

Efficiency, Sequenceability and Deal-Optimality in Fair Division of Indivisible Goods

Aurélie Beynier
Sorbonne Université, CNRS, LIP6
F-75005 Paris, France
aurelie.beynier@lip6.fr

Sylvain Bouveret
Univ. Grenoble Alpes, CNRS, LIG
Grenoble, France
sylvain.bouveret@imag.fr

Michel Lemaître
Formerly ONERA
Toulouse, France
michel.lemaître.31@gmail.com

Nicolas Maudet
Sorbonne Université, CNRS, LIP6
F-75005 Paris, France
nicolas.maudet@lip6.fr

Simon Rey
Sorbonne Université, ENS P-S
Paris, Cachan, France
srey@ens-paris-saclay.fr

Parham Shams
Sorbonne Université, CNRS, LIP6
F-75005 Paris, France
parham.shams@lip6.fr

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 as follows: at each stage, a designated agent picks one object among those that remain. Another intuitive way to obtain an allocation is to give objects to agents in the first place, and to let agents exchange them as long as such “deals” are beneficial. This paper investigates these notions, when agents have additive preferences over objects, and unveils surprising connections between them, and with other efficiency and fairness notions. In particular, we show that an allocation is sequenceable if and only if it is optimal for a certain type of deals, namely cycle deals involving a single object. Furthermore, any Pareto-optimal allocation is sequenceable, but not the converse. Regarding fairness, we show that an allocation can be envy-free and non-sequenceable, but that every competitive equilibrium with equal incomes is sequenceable. To complete the picture, we show how some domain restrictions may affect the relations between these notions. Finally, we experimentally explore the links between the scales of efficiency and fairness.

KEYWORDS

Multiagent Resource Allocation; Fair Division; Efficiency

ACM Reference Format:

Aurélie Beynier, Sylvain Bouveret, Michel Lemaître, Nicolas Maudet, Simon Rey, and Parham Shams. 2019. Efficiency, Sequenceability and Deal-Optimality in Fair Division of Indivisible Goods. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019)*, Montreal, Canada, May 13–17, 2019, IFAAMAS, 9 pages.

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, taking into account the agents’ preferences about the objects. This classical collective decision making problem has plenty of practical applications, among which the allocation of space resources [9, 30], of tasks to workers in crowdsourcing market systems [34], papers to reviewers [26] or courses to students [19].

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. 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 [23, 35]. A somewhat intermediate approach consists in allocating the objects to the agents using a *protocol*, which allows to build 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) [16], but have also been studied in the context of indivisible goods [14, 16].

In this paper, we focus on a particular allocation protocol: *sequences of sincere choices* (also known as *picking sequences*). This very simple 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 $(1, 2, 2, 1)$, agent 1 will choose first, then agent 2 will pick two objects in a row, and agent 1 will take the last object. This protocol, used in a lot of everyday situations, has been studied for the first time by Kohler and Chandrasekaran [29]. Later, Brams and Taylor [17] have studied a particular version of this protocol, namely alternating sequences, in which the sequence of agents is restricted to a balanced $((1, 2, 2, 1...))$ or strict $((1, 2, 1, 2...))$ alternation of agents. Bouveret and Lang [11] have further formalized this protocol, whose properties (especially related to game theoretic aspects) have later been characterized by Kalinowski *et al.* [27, 28]. Finally, Aziz *et al.* [4] have studied the complexity of problems related to finding whether a particular assignment (or bundle) is achievable by a particular class of picking sequences. Picking sequences have also been considered by Brams and King [15], that focus on a situation where the agents have ordinal preferences. They 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 elaborate on these ideas and analyze the links between sequences, certain types of deals among agents, and some efficiency and fairness properties, in a more general model in which the agents have numerical additive preferences on the objects. 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 (Section 3). Then, we show that in this slightly more general framework than the one by Brams and King, Pareto-optimality and sequenceability are not equivalent anymore (Section 4). By unveiling the connection between sequenceability and cycle deals among agents (Section 5), we obtain a rich “scale of efficiency” that allows us to characterize the degree of efficiency of a given allocation. Interestingly, some domain restrictions have significant effects on this hierarchy (Section 6). We also highlight (Section 7) a 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 (Section 8).

2 MODEL AND DEFINITIONS

The aim of the fair division of indivisible goods, also called MultiAgent Resource Allocation (MARA), is to allocate a finite set of objects $O = \{o_1, \dots, o_m\}$ to a finite set of agents $N = \{1, \dots, n\}$. A *sub-allocation* on $O' \subseteq O$ is a vector $\vec{\pi}^{O'} = \langle \pi_1^{O'}, \dots, \pi_n^{O'} \rangle$ of *bundles* of objects, such that $\forall i, \forall j$ with $i \neq j : \pi_i^{O'} \cap \pi_j^{O'} = \emptyset$ (a given object cannot be allocated to more than one agent) and $\bigcup_{i \in N} \pi_i^{O'} = O'$ (all the objects from O' are allocated). $\pi_i^{O'} \subseteq O'$ is called agent i 's *share* on O' . $\vec{\pi}^{O''}$ is a sub-allocation of $\vec{\pi}^{O'}$ when $\pi_i^{O''} \subseteq \pi_i^{O'}$ for each agent i . Any sub-allocation $\vec{\pi}^{O'}$ on the entire set of objects will be written $\vec{\pi}$ and just called *allocation*.

Any satisfactory allocation must take into account the agents' preferences on the objects. Here, we will make the classical assumption that these preferences are *numerically additive*. Each agent i has a *utility function* $u_i : 2^O \rightarrow \mathbb{R}^+$ measuring her satisfaction $u_i(\pi)$ when she obtains share π , which is defined as follows:

$$u_i(\pi) \stackrel{\text{def}}{=} \sum_{o_k \in \pi} w(i, o_k),$$

where $w(i, o_k)$ is the weight given by agent i to object o_k . This assumption, as restrictive as it may seem, is made by a lot of authors [8, 31, for instance] and is considered a good compromise between expressivity and conciseness.

Definition 2.1. An instance of the *additive multiagent resource allocation problem* (add-MARA instance) is a tuple $I = \langle N, O, w \rangle$, where N and O are as defined above and $w : N \times O \rightarrow \mathbb{R}^+$ is a mapping. Here, $w(i, o_k)$ is the weight given by agent i to object o_k .

We say that the agents' preferences are *strict on objects* if, for each agent i and each pair of objects $o_k \neq o_l$, we have $w(i, o_k) \neq w(i, o_l)$. Similarly, we say that the agents' preferences are *strict on shares* if, for each agent i and each pair of shares $\pi \neq \pi'$, we have $u_i(\pi) \neq u_i(\pi')$. Note that strict preferences on shares entail strict preferences on objects; the converse is false.

We will denote by $\mathcal{P}(I)$ the set of allocations for I . We will also use the notation $o_i o_j \dots$ as a shorthand for bundle $\{o_i, o_j, \dots\}$.

Definition 2.2. Given an agent i and a set of objects O' , we will write $\text{best}(O', i) = \text{argmax}_{o_k \in O'} w(i, o_k)$ the objects from O' having the highest weight for i (they will be called *top* objects of i).

A (sub-)allocation $\vec{\pi}^{O'}$ is said *frustrating* if no agent receives one of her top objects in $\vec{\pi}^{O'}$ (formally: $\text{best}(O', i) \cap \pi_i^{O'} = \emptyset$ for each agent i), and *non-frustrating* otherwise.

In the following, we will consider a particular way of allocating objects to agents: sequences of sincere choices.

Definition 2.3. Let $I = \langle N, O, w \rangle$ be an add-MARA instance. A *sequence of sincere choices* (or simply *sequence* when the context is clear) is a vector of N^m . We will denote by $\mathcal{S}(I)$ the set of possible sequences for the instance I .

Let $\vec{\sigma} \in \mathcal{S}(I)$ be a sequence of agents and let σ_t be the t^{th} agent of the sequence. $\vec{\sigma}$ is said to *generate* allocation $\vec{\pi}$ if and only if $\vec{\pi}$ can be obtained as a possible result of the non-deterministic¹ Algorithm 1 on input I and $\vec{\sigma}$.

Algorithm 1: Execution of a sequence

Input: an instance $I = \langle N, O, w \rangle$ and a sequence $\vec{\sigma} \in \mathcal{S}(I)$

Output: an allocation $\vec{\pi} \in \mathcal{P}(I)$

- 1 $\vec{\pi} \leftarrow$ empty allocation (such that $\forall i \in N : \pi_i = \emptyset$);
 - 2 $O_1 \leftarrow O$;
 - 3 **for** t **from** 1 **to** m **do**
 - 4 $i \leftarrow \sigma_t$;
 - 5 **choose** object $o_t \in \text{best}(O_t, i)$;
 - 6 $\pi_i \leftarrow \pi_i \cup \{o_t\}$;
 - 7 $O_{t+1} \leftarrow O_t \setminus \{o_t\}$
-

Definition 2.4. An allocation $\vec{\pi}$ is said to be *sequenceable* if there exists a sequence $\vec{\sigma}$ that generates $\vec{\pi}$, and *non-sequenceable* otherwise. For a given instance I , we will write $s(I)$ the binary relation defined by $(\vec{\sigma}, \vec{\pi}) \in s(I)$ if and only if $\vec{\pi}$ can be generated by $\vec{\sigma}$.

Example 2.5. Let I be the instance represented by the following weight matrix:²

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

For instance, sequence $\langle 2, 1, 2 \rangle$ generates two possible allocations: $\langle o_1, o_2 o_3 \rangle$ and $\langle o_2, o_1 o_3 \rangle$, depending on whether agent 2 chooses object o_1 or o_3 that she both prefers. Allocation $\langle o_1 o_2, o_3 \rangle$ can be generated by three sequences. Allocations $\langle o_1 o_3, o_2 \rangle$ and $\langle o_3, o_1 o_2 \rangle$ are non-sequenceable.

For any instance I , $|\mathcal{S}(I)| = |\mathcal{P}(I)| = n^m$. Also note that the number of objects allocated to an agent by a sequence is the number of times the agent appears in the sequence.

The notion of frustrating allocation and sequenceability were already implicitly present in the work by Brams and King [15], and sequenceability has been extensively studied by Aziz *et al.* [4] with a focus on sub-classes of sequences (*e.g.* alternating sequences). However, a fundamental difference is that in our setting, the preferences

¹The algorithm contains an instruction **choose** splitting the control flow into several branches, building all the allocations generated by $\vec{\sigma}$.

²In this example and the following ones, we represent instances by a matrix where the value at row i and column o_k represents the weight $w(i, o_k)$.

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 2.5 shows.

3 SEQUENCEABLE ALLOCATIONS

We have seen in Example 2.5 that some allocations are non-sequenceable. We will now formalize this and give a precise 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. We first start by noticing that in every sequenceable allocation, the first agent of the sequence gets a top object, so every frustrating allocation is non-sequenceable. However, being non-frustrating is not a sufficient condition for an allocation to be sequenceable, as the following example shows:

Example 3.1. Consider this instance:

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

In allocation $\vec{\pi} = \langle o_1 o_4, o_2 o_3 \rangle$, each agent receives her top object. However, after o_1 and o_2 have been allocated (they must be allocated first by all sequences generating $\vec{\pi}$), the dotted sub-allocation remains. This sub-allocation is obviously non-sequenceable because it is frustrating. Hence $\vec{\pi}$ is not sequenceable either.

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

PROPOSITION 1. *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an instance and $\vec{\pi}$ be an allocation of this instance. The two following statements are equivalent:*

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

PROOF. (B) \Rightarrow (A). Let us suppose that for all subsets of objects $\mathcal{O}' \subseteq \mathcal{O}$, at least one agent gets one of her top objects in $\vec{\pi}|_{\mathcal{O}'}$. We will show that $\vec{\pi}$ is sequenceable. Let $\vec{\sigma}$ be a sequence of agents and $\vec{\mathcal{O}} \in (2^{\mathcal{O}})^m$ be a sequence of bundles jointly defined as follows:

- $\mathcal{O}_1 = \mathcal{O}$ and σ_1 is an agent that receives one of her top objects in $\vec{\pi}|_{\mathcal{O}_1}$;
- $\mathcal{O}_{t+1} = \mathcal{O}_t \setminus \{o_t\}$, 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 $\vec{\pi}|_{\mathcal{O}_t}$, for $t \geq 1$.

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

(A) \Rightarrow (B) by contraposition. Let $\vec{\pi}$ be an allocation containing a frustrating sub-allocation $\vec{\pi}|_{\mathcal{O}'}$. Suppose that there exists a sequence $\vec{\sigma}$ generating $\vec{\pi}$. We can notice that in a sequence of sincere choices, when an object is allocated to an agent, all the objects that are strictly better for her have already been allocated at a previous step. Let $o_k \in \mathcal{O}'$, and let i be the agent that receives o_k in $\vec{\pi}$. Since $\vec{\pi}|_{\mathcal{O}'}$ is frustrating, there is another object $o_l \in \mathcal{O}'$ such that $w(i, o_l) > w(i, o_k)$. As we have seen, o_l is necessarily allocated before o_k in the execution of $\vec{\sigma}$. Let j be the agent who receives o_l . Using the same argument for j and o_l we find another object $o_p \in \mathcal{O}'$ allocated before o_l in the sequence. Iterating this argument, since \mathcal{O}' is finite, we will eventually find an object which has been encountered before. This creates a cycle in the precedence relation

of the objects in the execution of the sequence. Contradiction: no sequence can thus generate $\vec{\pi}$. \square

Besides characterizing 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.

PROPOSITION 2. *Let $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ be an instance and $\vec{\pi}$ be an allocation of this instance. We can decide in time $O(n \times m^2)$ if $\vec{\pi}$ is sequenceable.*

The proof is based on the execution of Algorithm 2. This algorithm is similar in spirit to the one proposed by Brams and King [15] but is more general because (i) it can involve non-strict preferences on objects, and (ii) it can conclude with non-sequenceability.

Algorithm 2: Sequencing an allocation

Input: $I = \langle \mathcal{N}, \mathcal{O}, w \rangle$ and $\vec{\pi} \in \mathcal{P}(I)$
Output: a sequence $\vec{\sigma}$ generating $\vec{\pi}$ or **NonSeq**

```

1  $(\vec{\sigma}, \mathcal{O}') \leftarrow (\langle \rangle, \mathcal{O})$ ;
2 for  $t$  from 1 to  $m$  do
3   if  $\exists i$  such that  $\text{best}(\mathcal{O}', i) \cap \pi_i \neq \emptyset$  then
4     Append  $i$  to  $\vec{\sigma}$ ;
5     let  $o_k \in \text{best}(\mathcal{O}', i) \cap \pi_i$ ;
6      $\mathcal{O}' \leftarrow \mathcal{O}' \setminus \{o_k\}$ ;
7   else return NonSeq ;
8 return  $\vec{\sigma}$ ;
```

PROOF. We show that Algorithm 2 returns a sequence $\vec{\sigma}$ generating the input allocation $\vec{\pi}$ if and only if there is one. Suppose that the algorithm returns a sequence $\vec{\sigma}$. 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 **NonSeq**. Then, at a given step t , $\forall i$, $\text{best}(\mathcal{O}', i) \cap \pi_i = \emptyset$. By definition, $\vec{\pi}|_{\mathcal{O}'}$ is therefore, at this step, a frustrating sub-allocation of $\vec{\pi}$. By Proposition 1, $\vec{\pi}$ 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 completed in time $O(m)$. This loop being executed m times, the algorithm runs in $O(n \times m^2)$. \square

4 PARETO-OPTIMALITY

An allocation is Pareto-optimal if no other allocation dominates it. In our context, allocation $\vec{\pi}'$ dominates allocation $\vec{\pi}$ 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 . 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?

This question has already been extensively discussed independently by Aziz *et al.* [3] and Bouveret and Lemaître [13]. We complete the discussion here to give more insights about the implications of the previous results in our framework.

Brams and King [15, Proposition 1] prove the equivalence between sequenceability and Pareto-optimality. However, they have

a different notion of Pareto-optimality, because the agents' preferences are given as linear orders over *objects*. To be able to compare bundles, these preferences are lifted on subsets using the *responsive set extension* \succ_{RS} . This extension leaves many bundles incomparable and leads to define *possible* and *necessary* Pareto-optimality. Brams and King's notion is 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 $u(o_k) > u(o_l)$ if and only if $o_k \succ o_l$). This characterization of responsive dominance yields the following reinterpretation of Brams and King's result: an allocation $\vec{\pi}$ is sequenceable if and only if for each other allocation $\vec{\pi}'$, there is a set u_1, \dots, u_n of additive utility functions, respectively compatible with \succ_1, \dots, \succ_n such that $u_i(\pi_i) > u_i(\pi'_i)$ for at least one agent i .

The latter notion of Pareto-optimality is very weak, because (unlike in our context) the set of additive utility functions is not fixed – we just have to find one that works. Under our stronger notion, the equivalence between sequenceability and Pareto-optimality no longer holds.³

Example 4.1. 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 $\vec{\pi} = \langle o_1, o_2, o_3 \rangle$ giving utilities $\langle 5, 3 \rangle$. $\vec{\pi}$ is then sequenceable but it is dominated by $\vec{\pi}' = \langle o_2, o_3, o_1 \rangle$, giving utilities $\langle 6, 8 \rangle$ (and generated by $\langle 2, 1, 1 \rangle$). Observe that, under ordinal linear preferences, $\vec{\pi}'$ would not dominate $\vec{\pi}$, but they would be incomparable.

The last example shows that a sequence of sincere choices does not necessarily generate a Pareto-optimal allocation. What about the converse? We can see, as a trivial corollary of the reinterpretation of Brams and King's result in our terminology, that the answer is positive *if the preferences are strict on shares*. The following result is more general, because it holds even without this assumption:

PROPOSITION 3 ([3, 13]). *Every Pareto-optimal allocation is sequenceable.*

As already noticed by Aziz *et al.* [3], the proof follows from an adaptation of Brams and King's Proposition 1 (necessity part of the proof) [15]. However, we find useful to give the proof, because it is more general than the previous one, and will be reused in subsequent results of this paper. Before giving this proof, we illustrate it on a concrete example [12, Example 5].

Example 4.2. Let us consider the following instance:

$$\begin{pmatrix} \dagger(12) & 15 & \dagger 11 & \textcircled{7} & 2 \\ 2 & 12 & \textcircled{7} & \dagger 15 & \dagger(11) \\ 15 & \dagger(20) & 9 & 2 & 1 \end{pmatrix}$$

The circled allocation $\vec{\pi}$ is not sequenceable: indeed, every sequence that could generate it should start with $\langle 3, 1, \dots \rangle$, leaving the frustrating sub-allocation $\vec{\rho}$ in a dotted box.

³Actually, since it is known [1, 21] that testing Pareto-optimality with additive preferences in coNP-complete, and that testing sequenceability is in P (Proposition 2), they cannot be equivalent unless P = coNP.

Let us consider agent 1 for instance. Since the suballocation is frustrating, she does not receive o_3 (which is her top object), but agent 2 does. This latter agent, however, does not get her top object, o_4 , because agent 1 receives it. Obviously, if agent 1 gives o_4 to agent 2 and receives o_3 in return, we have built a cycle in which each agent gives a regular object and receives a top one. Doing this, we have built an allocation strictly dominating $\vec{\pi}$.

PROOF. As stated in the example, we will now prove the contraposition of the proposition: every non-sequenceable allocation is dominated. Let $\vec{\pi}$ be a non-sequenceable allocation. From Proposition 1, in a non-sequenceable allocation, there is at least one frustrating sub-allocation. Let $\vec{\rho}$ be such a sub-allocation (that can be $\vec{\pi}$ itself). We will, from $\vec{\rho}$, build another sub-allocation dominating it. Let us choose an arbitrary agent i involved in $\vec{\rho}$, receiving an object not among her top ones in $\vec{\rho}$. Let o_i be a top object of i in $\vec{\rho}$, and let j ($\neq i$) be the unique agent receiving it in $\vec{\rho}$. Let o_j be a top object of j . We can notice that $o_j \neq o_i$ (otherwise j would obtain one of her top objects and $\vec{\rho}$ would not be frustrating). Let k be the unique agent receiving o_j in $\vec{\rho}$, and so on. Using this argument iteratively, we form a path starting from i 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 i be the first such agent and o_k be the last object seen before her in the sequence (i is the unique agent receiving o_k). We have built a cycle $i \xrightarrow{o_k} k \xrightarrow{o_{k-1}} k-1 \cdots i+1 \xrightarrow{o_i} i$ 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 $\vec{\rho}$ 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 $\vec{\rho}$ (and hence in $\vec{\pi}$): $(i \leftarrow o_k)(i+1 \leftarrow o_i) \cdots (k \leftarrow o_{k-1})$ are replaced by: $(i \leftarrow o_i)(i+1 \leftarrow o_{i+1}) \cdots (k \leftarrow o_k)$ where $(i \leftarrow o_i)$ means that o_i is attributed to i . The same substitutions operated in $\vec{\pi}$ yield an allocation $\vec{\pi}'$ that dominates $\vec{\pi}$. \square

COROLLARY 4.3. *No frustrating allocation can be Pareto-optimal (equivalently, in every Pareto-optimal allocation, at least one agent receives a top object).*

Proposition 3 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. We will see that this scale can be further refined.

5 CYCLE DEALS-OPTIMALITY

Pareto-optimality can be thought as a reallocation of objects among agents using improving *deals* [35, 37], as we have seen, to some extent, in the proof of Proposition 3. *Trading cycles* or *cycle deals* constitute a sub-class of deals, which is classical and used, *e.g.*, by Varian [39, page 79] and Lipton *et al.* [31, Lemma 2.2] in the context

of envy-freeness. Trying to link efficiency concepts with various notions of deals is thus a natural idea.

Definition 5.1. Let $\langle \mathcal{N}, \mathcal{O}, w \rangle$ be an add-MARA instance and $\vec{\pi}$ be an allocation of this instance. A (N, M) -cycle deal of $\vec{\pi}$ is a sequence of transfers of items $\mu = \langle (\mu_1, \mathcal{O}_1), \dots, (\mu_N, \mathcal{O}_N) \rangle$, where, for each $j \in \{1, \dots, N\}$, μ_j denotes the j^{th} agent involved in the sequence and $\mu_j \in \mathcal{N}$, $\mathcal{O}_j \subseteq \pi_j$, and $|\mathcal{O}_j| \leq M$. The allocation $\vec{\pi}[\leftarrow \mu]$ resulting from the application of μ to $\vec{\pi}$ is defined as follows:

- $\pi[\leftarrow \mu]_{\mu_j} = \pi_{\mu_j} \setminus \mathcal{O}_j \cup \mathcal{O}_{j-1}$ for $j \in \{2, \dots, N\}$;
- $\pi[\leftarrow \mu]_{\mu_1} = \pi_{\mu_1} \setminus \mathcal{O}_1 \cup \mathcal{O}_N$;
- $\pi[\leftarrow \mu]_i = \pi_i$ if $i \notin \{\mu_1, \dots, \mu_N\}$.

A cycle deal $\langle (\mu_1, \mathcal{O}_1), \dots, (\mu_N, \mathcal{O}_N) \rangle$ will be written

$$\mu_1 \xrightarrow{\mathcal{O}_1} \mu_2 \dots \mu_{N-1} \xrightarrow{\mathcal{O}_{N-1}} \mu_N \xrightarrow{\mathcal{O}_N} \mu_1.$$

In other words, in a cycle deal (we omit N and M when they are not necessary to understand the context), each agent gives a subset of at most M items from her share to the next agent in the sequence and receives in return a subset from the previous agent. $(N, 1)$ -cycle deals will be denoted by N -cycle deals. 2-cycle deals will be called *swap*-deals. Among these cycle deals, some are more interesting: those where each agent improves her utility by trading objects. More formally, a deal d will be called *weakly improving* if $u_i(\pi[\leftarrow d]_i) \geq u_i(\pi_i) \forall i \in \mathcal{N}$ with at least one of these inequalities being strict, and *strictly improving* if all these inequalities are strict.

Intuitively, if it is possible to improve an allocation by applying an improving cycle deal, then it means that this allocation is inefficient. Reallocating the items according to the deal will make everyone better-off. It is thus natural to derive a concept of efficiency from this notion of cycle-deal.

Definition 5.2. An allocation is said to be \succ - (N, M) -cycle optimal (resp. \geq - (N, M) -cycle optimal) if it does not admit any strictly (resp. weakly) improving (K, M) -cycle deal for any $K \leq N$.

We begin with easy observations. First, \geq -cycle optimality implies \succ -cycle optimality, and these two notions become equivalent when the preferences are strict on shares. Moreover, restricting the size of the cycles and the size of the bundles exchanged yield less possible deals and hence lead to weaker optimality notions.

Note that for $N' \leq N$ and $M' \leq M$ (at least one of these inequalities being strict), \succ - (N, M) -cycle-optimality and \geq - (N', M') -cycle-optimality are incomparable. These observations show that cycle-deal optimality notions form a (non-linear) hierarchy of efficiency concepts of diverse strengths. The natural question is whether they can be related to sequenceability and Pareto-optimality. Obviously, Pareto-optimality implies both \succ -cycle-optimality and \geq -cycle-optimality. An easy adaptation of the proof of Proposition 3 leads to the following stronger result:

PROPOSITION 4. *An allocation $\vec{\pi}$ is sequenceable if and only if it is \succ - n -cycle optimal (with $n = |\mathcal{N}|$).*

PROOF. Let $\vec{\pi}$ be a non-sequenceable allocation. Then by Proposition 1, there is at least one frustrating sub-allocation in $\vec{\pi}$. Using the same line of arguments as in the proof of Proposition 3 we can build a strictly improving k -cycle. Hence $\vec{\pi}$ is not \succ -cycle

optimal. Conversely, suppose that $\vec{\pi}$ admits a strictly improving k -cycle deal. Then obviously this cycle yields a sub-allocation that is frustrating. \square

The scale of efficiency introduced in Section 4 can then be refined with a hierarchy of \succ -cycle optimality notions below sequenceable allocations: Pareto-optimal \Rightarrow sequenceable $\Leftrightarrow \succ$ - n -cycle optimal $\Rightarrow \succ$ - $(n-1)$ -cycle optimal $\Rightarrow \dots \Rightarrow \succ$ -swap optimal.

As for \geq -cycle optimality, it forms a parallel hierarchy between Pareto-optimal and non-sequenceable allocations. Note that sequenceability does not involve any \geq - n -cycle-optimality. Thus, as soon as $k < n$, \geq - k -cycle optimality and sequenceability become incomparable notions.

For instance, for 3 agents, there exist allocations which are \geq -swap optimal but not sequenceable and the other way around:

$$\begin{pmatrix} 2 & \dagger 1 & \textcircled{2} & \dagger 3 \\ \dagger \textcircled{3} & 3 & 1 & \textcircled{2} \\ 1 & \textcircled{2} & \dagger 3 & 1 \end{pmatrix}$$

Here the circled allocation is \geq -swap optimal, but not sequenceable: there exists a strictly improving 3-cycle. At the same time, the dag allocation is sequenceable (by $\langle 2, 3, 1, 1 \rangle$), but not even \geq -swap optimal, since 1 and 2 may agree on a weakly improving deal.

The observations previously made in this section suggest that, in some situations, the most complex cycle deals could be required to reach Pareto-optimal allocations. This is indeed the case—we now make this claim more precise. Observe that to be involved in a weakly improving cycle deal, each agent must pass at least one item, thus for a (n, k) -cycle deal, we have that $k \leq m - (n - 1)$ (i.e. the “largest bundle” circulating in this cycle deal can be at most $m - n + 1$). The following generic example shows that it may be necessary to implement a $(n, m - n + 1)$ -cycle to reach a Pareto-optimal allocation.

	α_1	α_2	\dots	α_{n-1}	β_1	\dots	β_{m-n+1}
1	1	0	0	0	$1/(m-n+1)$	$1/(m-n+1)$	$1/(m-n+1)$
2	1	1	0	0	0	0	0
3	0	1	1	0	0	0	0
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
n	0	0	0	2	$1/(m-n+1)$	$1/(m-n+1)$	$1/(m-n+1)$

Initially, every agent $i = 1, \dots, n-1$ holds item α_i , while agent n holds $\beta_1, \dots, \beta_{m-n+1}$. Hence everyone enjoys utility 1. This allocation is dominated by the allocation where each agent $i = 2, \dots, n$ holds α_{i-1} while agent 1 holds $\beta_1, \dots, \beta_{m-n+1}$. In this allocation, the utilities of agents are instead $\langle 1, 1, \dots, 2 \rangle$. But to obtain α_{n-1} , agent n must get it from $n-1$ who would only release it if she gets α_{i-2} , etc. In the end, agent 1 will only release α_1 if she gets the full bundle $\beta_1, \dots, \beta_{m-n+1}$. Overall there are n agents involved in the deal, exchanging up to $m - n + 1$ items. By construction, it is easy to check that no simpler cycle deal (either in terms of number of items or number of agents) would allow to reach this allocation. Furthermore, there is clearly no other allocation Pareto-dominating the initial allocation.

However, it is important to note that cycle-deals may not be sufficient to reach Pareto-optimal outcomes when there are more items than agents.

Example 5.3. Consider the following example:

$$\begin{pmatrix} \textcircled{3} & \textcircled{6} & \dagger 6 & 0 & 6 & \dagger 4 \\ \dagger 2 & 0 & \textcircled{6} & \textcircled{3} & \dagger 7 & 0 \\ 0 & \dagger 5 & 0 & \dagger 4 & \textcircled{6} & \textcircled{3} \end{pmatrix}$$

Note that in the circled allocation, all agents enjoy the same utility, $\langle 9, 9, 9 \rangle$, and that it is Pareto-dominated by the dag allocation which induces the vector of utilities $\langle 10, 9, 9 \rangle$. We leave it to the reader to check that no swap deal, nor 3-cycle, would be weakly improving. In fact, the only way to reach the dag allocation from this initial allocation would require to implement *simultaneously* two $(3, 1)$ -cycle deals $(1 \rightarrow 2 \rightarrow 3 \text{ and } 3 \rightarrow 2 \rightarrow 1)$.

Finally, a corollary of Propositions 2 and 4 is that checking whether an allocation is \succ - n -cycle optimal can be made in polynomial time (by checking whether it is sequenceable).

More generally, we can observe that checking whether an allocation is (k, k') -cycle optimal can be done by iterating over all k -uples of agents⁴, and for each one iterating over all possible transfers involving less than k' objects. In total, there are $k! \binom{n}{k}$ k -uples of agents (which is upper-bounded by n^{k+1}). For each k -uple, there are at most $\left(\sum_{k''=0}^{k'} \binom{m}{k''}\right)^k$ possible transfers, which is again upper-bounded by $(1+m)^{kk'}$. Hence, in total, checking whether an allocation is (k, k') -cycle optimal can be done in time $O(n^{k+1} \times (1+m)^{kk'})$. This is polynomial in n and m if both k and k' are bounded (as for swap deals).

6 RESTRICTED DOMAINS

We now study the impact of several preference restrictions on the hierarchy of efficiency notions introduced in Section 5.

Strict preferences on objects. When the preferences are strict on objects, then obviously every sequence generates exactly one allocation. The following proposition is stronger and shows that the converse is also true:

PROPOSITION 5. *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, each agent has only one possible choice at her turn in the sequence of sincere choices 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 o_k, o_l . Suppose that exactly t objects are ranked above o_k and o_l . Then the sequence where agent 1 picks $t+1$ items in a row, and 2 picks the $m-t-1$ remaining ones obviously generates two allocations, depending on agent 1's choice at step $t+1$. \square

Same order preferences. We say that the agents have *same order preferences* [12] if there is a permutation $\eta : \mathcal{O} \mapsto \mathcal{O}$ such that for each agent i and each pair of objects o_k and o_l , if $\eta(o_k) < \eta(o_l)$ then $w(i, \eta(o_k)) \geq w(i, \eta(o_l))$.

PROPOSITION 6. *All the allocations of an instance with same order preferences are sequenceable (and actually cycle-deal optimal). Conversely, if all the allocations of an instance are sequenceable, then this instance has same order preferences.*

⁴We do not need to also run through all cycles of strictly less than k agents: such a cycle can be simulated just by appending at the end some agents whose role is just to pass the objects they receive to the next agent.

PROOF. Suppose that the agents have same order preferences, and let $\vec{\pi}$ be an arbitrary allocation. In every sub-allocation of $\vec{\pi}$ 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, $\vec{\pi}$ is sequenceable.

Conversely, let us assume for contradiction that there are two distinct objects o_k and o_l and two distinct agents i and j such that $w(i, o_k) > w(i, o_l)$ and $w(j, o_k) < w(j, o_l)$. The sub-allocation $\vec{\pi}^{\{o_k, o_l\}}$ such that $\pi_i^{\{o_k, o_l\}} = \{o_l\}$ and $\pi_j^{\{o_k, o_l\}} = \{o_k\}$ is frustrating. By Proposition 1, every allocation $\vec{\pi}$ containing this frustrating sub-allocation is non-sequenceable. \square

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

PROPOSITION 7. *For a given instance I , 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 is a consequence of Propositions 5 and 6.

Single-peaked preferences. An interesting domain restriction are single-peaked preferences [10, 22], which, beyond voting, is also relevant in resource allocation settings [6, 20]. Formally, in this context, single-peakness can be defined as follows.

There exists a linear order \succ over the set of objects \mathcal{O} . Let $top(i)$ be the preferred object of i . An agent i has *single-peaked preferences* wrt. \succ if, for any two objects $(o_k, o_l) \in \mathcal{O}$ such that either $top(i) \succ o_l \succ o_k$ or $o_k \succ o_l \succ top(i)$ (i.e. lying on the same "side" of the agent's peak), it is the case that i prefers o_l over o_k .

Interestingly, when preferences are single-peaked, the hierarchy of n -cycle optimality collapses at the second level:

PROPOSITION 8. *If the preferences are single-peaked and additive, then an allocation $\vec{\pi}$ is $\geq n$ -cycle optimal iff it is swap-optimal.*

PROOF. ([20, revisited]) First, note that $\geq n$ -cycle optimality trivially implies swap-optimality. Let us now show the converse.

Let us consider for the sake of contradiction an allocation $\vec{\pi}$ that is swap-optimal and such that there exists a $\geq k$ -cycle μ , with $k \leq n$. Without loss of generality, let us suppose that $\mu = \langle (\mu_1, \{o_1\}), \dots, (\mu_k, \{o_k\}) \rangle$. We show by induction on k , the length of μ , that such a cycle can not exist.

Base case: $k = 2$ A 1-cycle of length $k = 2$ is a swap-deal but as $\vec{\pi}$ is swap-optimal, no improving swap-deal exists in $\vec{\pi}$ hence the contradiction.

Induction step: Let us assume that for each k' such that $2 \leq k' \leq k-1$, no $\geq k'$ -cycle exists in $\vec{\pi}$ and let us show that no cycle of length k exists.

To exhibit a contradiction we will need to use the following necessary condition [7]: to be single-peaked, a profile U needs to be worst-restricted, i.e. for any triple of objects $O = (o_a, o_b, o_c) \in \mathcal{O}^3$ there always exists an object $o_j \in O$ such that there exists an agent i with $o_j \notin \text{argmin}_{o_k \in O} w(i, o_k)$ [36].

Because μ is a $\geq k$ -cycle, for all agent $\mu_i \neq \mu_1$ involved in μ we have $o_{i-1} \succ_{\mu_i} o_i$ and $o_k \succ_{\mu_1} o_1$. As no $\geq k'$ -cycle exists, with $k' < k$, for all agents $\mu_i \neq \mu_1$ involved in μ and for all objects o_l in

$\mu, o_l \neq o_i$ and $o_l \neq o_{i-1}$, we have $o_i >_{\mu_i} o_l$. Moreover for all objects o_l in $\mu, o_l \neq o_1$ and $o_l \neq o_k$, we have $o_1 >_{\mu_1} o_l$. If the preferences do not respect these conditions, a \geq - k' -cycle exists with $k' < k$.

Because the profile is worst-restricted, for all the triple of objects O in $\{o_1, \dots, o_k\}$, at most two resources of O can be ranked last among O by the agents. Let us call o_w one of these objects ranked last by agent μ_l and held by agent μ_w . Thanks to the previous paragraph, we know that $\text{best}(O, \mu_w) = o_{w-1}$ and so, because her preferences are single-peaked, μ_w puts o_{w+1} in last position among o_{w-1}, o_w, o_{w+1} . The same holds for agent μ_{w+1} who ranks o_{w-1} in last position among o_{w-1}, o_w, o_{w+1} (because $\text{top}(\mu_{w+1}) = o_w$). Therefore when we focus only on the three objects o_{w-1}, o_w, o_{w+1} , each of them is ranked last among them by one agent which violates the condition of worst-restriction. The contradiction is set, no \geq - k -cycle exists in $\vec{\pi}$. \square

Together with Proposition 4, Proposition 8 gives another interpretation of sequenceability in this domain:

COROLLARY 6.1. *If preferences are single-peaked (and additive), then an allocation $\vec{\pi}$ is sequenceable if and only if it is swap-optimal.*

Proposition 1 by Damamme *et al.* [20] is much stronger than our Corollary 6.1, as it shows that swap-optimality is actually equivalent to Pareto-efficiency *when each agent receives a single resource*. Unfortunately, in our context where each agent can receive several items, this is no longer the case, as the following example shows:

Example 6.2. Consider this instance, single-peaked with respect to $1 \triangleright \dots \triangleright 6$:

$$\begin{pmatrix} \dagger \textcircled{1} & \textcircled{2} & 3 & 4 & 5 & \dagger 6 \\ 1 & \dagger 3 & \textcircled{4} & \textcircled{5} & \dagger 6 & 2 \\ 1 & 2 & \dagger 4 & \dagger 5 & \textcircled{6} & \textcircled{3} \end{pmatrix}$$

The circled allocation is swap-optimal, but Pareto-dominated by the allocation marked with dag.

7 ENVY-FREENESS AND CEEI

The use of sequences of sincere choices can also be motivated by the search for a *fair* allocation protocol. Here, we will focus on two fairness properties and analyze their link with sequenceability.

The first of these notions is probably one of the most prominent fairness properties: envy-freeness [25, 38, 39].

Definition 7.1. Let I be an add-MARA instance and $\vec{\pi}$ be an allocation. $\vec{\pi}$ verifies the *envy-freeness property* (or is simply *envy-free*), when $u_i(\pi_i) \geq u_i(\pi_j), \forall (i, j) \in N^2$ (no agent strictly prefers the share of any other agent).

The notion of *competitive equilibrium* is an old and well-known concept in economics [24, 40]. If equal incomes are imposed among the stakeholders, this concept becomes the *competitive equilibrium from equal incomes* [32], yielding a very strong fairness concept that has been recently explored both in artificial intelligence and in economics [12, 19, 33].

Definition 7.2. Let $I = (N, O, w)$ be an add-MARA instance, $\vec{\pi}$ an allocation, and $\vec{p} \in [0, 1]^m$ a vector of prices. A pair $(\vec{\pi}, \vec{p})$ is said to form a *competitive equilibrium from equal incomes (CEEI)* if

$$\forall i \in N : \pi_i \in \text{argmax}_{\pi \subseteq O} \left\{ u_i(\pi) : \sum_{o_k \in \pi} p_k \leq 1 \right\}.$$

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 each object o_k is p_k .

We will say that allocation $\vec{\pi}$ is a CEEI if there exists a vector \vec{p} such that $(\vec{\pi}, \vec{p})$ forms a CEEI.

As Bouveret and Lemaître [12] and Brânzei *et al.* [18] have shown, with additive preferences, every CEEI allocation is envy-free. 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 9. *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 4.2 above, for which we can check that the circled allocation $\vec{\pi}$ is envy-free and non-sequenceable. \square

Concerning CEEI, it is already well known that any CEEI allocation is Pareto-optimal (hence sequenceable) if the preferences are strict on shares [12]. This is also a consequence of the First Welfare Theorem introduced by Babaioff *et al.* [5] for indivisible goods.

However, surprisingly, this result does not hold anymore if the preferences are not strict on shares, as the following example shows:

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

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

In spite of this negative result, we can still guarantee a certain level of efficiency for CEEI allocations:

PROPOSITION 10. *Every CEEI allocation is sequenceable.*

PROOF. We will show that no allocation can be at the same time non-sequenceable and CEEI. Let $\vec{\pi}$ be a non-sequenceable allocation. We can use the same terms and notations than in the proof of Proposition 3, especially concerning the dominance cycle.

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

$$\pi_i = \{o_k\} \cup \tau_i \quad \pi_{i+1} = \{o_i\} \cup \tau_{i+1} \quad \dots \quad \pi_k = \{o_{k-1}\} \cup \tau_k$$

whereas the allocation $\vec{\pi}'$ that dominates it, contains:

$$\pi'_i = \{o_i\} \cup \tau_i \quad \pi'_{i+1} = \{o_{i+1}\} \cup \tau_{i+1} \quad \dots \quad \pi'_k = \{o_k\} \cup \tau_k$$

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

Suppose that $\vec{\pi}$ is CEEI. This allocation must satisfy two kinds of constraints. First, $\vec{\pi}$ must satisfy the price constraint. If we write $p(\pi) \stackrel{\text{def}}{=} \sum_{o_k \in \pi} p_k$, we have, $\forall i \in C, p(\pi_i) \leq 1$ (1).

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




Figure 1: Distribution of the number of allocations by pair of (efficiency, fairness) criteria.

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

$$p\left(\bigcup_{j \in C} \pi_j\right) \leq |C| \quad \text{and} \quad p\left(\bigcup_{j \in C} \pi'_j\right) > |C|$$

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

8 EXPERIMENTS

We have exhibited in Sections 4 and 5 a “hierarchy of allocation efficiency” made of several steps: Pareto-optimal (PO), sequenceable (Seq), {cycle-deal-optimal}, non-sequenceable (NS). 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 by experimentally studying the distribution of allocations between the different levels. For cycle-deal optimality, we focus on the simplest type of deals, namely, $>$ -swap-deals. We thus have a linear scale of efficiency concepts, from the strongest to the weakest: PO \rightarrow Seq \rightarrow Swap \rightarrow NS. We also analyze the relation between efficiency and various notions of fairness by linking this latter scale with the 6-level scale of fairness introduced by Bouveret and Lemaître [12]: CEEI \rightarrow Envy-Freeness (EF) \rightarrow min-max share (mMS) \rightarrow proportionality (PFS) \rightarrow max-min share (MMS) \rightarrow NS. We generate 50 add-MARA instances involving 3 agents-8 objects, using two different models. For both models, a set of weights are uniformly drawn in the interval $\llbracket 0, 100 \rrbracket$ and the instances are then normalized. For the second model, these weights are reordered afterwards to make the preferences single-peaked. For each instance, we generate all 6561 allocations, and identify for each of them the *highest* level of fairness and efficiency satisfied. The average number of allocations with min-max interval is plotted as a box for each level on a logarithmic scale in Figure 1.

Note that some fairness and efficiency tests require to solve NP-hard or coNP-hard problems (MMS, mMS, and PO tests). These tests are delegated to an external ILP solver. This is especially interesting for the CEEI test which is known to be NP-hard [18],

and for which, to the best of our knowledge, no practical method had been described before. The implementation is available as a fully documented and tested Free Python library.⁵

We note several interesting facts. First, a majority of allocations do not have any efficiency nor fairness property (first black bar on the left). Second, the distribution of allocations on the efficiency scale seems to be related to the fairness criteria: a higher proportion of swap-optimal or sequenceable allocations are found among envy-free allocations than among allocations that do not satisfy any fairness property, and for CEEI allocations, there are even more Pareto-optimal allocations than just sequenceable ones. Lastly, the absence of vertical bar for swap-optimality in the experiments concerning single-peaked preferences confirms the results of Corollary 6.1: in this context, no allocation can be swap-optimal but not Sequenceable; hence, all the allocations that are swap-optimal are contained in the bars concerning sequenceable or Pareto-optimal allocations. Similarly, the absence of bars for swap-optimality and NS (non-sequenceable) in both graphs for the CEEI case confirms the result of Proposition 10.

9 CONCLUSION

In this paper, we have shown that picking sequences and cycle-deals can be reinterpreted to form a rich hierarchy of efficiency concepts. Many interesting questions remain open, such as the complexity of computing cycle-deals or the link between efficiency concepts and social welfare. One could also think of further extending the efficiency hierarchy by studying restrictions on possible sequences (e.g. alternating) or extending the types of deals to non-cyclic ones.

ACKNOWLEDGMENTS

This work is partially supported by the ANR project 14-CE24-0007-01 - CoCoRiCo-CoDec.

REFERENCES

- [1] Haris Aziz, Péter Biró, Jérôme Lang, Julien Lesca, and Jérôme Monnot. 2016. Optimal Reallocation Under Additive and Ordinal Preferences. In *Proceedings of*

⁵Available at: <https://gricad-gitlab.univ-grenoble-alpes.fr/bouveres/fairdiv>.

- the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'16). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 402–410.
- [2] Haris Aziz, Serge Gaspers, Simon Mackenzie, and Toby Walsh. 2015. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227 (2015), 71–92.
 - [3] Haris Aziz, Thomas Kalinowski, Toby Walsh, and Lirong Xia. 2016. Welfare of Sequential Allocation Mechanisms for Indivisible Goods. In *22nd European Conference on Artificial Intelligence (ECAI 2016)*, Gal A. Kaminka, Maria Fox, Paolo Bouquet, Eyke Hüllermeier, Virginia Dignum, Frank Dignum, and Frank van Harmelen (Eds.). IOS Press, 787–794.
 - [4] Haris Aziz, Toby Walsh, and Lirong Xia. 2015. Possible and necessary allocations via sequential mechanisms. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*. AAAI Press, 468–474.
 - [5] Moshe Babaioff, Noam Nisan, and Inbal Talgam-Cohen. 2017. Competitive equilibria with indivisible goods and generic budgets. *arXiv preprint arXiv:1703.08150* (2017).
 - [6] Sophie Bade. 2018. Matching with single-peaked preferences. *Journal of Economic Theory* 180 (2018), 81–99.
 - [7] Miguel A Ballester and Guillaume Haeringer. 2011. A characterization of the single-peaked domain. *Social Choice and Welfare* 36, 2 (2011), 305–322.
 - [8] Nikhil Bansal and Maxim Sviridenko. 2006. The Santa Claus problem. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing (STOC '06)*. ACM, New York, NY, USA, 31–40.
 - [9] Nicola Bianchessi, Jean-François Cordeau, Jacques Desrosiers, Gilbert Laporte, and Vincent Raymond. 2007. A Heuristic for the Multi-satellite, Multi-orbit and Multi-user Management of Earth Observation Satellites. *European Journal of Operational Research* 177, 2 (March 2007), 750–762.
 - [10] David Black. 1948. On the rationale of group decision-making. *The journal of political economy* (1948), 23–34.
 - [11] Sylvain Bouveret and Jérôme Lang. 2011. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, Toby Walsh (Ed.). IJCAI/AAAI, Barcelona, Spain, 73–78.
 - [12] Sylvain Bouveret and Michel Lemaître. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems* 30, 2 (2016), 259–290.
 - [13] Sylvain Bouveret and Michel Lemaître. 2016. Efficiency and Sequenceability in Fair Division of Indivisible Goods with Additive Preferences. Proceedings of the Sixth International Workshop on Computational Social Choice (COMSOC'16). (2016).
 - [14] Steven J. Brams, Marc D. Kilgour, and Christian Klamler. 2012. The undercut procedure: an algorithm for the envy-free division of indivisible items. *Social Choice and Welfare* 39, 2-3 (2012), 615–631.
 - [15] Steven J. Brams and Daniel King. 2005. Efficient Fair Division—Help the Worst off or Avoid Envy? *Rationality and Society* 17, 4 (2005), 387–421.
 - [16] Steven J. Brams and Alan D. Taylor. 1996. *Fair Division — From Cake-cutting to Dispute Resolution*. Cambridge University Press.
 - [17] Steven J. Brams and Alan D. Taylor. 2000. *The Win-win Solution. Guaranteeing Fair Shares to Everybody*. W. W. Norton & Company.
 - [18] Simina Brânzei, Hadi Hosseini, and Peter Bro Miltersen. 2015. Characterization and computation of equilibria for indivisible goods. In *International Symposium on Algorithmic Game Theory*. Springer, 244–255.
 - [19] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (dec 2011), 1061–1103.
 - [20] Anastasia Damamme, Aurélie Beynier, Yann Chevaleyre, and Nicolas Maudet. 2015. The Power of Swap Deals in Distributed Resource Allocation. In *The 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*. Istanbul, Turkey, 625–633.
 - [21] Bart de Keijzer, Sylvain Bouveret, Tomas Klos, and Yingqian Zhang. 2009. 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)*. Springer Verlag, Venice, Italy, 98–110.
 - [22] Edith Elkind, Martin Lackner, and Dominik Peters. 2016. Preference Restrictions in Computational Social Choice: Recent Progress. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*. 4062–4065.
 - [23] Ulle Endriss, Nicolas Maudet, Fariba Sadri, and Francesca Toni. 2006. Negotiating Socially Optimal Allocations of Resources. *Journal of Artificial Intelligence Research* 25 (2006), 315–348.
 - [24] Irving Fisher. 1892. *Mathematical Investigations in the Theory of Value and Prices, and Appreciation and Interest*. Augustus M. Kelley, Publishers.
 - [25] Duncan K. Foley. 1967. Resource Allocation and the Public Sector. *Yale Economic Essays* 7, 1 (1967), 45–98.
 - [26] Judy Goldsmith and Robert H Sloan. 2007. The AI conference paper assignment problem. In *Proc. AAAI Workshop on Preference Handling for Artificial Intelligence, Vancouver*. 53–57.
 - [27] Thomas Kalinowski, Nina Narodytska, and Toby Walsh. 2013. A Social Welfare Optimal Sequential Allocation Procedure. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-13)*, Francesca Rossi (Ed.). IJCAI/AAAI, Beijing, China, 227–233.
 - [28] Thomas Kalinowski, Nina Narodytska, Toby Walsh, and Lirong Xia. 2013. Strategic Behavior when Allocating Indivisible Goods Sequentially. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI-12)*. AAAI Press, Bellevue, WA, 452–458.
 - [29] David A Kohler and R Chandrasekaran. 1971. A class of sequential games. *Operations Research* 19, 2 (1971), 270–277.
 - [30] Michel Lemaître, Gérard Verfaillie, and Nicolas Bataille. 1999. 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)*. Stockholm, Sweden, 206–211.
 - [31] Richard Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On Approximately Fair Allocations of Divisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC-04)*. ACM, New York, NY, 125–131.
 - [32] Hervé Moulin. 2003. *Fair Division and Collective Welfare*. MIT Press.
 - [33] Abraham Othman, Tuomas Sandholm, and Eric Budish. 2010. Finding approximate competitive equilibria: efficient and fair course allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-10)*, Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen (Eds.). IFAAMAS, Toronto, Canada, 873–880.
 - [34] Senjuti Basu Roy, Joanna Lykouroutzou, Saravanan Thirumuruganathan, Sihem Amer-Yahia, and Gautam Das. 2015. Task assignment optimization in knowledge-intensive crowdsourcing. *VLDB J.* 24, 4 (2015), 467–491.
 - [35] Tuomas W. Sandholm. 1998. Contract Types for Satisficing Task Allocation: I. Theoretical Results. In *Proceedings of the AAAI Spring Symposium: Satisficing Models*, Sandip Sen (Ed.). AAAI Press, Menlo Park, California, 68–75.
 - [36] Amartya K Sen. 1966. A possibility theorem on majority decisions. *Econometrica* (1966), 491–499.
 - [37] Lloyd Shapley and Herbert Scarf. 1974. On cores and indivisibility. *Journal of mathematical economics* 1, 1 (1974), 23–37.
 - [38] Jan Tinbergen. 1953. *Redelijke Inkomensverdeling*. N. V. DeGulden Pers., Haarlem.
 - [39] Hal R. Varian. 1974. Equity, Envy and Efficiency. *Journal of Economic Theory* 9 (1974), 63–91.
 - [40] Léon Walras. 1874. *Éléments d'économie politique pure ou Théorie de la richesse sociale* (1st ed.). L. Corbaz.