac{s_i}{n} \rfloor$ passages at least, agent i receives an object of weight 1. ☐

5.1.2 Identical preferences

When agents give the same object the same weight (thus they have identical utility functions), our scale of criteria has only two levels : the max-min fair-share criterion, and the others merged in one.

Proposition 12. *If agents have identical preferences (for all* $i, j, l : w(j,l) = w(i,l)$*), then*

1. *there is always an allocation satisfying the max-min fair-share criterion, and in particular any min-optimal allocation satisfies it ;*

2. *if preferences are strict (i.e., for any agent, two distinct shares have different utilites), no allocation satisfies the proportional fair-share criterion, and thus none satisfies the three more demanding criteria ;*

3. *Let $\overrightarrow{\pi}$ be any admissible allocation. The following five propositions are equivalent : (i) each agent in $\overrightarrow{\pi}$ gets the same utility ; (ii) $\overrightarrow{\pi} \models$ CEEI ; (iii) $\overrightarrow{\pi} \models$ EF ; (iv) $\overrightarrow{\pi} \models$ mFS ; (v) $\overrightarrow{\pi} \models$ PFS.*

Proof. 1. Consider a min-optimal allocation $\overrightarrow{\pi}^\star$. Then for each agent i :

$$\begin{aligned} u_i^{\text{MFS}} &\stackrel{def}{=} \max_{\overrightarrow{\pi} \in \mathcal{F}} \min_{j \in \mathcal{A}} u_i(\pi_j) \\ &= \max_{\overrightarrow{\pi} \in \mathcal{F}} \min_{j \in \mathcal{A}} u_j(\pi_j) \\ &= \min_{j \in \mathcal{A}} u_j(\pi_j^\star) \leq u_i(\pi_i^\star) \end{aligned}$$

2. Because preferences are strict, for any allocation $\overrightarrow{\pi}$, the n numbers $u_i(\pi_i)$ are different. One of them at least is strictly smaller than their mean.
3. We prove implications in the order given in the proposition. Let $\overrightarrow{\pi}$ be an allocation in which agents get the same utility, which is $u_i(\mathcal{O})/n$ for agent i. Consider the following price vector : $p_l = nw(i,l)/u_i(\mathcal{O})$. Total price is n, and the price of every share of $\overrightarrow{\pi}$ is exactly 1. So each agent can buy any share of $\overrightarrow{\pi}$, and any share that would provide more utility costs necessarily more. Hence $\overrightarrow{\pi} \models$ CEEI. The other three implications follow from the scale of criteria (Section 3). The last one, closing the cycle ($\overrightarrow{\pi} \models$ PFS implies shares of equal utility) is easily proved. □

5.1.3 Same-order preferences

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 — we say for short : the instance is 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 (object 1 is one of the best, object 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 $[\![1, m]\!]$ 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 an allocation satisfying max-min fair-share for a given SOP add-MARA instance, then we can find one for every permutation derived from it:

[7]This property is sometimes known as full-correlation [8].

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

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 $\overrightarrow{\pi}^{\uparrow}$ be an allocation satisfying the max-min fair-share criterion for the SOP instance $\langle \mathcal{A}, \mathcal{O}, w^{\uparrow}\rangle$. Let $S = S_1, S_2, ...S_m$ be the sequence of agents defined as follows: S_l is the agent who receives object l in $\overrightarrow{\pi}^{\uparrow}$.

Such a sequence which depends on $\overrightarrow{\pi}^{\uparrow}$ always exists because the instance is SOP, and because each object is given to exactly one agent. As said before, the sequence S is called a *sequence of sincere choices*, after [10], which however uses it in a different way.

The key is to notice that the allocation $\overrightarrow{\pi}$ 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 $\overrightarrow{\pi}^{\uparrow}$. To see it, notice that before step p in the building of $\overrightarrow{\pi}$, 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 $\overrightarrow{\pi}^{\uparrow}$. Consequently, for each agent i and each object of $\pi_i^{\uparrow}$, there is an object in π_i which is weakly better for i: the utility of i weakly increases from $\overrightarrow{\pi}^{\uparrow}$ (in $\langle \mathcal{A}, \mathcal{O}, w^{\uparrow}\rangle$) to $\overrightarrow{\pi}$ (in $\langle \mathcal{A}, \mathcal{O}, w\rangle$).

Since the max-min fair-share 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 $\overrightarrow{\pi}^{\uparrow} \models \mathrm{MFS}$, and $\overrightarrow{\pi}$ makes everyone at least as well-off, we conclude $\overrightarrow{\pi} \models \mathrm{MFS}$. ☐

Because any 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 ones as far as the max-min fair-share criterion is concerned.[8] So, to prove that all instances of a given subset satisfy this criterion, we only need to prove that any SOP instances of that subset satisfies it. In the following, we will often consider only the SOP instances of the subsets of interest, and hence the results obtained for them will be valid also for all instances of the subset.

5.1.4 Weights defined by a scoring function

We consider here the case where agents express their preferences using exactly the same multiset of weights (formally, for all $(i, j) \in \mathcal{A}^2$, $\{\!\{w_{i,l} \mid l \in \mathcal{O}\}\!\} = \{\!\{w_{j,l} \mid l \in \mathcal{O}\}\!\}$ where $\{\!\{\}\!\}$ denotes a multiset). Equivalently we could say that agents use the same *scoring function*. A scoring function is simply a weakly decreasing function $g : [\![1, m]\!] \to \mathbb{R}^+$. It can be used to convert a purely ordinal expression of preferences into to a numerical one, in the following way. Consider that each agent ranks strictly the objects from 1 (the most prefered) to m (the least prefered). If $r(i, l)$ is the rank given to object l by agent i, then the weight $w(i, l)$ is defined as $g(r(i, l))$.

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

This framework, which is standard in social choice, allows to link a purely ordinal expression of preferences to a numerical one. It is the basis of well-known procedures in voting theory (plurality, veto, Borda scores for examples). It has been already considered in fair division of indivisible goods [8] and social choice [4].

Proposition 14. *Any add-MARA instance in which preferences are defined by the same scoring function is in $\mathcal{I}_{|\mathrm{MFS}}$, and any min-optimal allocation satisfies max-min fair-share in this case. The other conclusions (ii) and (iii) of Proposition 12 are also satisfied.*

Proof. By Proposition 13 it is enough to consider SOP instances, which are in this case instances with identical preferences. Then use Proposition 12. □

5.2 Restrictions upon the number of agents and objects

5.2.1 Two agents

The two agents case is particularly interesting because the famous cut-and-choose game give their max-min fair-share to both agents.

Proposition 15. *Any 2-agents add-MARA instance belongs to $\mathcal{I}_{|\mathrm{MFS}}$.*

Proof. Let $\overrightarrow{\pi}$ such that $\overrightarrow{\pi} = \mathrm{argmax}_{\overrightarrow{\pi}' \in \mathcal{F}} \min_{j \in \mathcal{A}} u_1(\pi'_j)$. By definition, π_1 and π_2 give agent 1 her max-min fair-share. But $u_2(\pi_1) + u_2(\pi_2) = u_2(\mathcal{O})$, so one of the two shares (suppose π_2) is such that $u_2(\pi_2) \geq \frac{1}{2} u_2(\mathcal{O}) = u_2^{\mathrm{PFS}}$. As $u_2^{\mathrm{PFS}} \geq u_2^{\mathrm{MFS}}$ (by Proposition 2), $u_2(\pi_2) \geq u_2^{\mathrm{MFS}}$, hence the allocation $\langle \pi_1, \pi_2 \rangle$ satisfies the max-min fair-share criterion.

Another pleasant proof : agent 1 cuts, so she will always get her max-min fair-share. Agent 2 chooses first, so she gets her min-max fair-share, therefore her max-min fair-share too (by Propositions 2 and 4). □

5.2.2 No more objects than agents

If there are strictly less objects than agents, the scale of criteria is reduced to only one level, and hence is of no help[9]. The case as many objects as agents puts in light the min-max fair-share criterion.

Proposition 16. *If there are strictly less objects than agents, any allocation satisfies the max-min fair-share criterion, but none satisfies the other criteria.*

If there are as many objects as agents, then

1. *any allocation which is a matching (giving to each agent one object) satisfies the max-min fair-share criterion.*

2. *any allocation satisfying the min-max fair-share criterion is a matching, envy-free, Pareto-efficient and CEEI.*

[9] The best resort in this case would be a normalized leximin-optimal allocation.

Proof. Case $m < n$: in any allocation, one agent at least receives no object, hence $u_i^{\mathrm{MFS}} = 0$ for all i. As $0 \leq u_i(\pi_i)$ for all i, each agent gets her max-min fair-share. Of course no allocation satisfies the proportional fair-share criterion.

Case $m = n$:
1. We have easily $u_i^{\mathrm{MFS}} = \min_{l \in \mathcal{O}} w(i,l)$, hence each agent receives her max-min fair-share in a matching.
2. We have also easily $u_i^{\mathrm{mFS}} = \max_{l \in \mathcal{O}} w(i,l)$. In an allocation satisfying the min-max fair-share criterion each agent receives a prefered object. The allocation is hence an envy-free matching. It is Pareto-efficient because for an agent to get strictly more utility, she necessarily has to take another object from another agent, strictly reducing this agent's utility. A price of 1 for each object provides a CEEI allocation, because any increase of utility must be paid more. □

5.2.3 Up to three more objects than agents

We prove that any instance with up to three more objects than agents satisfies the max-min fair-share criterion. We begin by the case $m = n + 1$.

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

Proof. Thanks to Proposition 13, we can restrict to SOP instances. Since objects n and $n+1$ are the worst ones, it is not difficult to see that all the shares from allocation $\langle \{1\}\{2\} \dots \{n-1\}\{n, n+1\} \rangle$ give to each agent her max-min fair-share. □

To continue with $m = n + 2$ and $m = n + 3$, we need first a convenient definition of the extension of an instance, that is adding p agents and q objects to a given instance, and an additional notation.

Definition 8. *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 $(i,l) \in \mathcal{A} \times \mathcal{O}$.*

Denote by $u_i^{\mathrm{MFS}}(I)$ the max-min fair-share of agent i in instance I.

Then we give some preliminary lemmas. The first one shows the behavior of $u_i^{\mathrm{MFS}}(I)$ when k additional agents and objects are added to I.

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 strictly increase from I to any (k,k)-extension of I. Formally : $u_i^{\mathrm{MFS}}(I_{+k,+k}) \leq u_i^{\mathrm{MFS}}(I)$ for all integer $k > 0$.*

Proof. Let I be an instance, and I' a $(1,1)$-extension of it. Start from an allocation $\overrightarrow{\pi}'$ of I' that gives her max-min fair-share to agent i in it, that is :

$$\min_{j=1}^{n+1}(u_i(\pi'_j)) = u_i^{\mathrm{MFS}}(I') \tag{2}$$

Removing from $\overrightarrow{\pi}'$ the share containing object $m+1$ yields a valid (possibly incomplete) allocation $\overrightarrow{\pi}$ for I. Hence,

$$\min_{j=1}^{n}(u_i(\pi_j)) \leq u_i^{\mathrm{MFS}}(I) \tag{3}$$

Removing a number from a set cannot strictly decrease its minimum, so we have

$$\min_{j=1}^{n+1}(u_i(\pi'_j)) \leq \min_{j=1}^{n}(u_i(\pi_j)) \tag{4}$$

By Equations 2 to 4 we conclude

$$u_i^{\mathrm{MFS}}(I') \leq u_i^{\mathrm{MFS}}(I) \tag{5}$$

Iterating the previous result from 1 to k completes the proof. □

The next (easy) lemma gives a very simple expression of max-min fair-share, proportional fair-share and min-max fair-share when all weights are equal.

Lemma 2. *Let I be an instance with n agents and m objects, all weights beeing equal to 1. Then $u^{\mathrm{MFS}}(I) = \lfloor \frac{m}{n} \rfloor$, $u^{\mathrm{PFS}}(I) = \frac{m}{n}$, $u^{\mathrm{mFS}}(I) = \lceil \frac{m}{n} \rceil$.*

Proof. First equality : $m = n\lfloor \frac{m}{n} \rfloor + r = (n-r)\lfloor \frac{m}{n} \rfloor + r(\lfloor \frac{m}{n} \rfloor + 1)$ with $0 \leq r < n$. So in an allocation satisfying the max-min fair-share criterion, $n - r$ agents receive $\lfloor \frac{m}{n} \rfloor$ and r agents receive $\lfloor \frac{m}{n} \rfloor + 1$. The second equality is just the definition of the proportional fair-share. For the third equality, the proof is similar to the first. □

The next lemma gives an upper bound of the max-min fair-share of an agent.

Lemma 3. *Let I be an add-MARA instance. For any agent i, we have $u_i^{\mathrm{MFS}}(I) \leq \lfloor \frac{m}{n} \rfloor \max_{l=1}^{m} w(i,l)$. In particular when $m < 2n$ then $u^{\mathrm{MFS}}(I) \leq \max_{l=1}^{m} w_l$.*

Proof. Follows from the first equality of Lemma 2 and the fact that $u^{\mathrm{MFS}}(I)$ is obviously a weakly increasing function of each $w(i,l)$. □

We can now use these lemmas to show the following more general result:

Lemma 4. *Let I be an add-MARA instance. If $I \in \mathcal{I}_{|\mathrm{MFS}}$ and $n \leq m \leq 2n$ then any (p,p)-extension of I is in $\mathcal{I}_{|\mathrm{MFS}}$.*

Proof. We prove the lemma for $p = 1$. The result for $p > 1$ is obtained by induction over p.

Take any instance I' with $n+1$ agents and $m+1$ objects. We have to prove that there is an allocation $\overrightarrow{\pi}'$ for I' such that $\overrightarrow{\pi}' \models \mathrm{MFS}$.

As explained in the discussion following Proposition 13, we consider SOP instances, and without loss of generality, we suppose that[10] $w(i,1) \geq w(i,2) .... \geq w(i,m) \geq w(i,m+1)$ for all i, $1 \leq i \leq n+1$.

[10]Actually, for the following demonstration, only the fact that $w(i,1)$ is greater or equal to other weights is necessary. But the full hypothesis (SOP instances) is required in the proof by recurrence for $p > 1$.

Restrict I' by removing agent $n+1$ and object 1. We obtain a (n, m)-instance I which, by hypothesis, has an allocation $\overrightarrow{\pi}$ such that $\overrightarrow{\pi} \models$ MFS. Extend it to

$$\overrightarrow{\pi}' = \langle \pi_1, \cdots, \pi_n, \{1\} \rangle \tag{6}$$

($\overrightarrow{\pi}'$ is $\overrightarrow{\pi}$ augmented with a new share built with the object 1 alone). $\overrightarrow{\pi}'$ is a valid allocation for I'. We show now that $\overrightarrow{\pi}' \models$ MFS. The hypothesis $n \leq m \leq 2n$ is equivalent to $\lfloor \frac{m+1}{n+1} \rfloor = 1$. Hence, by Lemma 3 :

$$u^{\mathrm{MFS}}(I') \leq w_1 \tag{7}$$

proving that agent $n+1$ obtains her max-min fair-share in $\overrightarrow{\pi}'$. As for other n first agents, Lemma 1 says that this is also the case for them in $\overrightarrow{\pi}'$ (they get the same share in I and I', hence same value, and their max-min fair share cannot strictly increase from I to I'). □

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

Proof. This is a direct consequence of Proposition 15 (showing that any instance with 2 agents and 4 objects belongs to $\mathcal{I}_{|\mathrm{MFS}}$), and Lemma 4. □

The case with n and $n+3$ objects can also be proved using Lemma 4, but for that we need to prove the base case with 3 agents and 6 objects.

Lemma 5. *Any add-MARA instance with* 3 *agents and* 6 *objects belongs to* $\mathcal{I}_{|\mathrm{MFS}}$.

Proof. As usual, we consider a SOP $(3, 6)$-instances I', and without loss of generality, we suppose that $w(i, 1) \geq w(i, 2) \geq w(i, 4) \geq w(i, 5) \geq w(i, 6)$ for all i. We will show that I' belongs to $\mathcal{I}_{|\mathrm{MFS}}$ by building an allocation $\overrightarrow{\pi}'$ such that $\overrightarrow{\pi}' \models$ MFS. We consider two subcases.

(i) If for an agent, say agent 3, we have $u_3^{\mathrm{MFS}} \leq w(3, 1)$, then give the share (1) to this agent, so she gets her max-min fair share. Remains a $(2, 5)$-instance I that belongs to $\mathcal{I}_{|\mathrm{MFS}}$ by Proposition 15. Hence there exists an allocation $\overrightarrow{\pi}$ which gives her max-min fair-share to both agents (1) and (2) in I. Extend $\overrightarrow{\pi}$ to $\overrightarrow{\pi}' = \langle \pi_1, \pi_2, \{1\} \rangle$ ($\overrightarrow{\pi}'$ is $\overrightarrow{\pi}$ augmented with a new share built with the object 1 alone). $\overrightarrow{\pi}'$ is a valid allocation for I'. Agent 3 gets her max-min fair-share in I', as said before. Agents 1 and 2 get the same share in I and I', hence the same value. By Lemma 1, their max-min fair share cannot strictly increase from I to I', so they also get their max-min fair-share in I'.

(ii) Otherwise, we have $w(i, 1) < u_i^{\mathrm{MFS}}$ and then $w(i, l) < u_i^{\mathrm{MFS}}$ for all i and l. So, any allocation satisfying the max-min fair-share criterion cannot include a share with a single object, and must have 2 objects in each share. It is not too difficult to show[11] that an allocation satisfying the max-min fair-share criterion is $\langle \{1, 6\}\{2, 5\}\{3, 4\} \rangle$ (up to a permutation). This allocation gives her max-min fair-share to all agents. □

[11] Let $\overrightarrow{\pi}^* = \langle \{1, 6\}\{2, 5\}\{3, 4\} \rangle$. Check that, for every allocation $\overrightarrow{\pi}$ with only 2 objects by share, for every share $\pi^* \in \overrightarrow{\pi}^*$ there is a share $\pi \in \overrightarrow{\pi}$ such that π has less or equal utility than π^*.

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

Proof. This is a direct consequence of Lemmas 4 and 5. □

6 Experiments

Tables 1 and 2 give some experimental results concerning our scale of criteria.

We generated 1000 couples of random instances, for n, the number of agents, ranging from 3 to 5, and for m, the number of objects, ranging from 1 to 11. In each couple of instance, the first is non SOP, the second is the SOP version of the first. In Table 1, weights are uniformly generated in $[0,1]$. In Table 2, weights are drawn from a Gaussian distribution, mean 0.5 and standard deviation 0.1.

The number on line n, m and column $\mathcal{C}$ gives the number of instances, out of 1000, which satisfy the criterion $\mathcal{C}$. The last column is not devoted, as could be expected, to the criterion CEEI, but to the criterion EFP which means envy-free and Pareto-efficient. In fact, it is computationally difficult to characterize exactly a CEEI instance in general (see [26, Section 3]), so in experiments we replaced CEEI by EFP[12].

Several facts can be noticed, which comfort our theoretical results.

- Main result : the scale of properties is really significative, of course when $n \leq m$. The numbers weakly decrease from left to right, and often strictly decrease, showing that the scale is not trivial.
- SOP instances are more conflicting than non SOP ones, in accordance with Lemma 13.
- In Table 1 (uniform distribution of weigths) 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.
- In Table 2 (gaussian distribution of weights, mean 0.5, standard deviation 0.1): instances where m is close to a multiple of n are less conflict-prone than others, which is not very surprising.
- All generated instances belong to $\mathcal{I}_{|\text{MFS}}$. This shows that it is actually very unlikely to find an instance not in $\mathcal{I}_{|\text{MFS}}$ (at least with uniform and gaussian generation of weights) even if such instances exist [28].

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 allocation satisfying the max-min fair-share criterion, things are surprisingly different for more general non-additive preferences.

[12] We saw in Section 3.5 that EFP is a necessary condition for CEEI when preferences are strict. We believe that CEEI and EFP are not equivalent in the context of this discrete model.

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

Table 1: Experimental results with a uniform distribution of weights.

Gauss		Non SOP instances					SOP instances				
n	m	MFS	PFS	mFS	EF	EFP	MFS	PFS	mFS	EF	EFP
3	1	1000	0	0	0	0	1000	0	0	0	0
3	2	1000	0	0	0	0	1000	0	0	0	0
3	3	1000	610	221	221	221	1000	2	0	0	0
3	4	1000	26	1	0	0	1000	0	0	0	0
3	5	1000	3	2	1	1	1000	3	2	2	2
3	6	1000	994	960	915	886	1000	647	218	218	218
3	7	1000	737	256	41	40	1000	185	55	45	44
3	8	1000	223	181	153	122	1000	223	151	123	123
3	9	1000	1000	1000	1000	1000	1000	999	967	870	839
3	10	1000	998	935	663	624	1000	908	738	653	645
3	11	1000	873	852	847	782	1000	873	847	829	807
4	1	1000	0	0	0	0	1000	0	0	0	0
4	2	1000	0	0	0	0	1000	0	0	0	0
4	3	1000	0	0	0	0	1000	0	0	0	0
4	4	1000	740	92	92	92	1000	0	0	0	0
4	5	1000	82	0	0	0	1000	0	0	0	0
4	6	1000	2	1	0	0	1000	0	0	0	0
4	7	1000	0	0	0	0	1000	0	0	0	0
4	8	1000	999	996	961	918	1000	767	86	86	85
4	9	1000	993	393	20	16	1000	267	21	14	13
4	10	1000	622	219	19	13	1000	114	24	14	13
4	11	1000	268	224	191	120	1000	268	216	157	140
5	1	1000	0	0	0	0	1000	0	0	0	0
5	2	1000	0	0	0	0	1000	0	0	0	0
5	3	1000	0	0	0	0	1000	0	0	0	0
5	4	1000	0	0	0	0	1000	0	0	0	0
5	5	1000	843	57	57	57	1000	0	0	0	0
5	6	1000	254	0	0	0	1000	0	0	0	0
5	7	1000	8	0	0	0	1000	0	0	0	0
5	8	1000	1	0	0	0	1000	0	0	0	0
5	9	1000	2	1	0	0	1000	2	0	0	0
5	10	1000	1000	1000	994	969	1000	854	57	57	56
5	11	1000	1000	608	6	6	1000	400	13	6	6

Table 2: Experimental results with a Gaussian distribution of weights. Mean = 0.5, standard deviation = 0.1

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 originally introduced in the context of fuzzy measures [18], and also used in the context of resource allocation [14].

Formally, we consider in this section k-additive multiagent 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, corresponding to the model introduced in Section 2), and thus forbids any preferential interdependence between objects. A 2-additive function allows such interdependence: for example, the weight $w(\{1, 2\})$ stands for the proper interest of the pair of objects $\{1, 2\}$ beyond these two individual objects: if $w(\{1, 2\}) > 0$, the value of this pair is more important than the intrinsic value of the two separated objects (which shows that they are complementary); if $w(\{1, 2\}) < 0$, they are substitutable.

As soon as we switch from 1-additive to 2-additive functions, finding an instance not belonging to $\mathcal{I}_{|\text{MFS}}$ (that is for which no allocation satisfying the max-min fair-share criterion exists) is not challenging anymore:

Example 8. *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-Add-MFS-Exist]) is even hard:

Proposition 20. [k-Add-MFS-Exist] *is* NP*-hard, for* $k \geq 2$ *and* $n \geq 3$.

Proof. NP-hardness can be proved by reduction from the partition problem (Problem 2). Let $\langle \{x_1, \ldots, x_n\}, s \rangle$ 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, $w(i, \{l\}) = s(x_l)$ and $w(i, \{l, n + m\}) = -3L$ for all $l \in [\![1, n]\!]$ and $m \in [\![1, 4]\!]$;
- $w(1, \{n + 1, n + 2\}) = w(1, \{n + 3, n + 4\}) = L$
- $w(2, \{n + 1, n + 3\}) = w(2, \{n + 2, n + 4\}) = L$

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

Let us compute the max-min fair share for each agent. Let us consider the allocation $(\{1, \ldots, n\}, \{n+1, n+2\}, \{n+3, n+4\})$. The evaluation of these three shares by agent 1 gives respectively $2L$, L, and L. Hence $u_1^{\mathrm{MFS}} \geq L$.

Let now $\overrightarrow{\pi}$ be a custom allocation. Three cases are possible. (i) There is a share π_i such that $\pi_i \supseteq \{l, m\}$, with $l \leq n$ and $m > n$. Then $u_1(\pi_i) \leq 0$, (ii) the objects $m > n$ are split into two different shares (say w.l.o.g π_1 and π_2) both containing no object from $l \leq n$. Then $\pi_1 \leq L$ or $\pi_2 \leq L$. (iii) the objects $m > n$ only appear in one share (say w.l.o.g π_1). In that case, the objects $l \leq n$ are split between the two shares π_2 and π_3. Since the utility function of agent i is additive on the objects of $\{1, \ldots, n\}$ and that $\sum_{l=1}^{n} w(1, \{l\}) = 2L$, we have $u_i(\pi_2) \leq L$ or $u_i(\pi_3) \leq L$. In cases (i), (ii) and (iii), $\min_{i \in \mathcal{A}} u_1(\pi_i) \leq L$. Hence $u_1^{\mathrm{MFS}} = K$. The other agents' case can be treated similarly.

Let $\overrightarrow{\pi}$ be once again a custom allocation. If $\overrightarrow{\pi}$ has a share π_i such that $\pi_i \supseteq \{l, m\}$, with $l \leq n$ and $m > n$, then as previously $u_i(\pi_i) \leq 0 < u_i^{\mathrm{MFS}}$. Now let us assume the contrary, and thus suppose there is a share (say w.l.o.g π_1) that only contains objects $m > n$. Then either $u_1(\pi_1) = 0$, which means that agent 1 does not receive her max-min fair share, or $u_1(\pi_1) > 0$ and thus $\pi_1 \supseteq \{n+1, n+2\}$, or $\pi_1 \supseteq \{n+3, n+4\}$. In any of the two last cases, we can easily see that if another share contains objects $m > n$ (hence contains only this kind of objects, according to our initial assumption), then the resulting utility for the concerned agent is 0, and hence the allocation will not give her her max-min fair-share.

Let us suppose then that π_2 and π_3 only contain objects $l \leq n$. $\overrightarrow{\pi}$ gives agents 2 and 3 their max-min fair share if and only if $u_2(\pi_2) \geq L$ and $u_3(\pi_3) \geq L$, which comes down to find a partition of the objects l into 2 subsets of value L, and hence a partition of $\mathcal{X}$ in the initial instance of [Partition]. □

It can be noticed that Proposition 20 only gives a NP-hardness result, as it is not known yet whether [k-Add-MFS-Exist] belongs to NP. We can only say that this problem belongs to Σ_2^{P}, because it can be solved by the same non-deterministic algorithm as in the additive case (see end of 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 [26]).

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[13] 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.

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] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. In *Proc. 13th ACM Conference on Electronic Commerce (EC-12)*, pages 197–214, June 2012.

[5] S. Bouveret, U. Endriss, and J. Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proceedings of the 3rd International Workshop on Computational Social Choice (COMSOC'10)*, Düsseldorf, Germany, September 2010.

[6] S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. In L. P. Kaelbling and A. Saffiotti, editors, *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI-05)*, Edinburgh, Scotland, Aug. 2005. Professional Book Center.

[13] In an ordinal setting, additivity should be replaced by its ordinal counterpart, separability.

[7] S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *Journal of Artificial Intelligence Research (JAIR)*, 32:525–564, 2008.

[8] S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proc. of IJCAI'11*, 2011.

[9] S. J. Brams, P. H. Edelman, and P. C. Fishburn. Fair division of indivisible items. *Theory and Decision*, 5(2), 2004.

[10] S. J. Brams and D. King. Efficient fair division – help the worst off or avoid envy? *Rationality and Society*, 17(4), 2005.

[11] S. J. Brams and A. D. Taylor. *The Win-win Solution. Guaranteeing Fair Shares to Everybody.* W. W. Norton & Company, 2000.

[12] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6), 2011.

[13] 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.

[14] 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.

[15] 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.

[16] S. Demko and T. P. Hill. Equitable distribution of indivisible items. *Mathematical Social Sciences*, 16:145–158, 1998.

[17] D. Foley. Resource allocation and the public sector. *Yale Econ. Essays*, 7(1), 1967.

[18] M. Grabisch. k-order additive discrete fuzzy measure and their representation. *Fuzzy Sets and Systems*, 92, 1997.

[19] D. K. Herreiner and C. Puppe. A simple procedure for finding equitable allocations of indivisible goods. *Social Choice and Welfare*, 19, 2002.

[20] T. P. Hill. Partinioning general probability measures. *The Annals of Probability*, 15(2):804–813, 1987.

[21] 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.

[22] R. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of EC'04*, 2004.

[23] E. Markakis and C.-A. Psomas. On worst-case allocations in the presence of indivisible goods. In *Proc. 7th International Workshop, WINE 2011*, Singapore, December 2011.

[24] J. S. Mill. *Utilitarianism.* University of Chicago Press, 1906.

[25] H. Moulin. *Cooperative Microeconomics, A Game-Theoretic Introduction.* Prentice Hall, 1995.

[26] A. Othman, T. Sandholm, and E. Budish. Finding approximate competitive equilibria: efficient and fair course allocation. In *Proceedings of AAMAS'10*, 2010.

[27] C. H. Papadimitriou. *Computational Complexity.* Addison–Wesley, 1994.

[28] A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proc. 14th ACM Conference on Economics and Computation (EC'14).* (Publication date as manuscript: February 2014).

[29] J. Rawls. *A Theory of Justice.* Harvard University Press, 1971.

[30] A. K. Sen. *Collective Choice and Social Welfare.* North-Holland, 1970.

[31] H. Steinhaus. The problem of fair division. *Econometrica*, 16(1), 1948.