Efficiency and Sequenceability in Fair Division of Indivisible Goods with Additive Preferences

Sylvain Bouveret and Michel Lemaître

Abstract

In fair division of indivisible goods, using sequences of sincere choices (or picking sequences) is a natural way to allocate the objects. The idea is the following: at each stage, a designated agent picks one object among those that remain. This paper, restricted to the case where the agents have numerical additive preferences over objects, revisits to some extent the seminal paper by Brams and King [10] which was specific to ordinal and linear order preferences over items. We point out similarities and differences with this latter context. In particular, we show that any Pareto-optimal allocation (under additive preferences) is sequenceable, but that the converse is not true anymore. This asymmetry leads naturally to the definition of a "scale of efficiency" having three steps: Pareto-optimality, sequenceability without Pareto-optimality, and non-sequenceability. Finally, we investigate the links between these efficiency properties and the "scale of fairness" we have described in an earlier work [8]: we first show that an allocation can be envy-free and non-sequenceable, but that every competitive equilibrium with equal incomes is sequenceable. Then we experimentally explore the links between the scales of efficiency and fairness.

1 Introduction

In this paper, we investigate fair division of indivisible goods. In this problem, a set of indivisible objects or goods has to be allocated to a set of agents (individuals, entities...), taking into account, as best as possible, the agents' preferences about the objects. This classical collective decision making problem has plenty of practical applications, among which the allocation of space resources [25, 5], of tasks to workers in crowdsourcing market systems, papers to reviewers [21] or courses to students [14].

This problem can be tackled from two different perspectives. The first possibility is to resort to a benevolent entity in charge of collecting in a *centralized* way the preferences of all the agents about the objects. This entity then computes an allocation that takes into account these preferences and satisfies some fairness (*e.g.* envy-freeness) and efficiency (*e.g.* Pareto-optimality) criteria, or optimizes a well-chosen social welfare ordering. The second possibility is to have a *distributed* point of view, *e.g.* by starting from an initial allocation and letting the agents negotiate to swap their objects [29, 15].

A somewhat intermediate approach consists in allocating the objects to the agents using a *protocol*, which can be seen as a way of building an allocation interactively by asking the agents a sequence of questions. Protocols are at the heart of works mainly concerning the allocation of divisible resources (cake-cutting) — see Brams and Taylor's seminal book [11] for a reference — but have also been studied in the context of indivisible goods [11, 9].

In this paper, we focus on a particular allocation protocol: *sequences of sincere choices* (also known as *picking sequences*). This very simple and natural protocol works as follows. A central authority chooses a sequence of agents before the protocol starts, having as many agents as the number of objects (some agents may appear several times in the sequence). Then, each agent appearing in the sequence is asked to choose in turn one object among those that remain. For instance, according to the sequence $\langle 1, 2, 2, 1 \rangle$, agent 1 is going to choose first, then agent 2 will pick two consecutive objects, and agent 1 will take the last object.

This simple protocol, actually used in a lot of everyday situations,[1] has been studied for the first time by Kohler and Chandrasekaran [24]. Later, Brams and Taylor [12] have studied a particular version of this protocol, namely alternating sequences, in which the sequence of agents is restricted to a balanced ($\langle 1, 2, 2, 1... \rangle$) or strict ($\langle 1, 2, 1, 2... \rangle$) alternation of agents. Bouveret and Lang [7] have further formalized this protocol, whose properties (especially related to game theoretic aspects) have been characterized by Kalinowski *et al.* [22, 23]. Finally, Aziz *et al.* [3] have studied the complexity of problems related to finding whether a particular assignment (or bundle) is achievable by a particular class of picking sequences.

On top of all these works specifically dedicated to this protocol, we can add the interesting work by Brams *et al.* [10]. This work, which focuses on a situation where the agents have ordinal preferences, is not specifically dedicated to picking sequences. However, the authors make an interesting link between this protocol and Pareto-optimality, showing, among others, that picking sequences always result in a Pareto-optimal allocation, but also that every Pareto-optimal allocation can be obtained in this way.

In this paper, we will elaborate on these ideas and analyze the links between sequences and some efficiency and fairness properties, in a more general model in which the agents have numerical additive preferences on the objects, without any further restriction. Our main contributions are the following. We give a formalization of the link between allocations and sequences of sincere choices, highlighting a simple characterization of the sequenceability of an allocation. Then, we show that in this slightly more general framework than the one by Brams *et al.*, surprisingly, Pareto-optimality and sequenceability are not equivalent anymore. As a consequence we can define a "scale of efficiency" that allows us to characterize the degree of efficiency of a given allocation. We also highlight an interesting link between sequenceability and another important economical concept: the competitive equilibrium from equal income (CEEI). Another contribution is the experimental exploration of the links between the scale of efficiency and fairness properties, which has led us to develop, among others, a practical method for testing if a given allocation passes the CEEI test — a practical problem which was still open [8].

The paper is organized as follows. Section 2 describes the model of our allocation problem and the allocation protocol based on sequences of sincere choices. Section 3 focuses on the problem consisting in deciding whether an allocation can be obtained by a sequence of sincere choices. In the next two sections, we analyze the relation between sequences of sincere choices and three classical properties: Pareto-optimality (Section 4), envy-freeness and competitive equilibrium from equal income (Section 5). Finally, we explore experimentally in Section 6 the links between the "scale of efficiency" mentioned above, and the "scale of fairness" that we have described in a previous work [8].

2   Model and Definitions

The aim the fair division of indivisible goods, also called MultiAgent Resource Allocation (MARA), is to allocate a finite set of *objects* $\mathcal{O} = \{1, \dots, M\}$ to a finite set of *agents* $\mathcal{A} = \{1, \dots, N\}$. A *sub-allocation* on $\mathcal{O}' \subseteq \mathcal{O}$ is a vector $\overrightarrow{\pi}^{|\mathcal{O}'} = \langle \pi_1^{|\mathcal{O}'}, \dots, \pi_N^{|\mathcal{O}'} \rangle$, such that $i \neq j \Rightarrow \pi_i^{|\mathcal{O}'} \cap \pi_j^{|\mathcal{O}'} = \emptyset$ (a given object cannot be allocated to more than one agent) and $\cup_{i \in \mathcal{A}} \pi_i^{|\mathcal{O}'} = \mathcal{O}'$ (all the objects from $\mathcal{O}'$ are allocated). $\pi_i^{|\mathcal{O}'} \subseteq \mathcal{O}'$ is called agent i's *share* on $\mathcal{O}'$. $\overrightarrow{\pi}^{|\mathcal{O}''}$ is a sub-allocation of $\overrightarrow{\pi}^{|\mathcal{O}'}$ when $\pi_i^{|\mathcal{O}''} \subseteq \pi_i^{|\mathcal{O}'}$ for each agent i. Any sub-allocation $\overrightarrow{\pi}^{|\mathcal{O}}$ on the entire set of objects will be denoted $\overrightarrow{\pi}$ and just called *allocation*.

Any satisfactory allocation must take into account the agents' preferences on the objects.

---

[1]Variants of this protocols are used, for instance, as drafting protocols by most sport leagues in North America, or in the board game *The Settlers of Catane* for allocating initial resources.

Here, we will make the classical assumption that these preferences are *numerically additive*. Each agent i has a *utility function* $u_i : 2^{\mathcal{O}} \to \mathbb{R}^+$ measuring her satisfaction $u_i(\pi)$ when she obtains share π, which is defined as follows:

$$u_i(\pi) \stackrel{def}{=} \sum_{\ell \in \pi} W(i, \ell),$$

where $W(i, \ell)$ is the weight given by agent i to object ℓ. This assumption, as restrictive as it may seem, is made by a lot of authors [26, 4, for instance] and is considered a good compromise between expressivity and conciseness.

If we put things together:

Definition 1. *An instance of the* additive multiagent resource allocation problem *(add-MARA instance for short)* $I = \langle \mathcal{A}, \mathcal{O}, W \rangle$ *is a tuple* $(\mathcal{A}, \mathcal{O}, W)$, *where:*

- $\mathcal{A} = \{1, \ldots, i, \ldots, N\}$ *is a set of N agents;*

- $\mathcal{O} = \{1, \ldots, \ell, \ldots M\}$ *is a set of M objects,*

- $W : \mathcal{A} \times \mathcal{O} \to \mathbb{R}^+$ *is a mapping, $W(i, \ell)$ being the weight given by agent i to object ℓ.*

We will denote by $\mathcal{P}(I)$ the set of allocations for I.

We will say that *the agents' preferences are strict on the objects* if, for each agent i and each pair of objects $\ell \neq m$, we have $W(i, \ell) \neq W(i, m)$. Similarly, we will say that *the agents' preferences are strict on the shares* if, for each agent i and each pair of shares $\pi \neq \pi'$, we have $u_i(\pi) \neq u_i(\pi')$. Finally, we will say that *the agents have same order preferences* [8] if there is a permutation $\eta : \mathcal{O} \mapsto \mathcal{O}$ such that for each agent i and each pair of objects ℓ and m, if $\eta(\ell) < \eta(m)$ then $W(i, \eta(\ell)) \geq W(i, \eta(m))$.

Observation 1. *Strict order on shares entails strict order on objects; the converse is false.*

The following definition will play a prominent role.

Definition 2. *Given an agent i and a set of objects O, let $\text{best}(O, i) = \text{argmax}_{\ell \in O} W(i, \ell)$ be the subset of objects in O having the highest weight for agent i (such objects will be called* top *objects of i). A (sub-)allocation $\overrightarrow{\pi}^{|\mathcal{O}'}$ is said* frustrating *if no agent receives one of her top objects in $\overrightarrow{\pi}^{|\mathcal{O}'}$ (formally: $\text{best}(\mathcal{O}', i) \cap \overrightarrow{\pi}_i^{|\mathcal{O}'} = \emptyset$ for each agent i), and* non-frustrating *if at least one agent receives one of her top objects in $\overrightarrow{\pi}^{|\mathcal{O}'}$.*

It should be emphasized that this notion of frustrating allocation was already present but implicit in the work by Brams and King [10]. Here, we bring this concept out because it will lead to a nice characterization of sequenceable allocations, as we will see later.

In the following, we will consider a particular way of allocating objects to agents: allocation by sequences of sincere choices. Informally the agents are asked in turn, according to a predefined sequence, to choose and pick a top object among the remaining ones.

Definition 3. *Let $I = \langle \mathcal{A}, \mathcal{O}, W \rangle$ be an add-MARA instance. A* sequence of sincere choices *(or simply* sequence *when the context is clear) is a vector of $\mathcal{A}^M$. We will denote by $\mathcal{S}(I)$ the set of possible sequences for instance I.*

Let $\overrightarrow{\sigma} \in \mathcal{S}(I)$. $\overrightarrow{\sigma}$ is said to *generate* allocation $\overrightarrow{\pi}$ if and only if $\overrightarrow{\pi}$ can be obtained as a possible result of the non-deterministic[2] Algorithm 1 on input I and $\overrightarrow{\sigma}$.

---

[2]The algorithm contains an instruction **choose** splitting the control flow into several branches, building all the allocations generated by $\overrightarrow{\sigma}$.

Algorithm 1: Execution of a sequence

Input: an instance $I = \langle \mathcal{A}, \mathcal{O}, W \rangle$ and a sequence $\overrightarrow{\sigma} \in \mathcal{S}(I)$
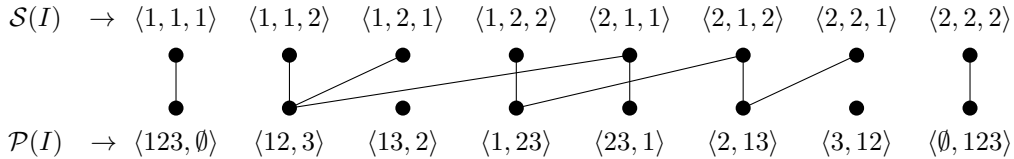**Output:** an allocation $\overrightarrow{\pi} \in \mathcal{P}(I)$

1 $\overrightarrow{\pi} \leftarrow$ empty allocation (such that $\forall i \in \mathcal{A} : \pi_i = \emptyset$);
2 $\mathcal{O}_1 \leftarrow \mathcal{O}$;
3 **for** t **from** 1 **to** M **do**
4 $\quad$ $i \leftarrow \sigma_t$;
5 $\quad$ **choose** object $o_t \in \text{best}(\mathcal{O}_t, i)$;
6 $\quad$ $\pi_i \leftarrow \pi_i \cup \{o_t\}$;
7 $\quad$ $\mathcal{O}_{t+1} \leftarrow \mathcal{O}_t \setminus \{o_t\}$

Definition 4. *An allocation $\overrightarrow{\pi}$ is said to be* sequenceable *if there exists a sequence $\overrightarrow{\sigma}$ that generates $\overrightarrow{\pi}$, and* non-sequenceable *otherwise. For a given instance I, we will denote by $s(I)$ the binary relation defined by $(\overrightarrow{\sigma}, \overrightarrow{\pi}) \in s(I)$ if and only if $\overrightarrow{\pi}$ can be generated by $\overrightarrow{\sigma}$.*

Example 1. *Let I be the instance represented by the following weight matrix:*[3]

$$\begin{pmatrix} 8 & 2 & 1 \\ 5 & 1 & 5 \end{pmatrix}$$

The binary relation $s(I)$ between $\mathcal{S}(I)$ and $\mathcal{P}(I)$ can be graphically represented as follows:

$\mathcal{S}(I) \quad \rightarrow \quad \langle 1,1,1 \rangle \quad \langle 1,1,2 \rangle \quad \langle 1,2,1 \rangle \quad \langle 1,2,2 \rangle \quad \langle 2,1,1 \rangle \quad \langle 2,1,2 \rangle \quad \langle 2,2,1 \rangle \quad \langle 2,2,2 \rangle$

$\mathcal{P}(I) \quad \rightarrow \quad \langle 123, \emptyset \rangle \quad \langle 12,3 \rangle \quad \langle 13,2 \rangle \quad \langle 1,23 \rangle \quad \langle 23,1 \rangle \quad \langle 2,13 \rangle \quad \langle 3,12 \rangle \quad \langle \emptyset,123 \rangle$

For instance, sequence $\langle 2,1,2 \rangle$ generates two possible allocations: $\langle 1,23 \rangle$ and $\langle 2,13 \rangle$, depending on whether agent 2 chooses object 1 or 3 that she both prefers. Allocation $\langle 12,3 \rangle$ can be generated by three sequences. Allocations $\langle 13,2 \rangle$ and $\langle 3,12 \rangle$ are non-sequenceable.

Once again, this notion of sequenceability is already implicitly present in the work by Brams and King [10], and has been extensively studied by Aziz *et al.* [3]. However, a fundamental difference is that in our setting, the preferences might be non strict on objects, which entails that the same sequence can yield different allocations (in the worst case, an exponential number), as Example 1 shows.

Observation 2. *For any instance I, $|\mathcal{S}(I)| = |\mathcal{P}(I)| = N^M$.*

Observation 3. *The number of objects allocated to an agent by a sequence is equal to the number of times the agent appears in the sequence. Formally: for all $(\overrightarrow{\sigma}, \overrightarrow{\pi}) \in s(I)$ and all agents i, $|\pi_i| = \sum_{\ell \in \mathcal{O}} [\sigma_\ell = i]$ where $[z = t]$ is 1 if the equality is verified, and 0 otherwise.*

3 Sequenceable allocations

In this section and in the following one, we will give a characterization of sequenceable allocations, that is, we will try to identify under which conditions an allocation is achievable by the execution of a sequence of sincere choices. The question has already been extensively studied by Aziz *et al.* [3], but in a quite different context — namely, ordinal strict preferences on objects — and with a particular focus on sub-classes of sequences (*e.g.* alternating sequences). As we will show, the properties are not completely similar in our context.

---

[3]In this example and in the following ones, we will represent instances by a matrix in which each value at row i and column ℓ represents the weight $W(i, \ell)$. We will also use notation $ab...$ as a shorthand for $\{a, b, ...\}$

3.1 Characterization

We have seen in Example 1 that some allocations are non-sequenceable. We will now formalize this and give a precise characterization of sequenceable allocations. We first start by noticing that in every sequenceable allocation, the first agent of the sequence gets a top object, which yields the following remark:

Observation 4. *Every frustrating allocation is non-sequenceable.*

Example 2. *In the following instance, the circled allocation $\langle 23, 1 \rangle$ is non-sequenceable because it is frustrating.*

$$\begin{pmatrix} 2 & \textcircled{1} & \textcircled{1} \\ \textcircled{1} & 2 & 2 \end{pmatrix}$$

However, it is possible to find a non-sequenceable allocation that gives her top object to one agent (as allocation $\langle 13, 2 \rangle$ in Example 1) or even to all, as the following example shows.

Example 3. *Consider this instance:*

$$\begin{pmatrix} \textcircled{9} & 8 & 2 & \textcircled{1} \\ 2 & \textcircled{5} & \textcircled{1} & 4 \end{pmatrix}$$

In the circled allocation $\overrightarrow{\pi} = \langle 14, 23 \rangle$, every agent receives her top object. However, after objects 1 and 2 have been allocated (they must be allocated first by all sequence generating $\overrightarrow{\pi}$), the sub-allocation shown in a dotted box above remains. This sub-allocation is obviously non-sequenceable because it is frustrating. Hence $\overrightarrow{\pi}$ is not sequenceable either.

This property of containing a frustrating sub-allocation exactly characterizes the set of non-sequenceable allocations:

Proposition 1. *Let $I = \langle \mathcal{A}, \mathcal{O}, W \rangle$ be an instance and $\overrightarrow{\pi}$ be an allocation of this instance. The two following statements are equivalent:*

(A) $\overrightarrow{\pi}$ is sequenceable.

(B) No sub-allocation of $\overrightarrow{\pi}$ is frustrating (in every sub-allocation, at least one agent receives a top object).

Proof. (B) implies (A). Let us suppose that for all subset of objects $\mathcal{O}' \subseteq \mathcal{O}$ there is at least one agent obtaining one of her top objects in $\overrightarrow{\pi}^{|\mathcal{O}'}$. We will constructively show that $\overrightarrow{\pi}$ is sequenceable. Let $\overrightarrow{\sigma}$ be a sequence of agents and $\overrightarrow{\mathcal{O}} \in (2^{\mathcal{O}})^M$ be a sequence of sets of objects jointly defined as follows:

- $\mathcal{O}_1 = \mathcal{O}$ and σ_1 is an agent that receives one of her top objects in $\overrightarrow{\pi}^{|\mathcal{O}_1}$;

- $\mathcal{O}_{t+1} = \mathcal{O}_t \setminus \{o_t\}$, where $o_t \in \text{best}(\mathcal{O}_t, \sigma_t)$ and σ_t is an agent that receives one of her top objects in $\overrightarrow{\pi}^{|\mathcal{O}_t}$, for $t \geq 1$.

From the assumption on $\overrightarrow{\pi}$, we can check that the sequence $\overrightarrow{\sigma}$ is perfectly defined. Moreover, $\overrightarrow{\pi}$ is one of the allocations generated by $\overrightarrow{\sigma}$.

(A) implies (B) by contraposition. Let $\overrightarrow{\pi}$ be an allocation containing a frustrating sub-allocation $\overrightarrow{\pi}^{|\mathcal{O}'}$. Suppose that there exists a sequence $\overrightarrow{\sigma}$ generating $\overrightarrow{\pi}$. We can notice that in Algorithm 1, when an object is allocated to an agent, all the objects which are strictly better for her have already been allocated at a previous step. Let $\ell \in \mathcal{O}'$, and let i be the agent that receives ℓ in $\overrightarrow{\pi}$. Since $\overrightarrow{\pi}^{|\mathcal{O}'}$ is frustrating, there is another object $m \in \mathcal{O}'$ such that $W(i, m) > W(i, \ell)$. From the previous remark, m is necessarily allocated before ℓ in

the execution of sequence $\overrightarrow{\sigma}$. We can deduce, from the same line of reasoning on m and agent j that receives it, that there is another object p allocated before m in the execution of the sequence. The set $\mathcal{O}'$ being finite, using the same argument iteratively, we will necessarily find an object which has already been encountered before. This leads to a cycle in the precedence relation of the objects in the execution of the sequence. Contradiction: no sequence can thus generate $\overrightarrow{\pi}$. $\qquad\square$

Beyond the fact that it characterizes a sequenceable allocation, the proof of Proposition 1 gives a practical way of checking if an allocation is sequenceable, and, if it is the case, of computing a sequence that generates this allocation. This yields the following result:

Proposition 2. *Let $I = \langle \mathcal{A}, \mathcal{O}, W \rangle$ be an instance and $\overrightarrow{\pi}$ be an allocation of this instance. We can decide in time $O(N \times M^2)$ if $\overrightarrow{\pi}$ is sequenceable.*

The proof of this proposition is based on the execution of a similar algorithm to the one which is used by Brams and King [10] in the proof of their Proposition 1 (necessity). However, our algorithm is more general because (i) it can involve non-strict preferences on objects, and (ii) it can conclude with non-sequenceability.

Proof. We show that Algorithm 2 returns a sequence $\overrightarrow{\sigma}$ generating the input allocation $\overrightarrow{\pi}$ if and only if there is one. Suppose that the algorithm returns a sequence $\overrightarrow{\sigma}$. Then, by definition of the sequence (in the loop from line 2 to line 7), at each step t, $i = \sigma_t$ can choose an object in π_i, that is one of her top objects. Conversely, suppose the algorithm returns *Non-sequenceable*. Then, at a given step t, $\forall i$, $\mathrm{best}(\mathcal{O}', i) \cap \pi_i = \emptyset$. By definition, $\overrightarrow{\pi}^{|\mathcal{O}'}$ is therefore, at this step, a frustrating sub-allocation of $\overrightarrow{\pi}$. By Proposition 1, $\overrightarrow{\pi}$ is thus non-sequenceable. The loop from line 2 to line 7 runs in time $O(N \times M)$, because searching for the top objects in the preferences of each agent can be made in time $O(M)$. This loop being executed M times, the algorithm globally runs in time $O(N \times M^2)$. $\qquad\square$

Algorithm 2: Sequencing of an allocation

    Input: an instance $I = \langle \mathcal{A}, \mathcal{O}, W \rangle$ and an allocation $\overrightarrow{\pi} \in \mathcal{P}(I)$
    Output: a sequence $\overrightarrow{\sigma}$ generating $\overrightarrow{\pi}$ or *Non-sequenceable*
1 $(\overrightarrow{\sigma}, \mathcal{O}') \leftarrow (\langle\rangle, \mathcal{O})$;
2   **for** t **from** *1* **to** M **do**
3   **if** $\exists i$ *such that* $\mathrm{best}(\mathcal{O}', i) \cap \pi_i \neq \emptyset$ **then**
4 $\overrightarrow{\sigma} \leftarrow \overrightarrow{\sigma} \cdot i$;
5   let $\ell \in \mathrm{best}(\mathcal{O}', i) \cap \pi_i$;
6 $\mathcal{O}' \leftarrow \mathcal{O}' \setminus \{\ell\}$;
7   **else return** *Non-sequenceable* ;
8   **return** $\overrightarrow{\sigma}$;

3.2   Strict preferences on objects

We now characterize the instances for which the relation $s(I)$ is a mapping.

Proposition 3. *Preferences are strict on objects iff $s(I)$ is a mapping from $\mathcal{S}(I)$ to $\mathcal{P}(I)$.*

Proof. If preferences are strict on objects, then each agent has only one possible choice at line 5 of Algorithm 1 and hence every sequence generates one and only one allocation.

Conversely, if preferences are not strict on objects, at least one agent (suppose w.l.o.g. agent 1) gives the same weight to two different objects. Suppose that there is at least t

objects ranked above. Then obviously, the following sequence $\underbrace{111...111}_{t+1 \text{ times}} \underbrace{222...222}_{M-t-1 \text{ times}}$ generates two allocations, depending on agent 1's choice at step $t+1$. $\qquad\square$

3.3 Same order preferences

Proposition 4. *All the allocations of an instance with same order preferences are sequenceable. Conversely, if all the allocations of an instance are sequenceable, then this instance has same order preferences.*

Proof. Let I be an instance with same order preferences, and let $\overrightarrow{\pi}$ be an arbitrary allocation. In every sub-allocation of $\overrightarrow{\pi}$ at least one agent obtains a top object (because the preference order is the same among agents) and hence cannot be frustrating. By Proposition 1, $\overrightarrow{\pi}$ is sequenceable.

Conversely, let us assume for contradiction that I is an instance not having same order preferences. Then there are two distinct objects ℓ and m and two distinct agents i and j such that $W(i,\ell) \geq W(j,\ell)$ and $W(i,m) \leq W(j,m)$, one of the two inequalities being strict (assume w.l.o.g. the first one). The sub-allocation $\overrightarrow{\pi}^{|\{\ell,m\}}$ such that $\pi_i^{|\{\ell,m\}} = \{m\}$ and $\pi_j^{|\{\ell,m\}} = \{\ell\}$ is frustrating. By Proposition 1, every allocation $\overrightarrow{\pi}$ containing this frustrating sub-allocation (hence such that $m \in \pi_i$ and $\ell \in \pi_j$) is non-sequenceable. $\qquad\square$

Let us now characterize the instances for which $s(I)$ is a one-to-one correspondence.

Proposition 5. *For a given instance, the following two statements are equivalent.*

(A) Preferences are strict on objects and in the same order.

(B) The relation $s(I)$ is a one-to-one correspondence.

The proof of this proposition is an easy consequence of Propositions 3 and 4.

4 Pareto-optimality

An allocation is Pareto-optimal if there is no other allocation dominating it. In our context, allocation $\overrightarrow{\pi}'$ dominates allocation $\overrightarrow{\pi}$ if for all agent i, $u_i(\pi_i') \geq u_i(\pi_i)$ and $u_j(\pi_j') > u_j(\pi_j)$ for at least one agent j. Pareto-optimality formalizes the idea of efficiency: when an allocation is Pareto-optimal, one cannot strictly increase the utility of one agent without strictly decreasing another one's.

When an allocation is generated from a sequence, in some sense, a weak form of efficiency is applied to build the allocation: each successive (picking) choice is "locally" optimal. This raises a natural question: is every sequenceable allocation Pareto-optimal?

Brams *et al.* [10, Proposition 1] answer positively by proving the equivalence between sequenceability and Pareto-optimality. However, they have a different notion of Pareto-optimality, because they only have partial ordinal information about the agents' preferences. More precisely, in Brams and King's model, the agents' preferences are given as linear orders over *objects* (*e.g.* $1 \succ 2 \succ 3 \succ 4$). To be able to compare bundles, these preferences are lifted on subsets using the *responsive set extension* $\succ_{RS}$, which is similar to the one defined in the work by Aziz *et al.* [2], and to the one defined by Bouveret *et al.* [6] for SCI-nets.[4] This extension leaves many bundles incomparable (*e.g.* 14 and 23 if we consider the order $1 \succ 2 \succ 3 \succ 4$). It leads Bouveret *et al.* [6] to define, among others, two modes of

---

[4]It is actually not completely clear in Brams and King's paper whether or not their notion of dominance extends to bundles of different sizes, but it seems to be implicitly the case, using monotonicity.

Pareto-optimality: *possible* and *necessary* Pareto-optimality. Brams and King's notion of Pareto-optimality exactly corresponds to possible Pareto-optimality.

Aziz *et al.* [2] show that, given a linear order $\succ$ on objects and two bundles π and π', $\pi \succ_{RS} \pi'$ if and only if $u(\pi) > u(\pi')$ for *all* additive utility functions u compatible with $\succ$ (that is, such that $u(\ell) > u(m)$ if and only if $\ell \succ m$). This characterization of responsive dominance yields the following reinterpretation of Brams and King's result:

Proposition 6 (Brams and King [10]). *Let $\langle \succ_1, \ldots, \succ_N \rangle$ be the profile of agents' ordinal preferences (represented as linear orders). Allocation $\overrightarrow{\pi}$ is sequenceable if and only if for each other allocation $\overrightarrow{\pi}'$, there is a sequence $u_1, \ldots, u_N$ of additive utility functions, respectively compatible with $\succ_1, \ldots, \succ_N$ such that $u_i(\overrightarrow{\pi}) > u_i(\overrightarrow{\pi}')$ for at least one agent i.*

The latter notion of Pareto-optimality is very weak, because the additive utility function is not fixed — we just have to find one that works. In our context where the utility function is fixed and hence leads to a much stronger notion of Pareto-optimality,[5] there is no reason to suppose that Pareto-optimality is equivalent to sequenceability anymore. And it turns out that it is indeed not the case, as the following example shows.[6]

Example 4. *Let us consider the following instance:*

$$\begin{pmatrix} 5 & 4 & 2 \\ 8 & 2 & 1 \end{pmatrix}$$

The sequence $\langle 1, 2, 2 \rangle$ generates allocation $A = \langle 1, 23 \rangle$ giving utilities $\langle 5, 3 \rangle$. A is dominated by $B = \langle 23, 1 \rangle$, giving utilities $\langle 6, 8 \rangle$ (and generated by $\langle 2, 1, 1 \rangle$). Observe that, under ordinal linear preferences, B would not dominate A, but they would be incomparable.

The last example shows that a sequence of sincere choices does not necessarily generate a Pareto-optimal allocation (even when the preferences on objects are in the same order, as the example shows). What about the converse? We can see, as a trivial corollary of the latter reinterpretation of Brams and King's result, that the answer is positive *if the preferences are strict on shares*. The following result is more general:

Proposition 7. *Every Pareto-optimal allocation is sequenceable.*

Before giving the formal proof, we illustrate it on a concrete example [8, Example 5].

Example 5. *Let us consider the following instance:*

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

The circled allocation $\overrightarrow{\pi}$ is not sequenceable: indeed, every sequence that could generate it should start with $\langle 3, 2, \ldots \rangle$, leaving the frustrating sub-allocation $\overrightarrow{\rho}$ in a dotted box.

Let us now choose an arbitrary agent who does not receive a top object in $\overrightarrow{\rho}$, for instance agent $a_1 = 2$. Let $o_1 = 3$ be her top object (of weight 11 in this case). The agent receiving o_1 in $\overrightarrow{\rho}$ is $a_2 = 1$. This agent prefers object $o_2 = 4$ (of weight 15), held by a_1, already encountered. We have built a cycle $(a_1, o_1) \to (a_2, o_2) \to (a_1, o_1)$, in other words $(2, 3) \to (1, 4) \to (2, 3)$, that tells us exactly how to build another sub-allocation dominating $\overrightarrow{\rho}$. This sub-allocation can be built by replacing in $\overrightarrow{\rho}$ the attributions $(a_1 \leftarrow o_2)(a_2 \leftarrow o_1)$ by the

---

[5]Interestingly, in a very recent paper, Aziz *et al.* [1] have a similar discussion about Pareto-optimality.

[6]Actually, it was already known from a result by de Keijzer *et al.* [17] and Aziz *et al.* [1] that testing Pareto-optimality with additive preferences in coNP-complete. Since we know from Proposition 2 that testing sequenceability is polynomial, Pareto-optimality cannot be equivalent to it unless P = coNP.

attributions $(a_1 \leftarrow o_1)(a_2 \leftarrow o_2)$. Hence, each agent involved in the cycle obtains a strictly better object than the previous one. Doing the same substitutions in the initial allocation $\overrightarrow{\pi}$ yields an allocation $\overrightarrow{\pi}'$ that dominates $\overrightarrow{\pi}$ (marked with † in the matrix W above).

Now we will give the formal proof.[7]

Proof. As stated in the example, we will now prove the contraposition of the proposition: every non-sequenceable allocation is dominated. Let $\overrightarrow{\pi}$ be a non-sequenceable allocation. From Proposition 1, in a non-sequenceable allocation, there is at least one frustratring sub-allocation. Let $\overrightarrow{\rho}$ be such a sub-allocation (that can be $\overrightarrow{\pi}$ itself). We will, from $\overrightarrow{\rho}$, build another sub-allocation dominating it. Let us choose an arbitrary agent a_1 involved in $\overrightarrow{\rho}$, receiving an object not among her top ones in $\overrightarrow{\rho}$. Let o_1 be a top object of a_1 in $\overrightarrow{\rho}$, and let a_2 ($\neq a_1$) be the unique agent receiving it in $\overrightarrow{\rho}$. Let o_2 be a top object of a_2. We can notice that $o_2 \neq o_1$ (otherwise a_2 would obtain one of her top objects and $\overrightarrow{\rho}$ would not be frustrating). Let a_3 be the unique agent receiving o_2 in $\overrightarrow{\rho}$, and so on. Using this argument iteratively, we form a path starting from a_1 and alternating agents and objects, in which two successive agents and objects are distinct. Since the number of agents and objects is finite, we will eventually encounter an agent which has been encountered at a previous step of the path. Let a_i be the first such agent and o_k be the last object seen before her in the sequence (a_i is the unique agent receiving o_k). We have built a cycle $(a_i, o_i) \rightarrow (a_{i+1}, o_{i+1}) \cdots (a_k, o_k) \rightarrow (a_i, o_i)$ in which all the agents and objects are distinct, and that has at least two agents and two objects. From this cycle, we can modify $\overrightarrow{\rho}$ to build a new sub-allocation by giving to each agent in the cycle a top object instead of another less preferred object, all the agents not appearing in the cycle being left unchanged. More formally, the following attributions in $\overrightarrow{\rho}$ (and hence in $\overrightarrow{\pi}$): $(a_i \leftarrow o_k)(a_{i+1} \leftarrow o_i) \cdots (a_k \leftarrow o_{k-1})$ are replaced by: $(a_i \leftarrow o_i)(a_{i+1} \leftarrow o_{i+1}) \cdots (a_k \leftarrow o_k)$ where $(a \leftarrow o)$ means that o is attributed to a. The same substitutions operated in $\overrightarrow{\pi}$ yield an allocation $\overrightarrow{\pi}'$ that dominates $\overrightarrow{\pi}$. □

Corollary 1. *No frustrating allocation can be Pareto-optimal (equivalently, in every Pareto-optimal allocation, at least one agent receives a top object).*

Proposition 7 implies that there exists, for a given instance, three classes of allocations: (1) non-sequenceable (therefore non Pareto-optimal) allocations, (2) sequenceable but non Pareto-optimal allocations, and (3) Pareto-optimal (hence sequenceable) allocations. These three classes define a "scale of efficiency" that can be used to characterize the allocations. What is interesting and new here is the intermediate level.

5   Envy-Freeness and CEEI

In the previous section we have investigated the link existing between efficiency and sequenceability. The use of sequence of sincere choices can also be motivated by the search for a *fair* allocation protocol. Hence, investigating the link between sequenceability and fairness is a natural next step. In this section, we will focus on two fairness properties, envy-freeness and competitive equilibrium from equal income and analyze their link with sequenceability.

Envy-freeness [19, 30] is probably one of the most prominent fairness properties. It can be formally expressed in our model as follows.

---

[7]This proof is similar to Brams and King's [10, Proposition 1, necessity]. However we give it entirely because it is more general and will be reused in Proposition 9. The central idea of trading cycle is classical and is used, *e.g.*, by Varian [30, page 79] and Lipton *et al.* [26, Lemma 2.2] in the context of envy-freeness.

Definition 5. *Let I be an add-MARA instance and $\overrightarrow{\pi}$ be an allocation. $\overrightarrow{\pi}$ verifies the envy-freeness property (or is simply envy-free), when $u_i(\pi_i) \geq u_i(\pi_j)$, $\forall(i,j) \in \mathcal{A}^2$ (no agent strictly prefers the share of any other agent).*

The notion of *competitive equilibrium* is an old and well-known concept in economics [31, 18]. If equal incomes are imposed among the stakeholders, this concept becomes the *competitive equilibrium from equal incomes* [27, page 177, for example], yielding a very strong fairness concept which has been recently explored in artificial intelligence [28, 14, 8]. Here is the definition of this concept adapted to our model.

Definition 6. *Let $I = (\mathcal{A}, \mathcal{O}, W)$ be an add-MARA instance, $\overrightarrow{\pi}$ an allocation, and $\overrightarrow{p} \in [0,1]^M$ a vector of prices. A pair $(\overrightarrow{\pi}, \overrightarrow{p})$ is said to form a competitive equilibrium from equal incomes (CEEI) if*

$$\forall i \in \mathcal{A} : \pi_i \in \operatorname{argmax}_{\pi \subseteq \mathcal{O}} \left\{ u_i(\pi) : \sum_{\ell \in \pi} p_\ell \leq 1 \right\}.$$

In other words, π_i is one of the maximal shares that i can buy with a budget of 1, provided that the price of each object ℓ is p_ℓ.

We will say that allocation $\overrightarrow{\pi}$ satisfies the CEEI test (is a CEEI allocation for short) if there exists a vector $\overrightarrow{p}$ such that $(\overrightarrow{\pi}, \overrightarrow{p})$ forms a CEEI.

As we have shown in a previous work [8], every CEEI allocation is envy-free in the model we use (additive, numerical preferences). In this section, we investigate the question of whether an envy-free or CEEI allocation is necessarily sequenceable.

For envy-freeness, the answer is negative.

Proposition 8. *There exist non-sequenceable envy-free allocations, even if the agents' preferences are strict on shares.*

Proof. A counterexample with strict preferences on shares is given in Example 5 above, for which we can check that the circled allocation $\overrightarrow{\pi}$ is envy-free and non-sequenceable. $\square$

Interestingly, for CEEI, however, the answer is positive.

Proposition 9. *Every CEEI allocation is sequenceable.*

It should be noted that we already know that every CEEI allocation is Pareto-optimal if the preferences are strict on shares [8, Proposition 8]. From Observation 1 and Proposition 7, if the preferences are strict on shares, then every CEEI allocation is sequenceable. Proposition 9 is more general: no assumption is made on the strict order on shares (nor on objects). Note that a CEEI allocation can be ordinally necessary Pareto-dominated, as the following example shows.

$$\begin{pmatrix} \dagger② & \dagger 3 & 3 & ② \\ 2 & 3 & \dagger④ & 1 \\ 0 & ④ & 2 & \dagger 4 \end{pmatrix}$$

The circled allocation is CEEI (with prices 0.5, 1, 1, 0.5) but is ordinally necessary (hence also additively) dominated by the allocation marked with $\dagger$.

Proof. We will show that no allocation can be at the same time non-sequenceable and CEEI. Let $\overrightarrow{\pi}$ be a non-sequenceable allocation. We can use the same terms and notations than in the proof of Proposition 7, especially concerning the dominance cycle.

Let $\mathcal{C}$ be the set of agents concerned by the cycle. $\overrightarrow{\pi}$ contains the following shares:

$$\pi_{a_i} = \{o_k\} \cup \tau_i \quad \pi_{a_{i+1}} = \{o_i\} \cup \tau_{i+1} \quad .... \quad \pi_{a_k} = \{o_{k-1}\} \cup \tau_k$$

whereas the allocation $\overrightarrow{\pi}'$ that dominates it, contains the following shares:

$$\pi'_{a_i} = \{o_i\} \cup \tau_i \quad \pi'_{a_{i+1}} = \{o_{i+1}\} \cup \tau_{i+1} \quad .... \quad \pi'_{a_k} = \{o_k\} \cup \tau_k$$

the other shares being unchanged from $\overrightarrow{\pi}$ to $\overrightarrow{\pi}'$.

Suppose that $\overrightarrow{\pi}$ is CEEI. This allocation must satisfy two kinds of constraints. First, $\overrightarrow{\pi}$ must satisfy the price constraint. If $p(\pi) \overset{def}{=} \sum_{\ell \in \pi} p_\ell$, we have, $\forall i \in \mathcal{C}, p(\pi_i) \le 1$ (1).

Next, $\overrightarrow{\pi}$ must be optimal: every share having a higher utility for an agent than her share in $\overrightarrow{\pi}$ costs strictly more than 1. Provided that $\forall i \in \mathcal{C} : u_i(\pi'_i) > u_i(\pi_i)$ (because $\overrightarrow{\pi}'$ substitutes more preferred objects to less preferred objects in $\overrightarrow{\pi}$), this constraint can be written as $\forall i \in \mathcal{C}, p(\pi'_i) > 1$ (2).

By summing equations (1) and (2), provided that all shares are disjoint, we obtain

$$p(\cup_{j \in \mathcal{C}} \pi_j) \le |\mathcal{C}| \quad \text{and} \quad p(\cup_{j \in \mathcal{C}} \pi'_j) > |\mathcal{C}|$$

Yet, $\cup_{j \in \mathcal{C}} \pi_j = \cup_{i \in \mathcal{C}} \pi'_j$ (because the allocation $\overrightarrow{\pi}'$ is obtained from $\overrightarrow{\pi}$ by simply swapping objects between agents in $\mathcal{C}$). The two previous equations are contradictory. □

6 Experiments

We have exhibited in Section 4, a "scale of allocation efficiency", made of three steps: non-sequenceable (NS), sequenceable and non Pareto-optimal (SnP), and Pareto-optimal (PO). A natural question is to know, for a given instance, which proportion of allocations are located at each level of the scale. We give a first answer in this section by experimentally characterizing the distribution of allocations between the different levels. Moreover, we analyze the relation between fairness and efficiency by linking this scale of efficiency with the scale of fairness introduced in a previous work [8]. This scale of fairness has six levels, from the weakest to the strongest property: no criterion satisfied (–), maxmin share (MFS), proportionality (PFS), minmax share (mFS), envy-freeness (EF) and CEEI.

Our experimental protocol is the following. We have generated 100 add-MARA instances with 3 agents and 10 objects, reusing the uniform and Gaussian random allocation protocol described in our previous work [8].[8] For each of these instances, we have generated the set of possible allocations ($3^{10} = 59049$) and identified the highest level of fairness and the level of efficiency of each one (using a brute-force approach with easy cuts for the NP-complete Pareto-efficiency test). We want to emphasize that a particular difficulty was here to implement the CEEI test. This problem has been proved to be coNP-hard [13, Theorem 49] and to the best of our knowledge no practical method have been described yet, even for small numbers of objects and agents. The method we have developed for this problem is described in appendix and works in practice for problems up to 5 agents and 20 objects.

Figure 1 gives a graphical overview of the results. In this figure, the histograms represent the average value (on all the instances) of the set of allocations by pair of criteria (the min-max interval has been represented with error bars), using a logarithmic scale. The three curves represent the proportion, by fairness criterion, of the number of allocations satisfying the three efficiency criteria, using a linear scale.

We can observe several interesting facts. First, a huge majority of allocations do not have any efficiency nor fairness property (first black bar on the left). Secondly, the distribution of the allocations on the scale of efficiency seems to be correlated to the fairness criteria: a higher proportion of sequenceable allocations can be found among the envy-free allocations

---

[8]Note that the generated utilities are normalized, but this has no impact on the results since all the criteria tested are independent of the utility scales.
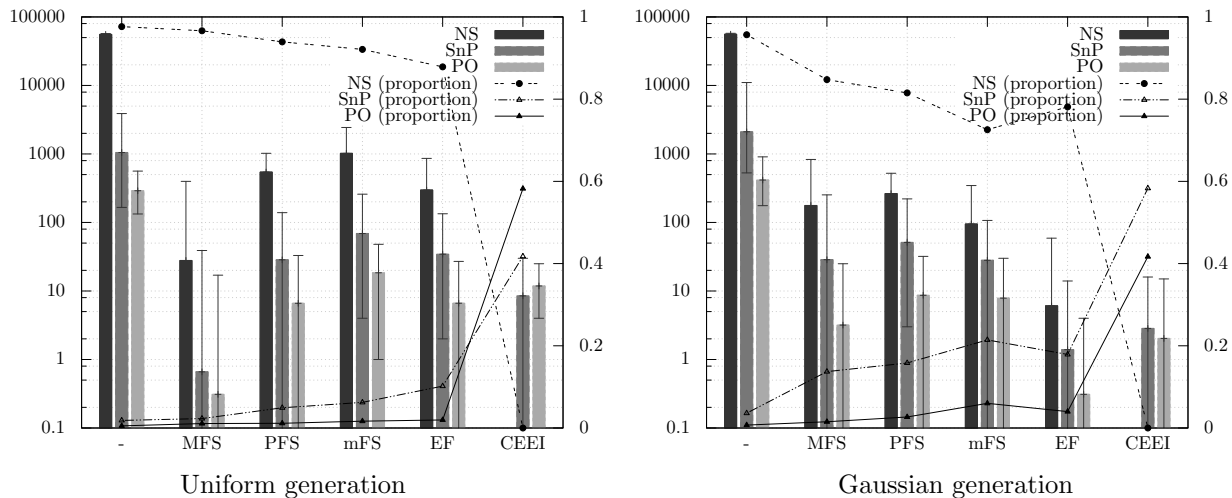
Figure 1: Distribution of the number of allocations by pair of (efficiency, fairness) criteria.

than among the allocations that do not satisfy any fairness property, and for CEEI allocations and uniform generation, there are even more Pareto-optimal allocations that just sequenceable ones. Third, there is a higher proportion of allocations satisfying fairness and efficiency properties for the uniform generation model. Even if this seems logical for fairness (because a "similar" attraction for the objects — as in the Gaussian model — is more likely to create conflicts among agents), the explanation is not so clear for efficiency.

7 Conclusion

Sequences of sincere choices are arguably a remarkable protocol for allocating a set of indivisible goods to agents. This protocol, known for ages, has received a lot of interest in recent years by researchers both in economics and computer science. In this paper, following the work by Brams *et al.* [10], we have shown that this protocol, beyond being appealing in practice, also has a theoretical interest in the context of numerical additive preferences. Namely, it can be used to characterize the efficiency of an allocation by defining an intermediate level between Pareto-optimality and no efficiency at all. Moreover, we have introduced the simple notion of frustrating (sub-)allocation and shown that it can be used to exactly characterize the set of sequenceable allocations. We also have characterized the set of instances for which there is an exact one-to-one relation between allocations and sequences. Finally, we have emphasized some links between fairness properties (especially CEEI) and efficiency criteria. Although being technically simple, we believe that these results are new and shed an interesting light on sequences of sincere choices and Pareto-optimality.

This work opens up to some interesting questions, such as the impact of restrictions on sequences like alternating sequences. Another interesting topic is the relation between sequences and social welfare orderings, with questions such as the loss of social welfare incurred by the execution of a sequence compared to the optimal allocation (price of sequenceability).

References

[1] Haris Aziz, Peter Biro, Jérôme Lang, Julien Lesca, and Jérôme Monnot. Optimal reallocation under additive and ordinal preferences. In *Proceedings of AAMAS'2016*, 2016.

[2] Haris Aziz, Serge Gaspers, Simon Mackenzie, and Toby Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.

[3] Haris Aziz, Toby Walsh, and Lirong Xia. Possible and necessary allocations via sequential mechanisms. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*, pages 468–474. AAAI Press, 2015.

[4] Nikhil Bansal and Maxim Sviridenko. The santa claus problem. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, STOC '06, pages 31–40, New York, NY, USA, 2006. ACM.

[5] Nicola Bianchessi, Jean-François Cordeau, Jacques Desrosiers, Gilbert Laporte, and Vincent Raymond. A heuristic for the multi-satellite, multi-orbit and multi-user management of earth observation satellites. *European Journal of Operational Research*, 177(2):750–762, March 2007.

[6] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In Helder Coelho, Rudi Studer, and Michael Wooldridge, editors, *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI-10)*, pages 387–392, Lisbon, Portugal, August 2010. IOS Press.

[7] Sylvain Bouveret and Jérôme Lang. A general elicitation-free protocol for allocating indivisible goods. In Toby Walsh, editor, *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 73–78, Barcelona, Spain, July 2011. IJCAI/AAAI.

[8] Sylvain Bouveret and Michel Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290, March 2016.

[9] Steven J. Brams, Marc D. Kilgour, and Christian Klamler. The undercut procedure: an algorithm for the envy-free division of indivisible items. *Social Choice and Welfare*, 39(2-3):615–631, 2012.

[10] Steven J. Brams and Daniel King. Efficient fair division—help the worst off or avoid envy? *Rationality and Society*, 17(4):387–421, 2005.

[11] Steven J. Brams and Alan D. Taylor. *Fair Division — From Cake-cutting to Dispute Resolution*. Cambridge University Press, 1996.

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

[13] Simina Brânzei. *Computational Fair Division*. PhD thesis, Department of Computer Science, Aarhus Universitet, Denmark, February 2015.

[14] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, dec 2011.

[15] Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. On maximal classes of utility functions for efficient one-to-one negociation. In Leslie Pack Kaelbling and Alessandro Saffiotti, editors, *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI-05)*, Edinburgh, Scotland, August 2005. Professional Book Center.

[16] George B. Dantzig and Curtis B. Eaves. Fourier-Motzkin Elimination and its Dual. *Journal of Combinatorial Theory*, A(14):288–297, 1973.

[17] Bart de Keijzer, Sylvain Bouveret, Tomas Klos, and Yingqian Zhang. On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT'09)*, Lecture Notes in Artificial Intelligence, pages 98–110, Venice, Italy, October 2009. Springer Verlag.

[18] Irving Fisher. *Mathematical Investigations in the Theory of Value and Prices, and Appreciation and Interest.* Augustus M. Kelley, Publishers, 1892.

[19] Duncan K. Foley. Resource allocation and the public sector. *Yale Economic Essays*, 7(1):45–98, 1967.

[20] Jean-Baptiste Joseph Fourier. Solution d'une question particuliere du calcul des inégalités. *Nouveau Bulletin des Sciences par la Société philomatique de Paris*, 99:100, 1826.

[21] Judy Goldsmith and Robert H Sloan. The AI conference paper assignment problem. In *Proc. AAAI Workshop on Preference Handling for Artificial Intelligence, Vancouver*, pages 53–57, 2007.

[22] Thomas Kalinowski, Nina Narodytska, and Toby Walsh. A social welfare optimal sequential allocation procedure. In Francesca Rossi, editor, *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-13)*, Beijing, China, August 2013. IJCAI/AAAI.

[23] Thomas Kalinowski, Nina Narodytska, Toby Walsh, and Lirong Xia. Strategic behavior when allocating indivisible goods sequentially. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI-12)*, Bellevue, WA, July 2013. AAAI Press.

[24] David A Kohler and R Chandrasekaran. A class of sequential games. *Operations Research*, 19(2):270–277, 1971.

[25] Michel Lemaître, Gérard Verfaillie, and Nicolas Bataille. Exploiting a common property resource under a fairness constraint: a case study. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 206–211, Stockholm, Sweden, July 1999.

[26] Richard Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of divisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC-04)*, pages 125–131, New York, NY, May 2004. ACM.

[27] Hervé Moulin. *Fair Division and Collective Welfare.* MIT Press, 2003.

[28] Abraham Othman, Tuomas Sandholm, and Eric Budish. Finding approximate competitive equilibria: efficient and fair course allocation. In Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen, editors, *Proceedings of the 9th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-10)*, pages 873–880, Toronto, Canada, May 2010. IFAAMAS.

[29] Tuomas W. Sandholm. Contract types for satisficing task allocation: I. theoretical results. In Sandip Sen, editor, *Proceedings of the AAAI Spring Symposium: Satisficing Models*, pages 68–75, Menlo Park, California, 1998. AAAI Press.

[30] Hal R. Varian. Equity, Envy and Efficiency. *Journal of Economic Theory*, 9:63–91, 1974.

[31] Léon Walras. *Éléments d'économie politique pure ou Théorie de la richesse sociale*. L. Corbaz, 1st edition, 1874.

Sylvain Bouveret
Laboratoire d'Informatique de Grenoble (LIG)
Grenoble INP, Universit Grenoble-Alpes
Grenoble, France
Email: sylvain.bouveret@imag.fr

Michel Lemaître
Formerly ONERA Toulouse, France
Email: michel.lemaitre.31@gmail.com

Appendix

This appendix presents an exact (complete) and practical method to decide whether a given allocation satisfies the CEEI test (Definition 6). We know that this problem is coNP-hard [13, Theorem 49]. The method we propose relies on linear programming techniques. Our ambition is limited to be able to tackle small instances (say no more than 5 agents and 20 objects) with standard LP solvers.

The starting point is the following. The problem comes down to decide if the following system of constraints is satisfiable. This system, denoted by S, depends on the allocation $\overrightarrow{\pi}$ at stake. Its variables are the object prices $(p_\ell)_{\ell=1}^M$ and it comprises three sets of constraints:

- a first set of inequalities defining the domain of prices:

$$0 \leq p_\ell \leq 1, \text{ for all } \ell \in [\![1, M]\!];$$

- a second set of N non strict inequalities, modelling the fact that each agent's share is affordable:

$$\sum_{\ell=1}^M A_{i\ell} p_\ell \leq 1, \text{ for all } i \in [\![1, N]\!],$$

with $A_{i\ell} = 1$ if $\ell \in \pi_i$, 0 otherwise;

- a third set of q strict inequalities expressing the optimality of $\overrightarrow{\pi}$ given the prices: any better share for an agent costs strictly more than the given budget. Namely: for each agent i and each share π' such that $u_i(\pi') > u_i(\pi_i)$,

$$\sum_{\ell=1}^M B_{k\ell} p_\ell > 1,$$

where k is an index, different for each pair (i, π'), and $B_{k\ell} = 1$ if $\ell \in \pi'$, 0 otherwise.

Note that these constraints cannot be properly handled by a (standard) LP tool, because the inequalities of the third set are strict. The Fourier-Motzkin elimination method [20, 16] could do the job, as it can manage strict inequality constraints. However, its drawback is its exponential number of steps.

Another possibility is to build, from S, a new system S' of non-strict linear inequalities which is equivalent to S as far as satisfiability is concerned. This system uses the same coefficients $A_{i\ell}$ and $B_{k\ell}$, and is defined as follows.

- replace in each inequality each variable p_ℓ by a new variable p'_ℓ, with domains

$$0 \leq p'_\ell \text{ for all } \ell \in [\![1, M]\!];$$

- replace in the second set of non strict inequalities, the budget bound 1 by a new variable d:

$$\sum_{\ell=1}^M A_{i\ell} p'_\ell \leq d, \text{ for all } i \in [\![1, N]\!];$$

- replace the third set of strict inequalities by the new following set of non-strict ones:

$$\sum_{\ell=1}^M B_{k\ell} p'_\ell \geq d + 1, \text{ for all } k \in [\![1, q]\!],$$

Proposition 10. *S is satisfiable if and only if S' is satisfiable. Consequently, an allocation satisfies the CEEI test if and only if the corresponding system S' is satisfiable. Moreover, in this case, giving to each object ℓ the price $\dfrac{\bar{p'_\ell}}{\bar{d}}$, with $\bar{p'_\ell}$ and $\bar{d}$ the values of p'_ℓ and d in the solution of S' yields a CEEI with respect to the initial allocation $\overrightarrow{\pi}$.*

Proof. ($\Longrightarrow$) Suppose S is satisfiable. Then there is a rational solution.[9] Let $(\bar{p}_\ell)_{\ell=1}^M$ be this solution values, and let $\bar{d}$ be the least common multiple of the $\bar{p}_\ell$ denominators. On the one hand we have:

$$\sum_{\ell=1}^M A_{i\ell}\bar{p}_\ell \leq 1 \implies \sum_{\ell=1}^M A_{i\ell}\bar{d}\bar{p}_\ell \leq \bar{d}$$

and on the other hand:

$$\sum_{\ell=1}^M B_{k\ell}\bar{p}_\ell > 1 \implies \sum_{\ell=1}^M B_{k\ell}\bar{d}\bar{p}_\ell > \bar{d}$$

$$\implies \sum_{\ell=1}^M B_{k\ell}\bar{d}\bar{p}_\ell \geq \bar{d}+1,$$

because $\displaystyle\sum_{\ell=1}^M B_{k\ell}\bar{d}\bar{p}_\ell$ and $\bar{d}$ are integers.

Hence S' is satisfiable if S is, with values $\bar{d}\bar{p}_\ell$ for p'_ℓ and $\bar{d}$ for d.

($\Longleftarrow$) Suppose S' satisfiable, and let $\bar{p'}_\ell$ and $\bar{d}$ be the values of a solution. Then, on the one hand we have:

$$\sum_{\ell=1}^M A_{i\ell}\bar{p'}_\ell \leq \bar{d} \implies \sum_{\ell=1}^M A_{i\ell}\frac{\bar{p'}_\ell}{\bar{d}} \leq 1$$

and on the other hand:

$$\sum_{\ell=1}^M B_{k\ell}\bar{p'}_\ell \geq \bar{d}+1 \implies \sum_{\ell=1}^M B_{k\ell}\bar{p'}_\ell > \bar{d}$$

$$\implies \sum_{\ell=1}^M B_{k\ell}\frac{\bar{p'}_\ell}{\bar{d}} > 1$$

Hence S is satisfiable if S' is, with values $\dfrac{\bar{p'}_\ell}{\bar{d}}$. $\qquad\square$

The system S' only has non-strict linear inequalities and a standard LP tool can be used to test its satisfiability (there is no need for integer linear programming). Note that it does not contradict the coNP-hardness of the CEEI test since the number of constraints is not polynomially bounded.

We now conclude with some practical remarks. It is known [8, Proposition 7] that any allocation satisfying the CEEI test is envy-free. Moreover, we know from Proposition 9 that every CEEI allocation is sequenceable. Since both properties can be tested in polynomial

---

[9]Because coefficients are rational. It is a known result, which derives for example from the Fourier-Motzkin elimination method: the only required operations to build the solution (if it exists) are the standard arithmetic operators.

time (envy-freeness can be tested $O(N^2M)$, sequenceability in $O(NM^2)$), we can use them in a preliminary step to filter out allocations that do not pass them. Finally, it can be noticed that the third set of inequalities often contains a lot of redundant constraints that can be eliminated to simplify the overall system.