Trends in Computational Social Choice

12

Cite as: Evangelos Markakis. Approximation Algorithms and Hardness Results for Fair Division with Indivisible Goods. In Ulle Endriss (editor), *Trends in Computational Social Choice*, chapter 12, pages 231–247. AI Access, 2017.

http://www.illc.uva.nl/COST-IC1205/Book/

CHAPTER 12

Approximation Algorithms and Hardness Results for Fair Division with Indivisible Goods

Evangelos Markakis

12.1 Introduction

Fair division problems have attracted the attention of various scientific disciplines in the last decades, including among others, mathematics, economics, computer science, and political science. Ever since the first attempt for a mathematical treatment by Steinhaus, Banach, and Knaster (Steinhaus, 1948), many interesting and challenging questions have emerged. Over the years, a vast literature has developed, see e.g., Brams and Taylor (1996), Robertson and Webb (1998), and Moulin (2003), considering several notions of fairness. For more recent surveys, see also Bouveret et al. (2016), and Procaccia (2016).

The objective in fair division is to allocate a set of resources to a set of agents in a way that leaves every agent satisfied, to the extent that is feasible. To model the preferences of the involved agents, the standard assumption in the literature is that every agent is associated with a valuation function on the set of resources. In most settings, valuation functions are further restricted to be additive functions, but non-additive scenarios are also discussed in some works. Under this setup, various solution concepts have been proposed as to what constitutes a fair allocation, including e.g., *proportionality* and *envy-freeness*, along with several variants, strengthenings and relaxations.

The models that have been studied so far, can essentially be grouped into two classes. The first one concerns *continuous* models, where it is assumed that the resources are infinitely divisible. This is appropriate when items can be split into arbitrarily smaller pieces, or in settings where the resources correspond to the percentage of time an agent can use a shared good, or even further in land division. The second class contains *discrete* models, where the resources are seen as *indivisible* goods. This means that the items to be allocated cannot be divided further; each item has to be entirely allocated to one agent. Although the same solution concepts apply well to both models, the picture is quite different when it comes to existence and computation of fair outcomes. Under continuous models, we can guarantee existence of most standard fairness concepts, and in some cases, we can also have very efficient algorithms, such as the algorithm of Even and Paz (1984) for proportional allocations. On the contrary, in the presence of indivisible goods, we cannot guarantee existence anymore for most fairness concepts, and it is even NP-hard to decide whether a given instance admits a fair allocation.

Undoubtedly, establishing NP-hardness is not encouraging news. However, as with many other optimization problems, one can still resort to algorithms that produce approximation guarantees with respect to the criteria under consideration. Clearly, the solution we would expect from an approximation algorithm depends on the fairness concept that we try to approximate. As an example, if envy is our fairness criterion, then our goal would be an algorithm that always returns an allocation such that the pairwise envy is never more than a small factor away from the optimal envy achievable. Such algorithms provide to us solutions that we could be willing to settle with, given the difficulty of finding an exact optimal solution.

Motivated by this discussion, the purpose of this chapter is to highlight the role of approximation algorithms and inapproximability results for discrete models, towards understanding what we can hope to achieve by polynomial-time algorithms. Consequently, the exposition will be based on presenting recent results published within the last years, along with open problems and trends in allocating indivisible items.

The structure of the remaining chapter is as follows. In Section 12.2, we provide the necessary definitions and we present the fairness concepts that we are interested in. In Section 12.3, we demonstrate problems that admit polynomial-time algorithms and also establish why most interesting problems are NP-hard. Finally in Section 12.4, we present approximation algorithms along with hardness of approximation results, classified according to the fairness criteria under consideration. We conclude with interesting open questions in Section 12.5.

12.2 Preliminaries

In this section, we first provide the general setup and the notation that we will be using. We then define the notions of fairness that will form the backbone of this exposition.

12.2.1 Notation and Terminology

We assume we have a set of n agents, $N = \{1, 2, \ldots, n\}$, and a set of m indivisible goods $M = \{1, 2, \ldots, m\}$. For every agent $i \in N$, the preferences for the items are expressed through a valuation function $v_i(\cdot)$. In its general form, this is a monotone set function $v_i : 2^M \rightarrow \mathbb{R}$, defined on subsets of items, such that for $S \subseteq M$, $v_i(S)$ denotes the value derived by agent i, when he obtains the set S. Monotonicity simply means that for $S \subseteq T \subseteq M$, we have that $v_i(S) \leqslant v_i(T)$. Following the usual assumptions in the majority of the fair division literature, and unless otherwise stated, we consider that each agent has an *additive* valuation function. Hence, for every $S \subseteq M$, $v_i(S) = \sum_{j \in S} v_i(\{j\})$. For $j \in M$, we will use

v_{ij} instead of $v_i(\{j\})$ for simplicity. Note that monotonicity for additive functions implies that $v_{ij} \geqslant 0$ for every $i \in N$ and $j \in M$.

Under additive valuations, it suffices to specify the value of an agent i for each individual item, in order to fully specify the valuation function. Hence, we can represent the relevant information for agent i by the vector $v_i = (v_{i1}, v_{i2}, \ldots, v_{im})$. The input then to any algorithm for fair division can be encoded by the matrix $V = (v_{ij})$, where for each $i \in N$, the i-th row of V corresponds to the valuation of agent i.

We are interested in solutions that allocate the whole set of goods M to the agents. An allocation of the goods to the agents is therefore a partition denoted by $S = (S_1, ..., S_n)$, where S_i is the subset allocated to agent i, $S_i \cap S_j = \emptyset$ and $\bigcup_{i \in N} S_i = M$. Given any set T, we denote by $\Pi_n(T)$ the set of all partitions of T into n bundles. Hence, all the solution concepts we will encounter, seek to produce an element of $\Pi_n(M)$.

12.2.2 Fairness Concepts

Clearly, one cannot hope to have a unique, universally accepted notion of fairness that can be equally applicable to all problems. Several notions have emerged throughout the years as to what can be considered a fair allocation. As it is infeasible to enumerate all these concepts in this exposition, we will only focus on the notions that are most relevant to recent algorithmic research in this area.

We start with two of the most dominant solution concepts in fair division, namely proportionality and envy-freeness.

Definition 12.1. *An allocation $S = (S_1, ..., S_n)$ is*

1. *proportional, if $v_i(S_i) \geqslant \frac{1}{n} v_i(M)$, for every $i \in N$.*
2. *envy-free, if for every $i, j \in N$, $v_i(S_i) \geqslant v_i(S_j)$.*

Proportionality was considered in the very first work on fair division (Steinhaus, 1948). Envy-freeness seems to have been initially suggested by Gamow and Stern (1958), and was considered more formally later by Foley (1967) and Varian (1974).

It can be easily seen that envy-freeness is a stricter notion than proportionality, i.e., harder to achieve. But even for proportionality, existence cannot be guaranteed under indivisible goods. This gives rise to considering relaxations of these definitions, with the hope of obtaining more positive results.

One first such relaxation is the concept of envy-freeness up to one good, where each player may envy another player, but only by an amount which does not exceed the value of a single item in the other player's bundle. Formally:

Definition 12.2. *An allocation $S = (S_1, ..., S_n)$ is envy-free up to one good, if for every pair of agents $i, j \in N$, there exists an item $a_j \in S_j$, such that*

$$v_i(S_i) \geqslant v_i(S_j \setminus \{a_j\})$$

Note that envy-freeness up to one good is a relaxation of envy-freeness but not of proportionality. There exist allocations that are proportional but not envy-free up to one good, and vice versa.

A relaxation of proportionality, namely the notion of maximin share allocations, was recently proposed by Budish (2011), building on concepts by Moulin (1990). Motivated by the question of what can we guarantee in the worst case to the agents, the rationale of this concept is to think of a generalization of the well-known cut-and-choose protocol to multiple agents as follows: suppose that agent i is asked to partition the goods into n bundles and then the rest of the agents choose a bundle before i. In the worst case, agent i will be left with her least valuable bundle. Hence, a risk-averse agent would choose a partition that maximizes the minimum value of a bundle in the partition. This value is called the maximin share of agent i, and for $n = 2$, it is precisely what he could guarantee to himself in the discrete form of the cut-and-choose protocol, by being the cutter. The objective then is to find an allocation where every agent receives at least his maximin share.

Definition 12.3. *Given a set of n agents, a set of goods M, and a valuation matrix V,*

1. *the maximin share of an agent i, is:*

$$\boldsymbol{\mu}_i := \boldsymbol{\mu}_i(n, M, v_i) = \max_{S \in \Pi_n(M)} \min_{S_j \in S} v_i(S_j) \,.$$

2. *an allocation $S = (S_1, ..., S_n) \in \Pi_n(M)$ is called a maximin share (MMS) allocation, if $v_i(S_i) \geqslant \boldsymbol{\mu}_i$, for every agent $i \in N$.*

Note that the maximin share of an agent does not depend on the whole valuation matrix V, but solely on v_i.

A related approach on determining worst case guarantees was taken by Hill (1987). His work examined the guarantee that a player can have as a function of two parameters: the number of players and the maximum value of an item across all players. As a result, the following function was identified, under the assumption that the valuation functions are normalized such that $\sum_{j \in M} v_{ij} = 1$ for every $i \in N$.

Definition 12.4. *Given any integer $n \geqslant 2$, let $V_n : [0, 1] \to \left[0, n^{-1}\right]$ be the nonincreasing function satisfying $V_n(\alpha) = 1/n$ for $\alpha = 0$, and for $\alpha > 0$:*

$$V_n(\alpha) = \begin{cases} 1 - k(n-1)\alpha & \text{if } \alpha \in I(n,k) \\ 1 - \frac{(k+1)(n-1)}{(k+1)n-1} & \text{if } \alpha \in NI(n,k) \end{cases}$$

where for $k \geqslant 1$, $I(n,k) = \left[\frac{k+1}{k((k+1)n-1)}, \frac{1}{kn-1}\right]$, and $NI(n,k) = \left(\frac{1}{(k+1)n-1}, \frac{k+1}{k((k+1)n-1)}\right)$.

Although this function seems complicated and hard to motivate at first sight, we will establish that it has a very interesting interpretation, and plays an important role in understanding what is feasible to achieve with indivisible goods.

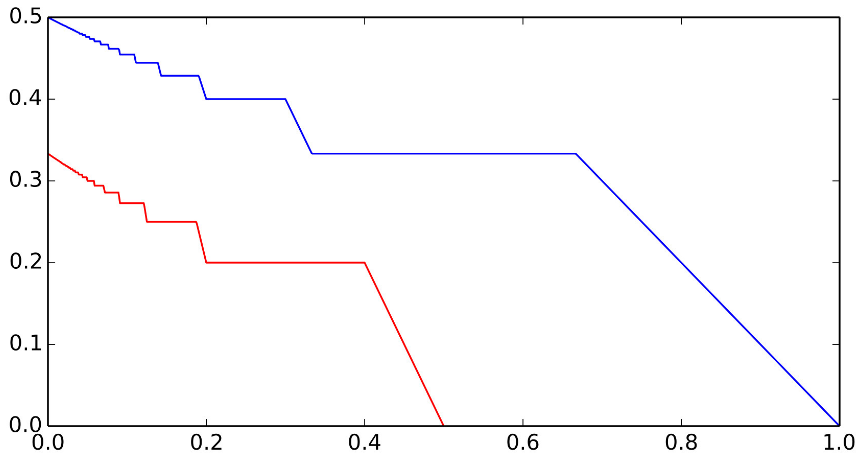

Figure 12.1: The function $V_n(\cdot)$ for $n = 2$ (upper curve) and $n = 3$ (lower curve).

Definition 12.5. *We will call an allocation $S = (S_1, ..., S_n)$ V_n-fair, if for every agent i, it holds that $v_i(S_i) \geqslant V_n(\alpha_i)$, where $\alpha_i = \max_{j\in M} v_{ij}$.*

In Figure 12.1, one can see the function $V_n(\cdot)$ for $n = 2$ and $n = 3$. For larger n, it has a similar form. The function alternates between segments where it is strictly decreasing (and the decrease is linear in α), and segments where it is constant and equal to the value at the left endpoint of the previous decreasing segment. For example, for $n = 2$, looking at the function from right to left, we can see that the rightmost decreasing segment is $[2/3, 1]$, and the function is equal to $1 - \alpha$ within this interval. This means that for two agents, if there exists an item with value $\alpha \geqslant 2/3$ for an agent, we can only guarantee to him an allocation with value $1 - \alpha$, since in worst case the other person will receive that item. In a similar manner, we see that $V_2(1/3) = 1/3$. This means that as long as the maximum value of a good for an agent is at most $1/3$, we can always find an allocation guaranteeing at least $1/3$ to him. The worst case for this is when we have 3 items of value $1/3$ each for both players.

It is not hard to see that this function forms a lower bound for $\boldsymbol{\mu}_i$. Therefore, V_n-fairness is a relaxation of maximin share fairness, which is itself a relaxation of proportionality.

Lemma 12.1. *For every additive valuation function v_i for which $\alpha_i = \max_{j\in M} v_{ij}$, it holds that*

$$\frac{1}{n} v_i(M) \geqslant \boldsymbol{\mu}_i(n, M, v_i) \geqslant V_n(\alpha_i)$$

More recently, Gourvès et al. (2015) introduced a refined version of the function $V_n(\cdot)$, which leads to slightly stronger guarantees in some cases. As this definition is even more complex, we refer the reader to Gourvès et al. (2015) for the exact definition of their refinement.

An Example. To illustrate some of these concepts, let us consider an example. Suppose we have the following instance with three agents and five items:

	a	b	c	d	e
Agent 1	$1/2$	$1/2$	$1/3$	$1/3$	$1/3$
Agent 2	$1/2$	$1/4$	$1/4$	$1/4$	0
Agent 3	$1/2$	$1/2$	1	$1/2$	$1/2$

If $M = \{a, b, c, d, e\}$ is the set of items, one can see that $\boldsymbol{\mu}_1(3, M, v_1) = 1/2$, $\boldsymbol{\mu}_2(3, M, v_2) = 1/4$, $\boldsymbol{\mu}_3(3, M, v_3) = 1$. E.g., for agent 1, no matter how she partitions the items into three bundles, the worst bundle will be worth at most $1/2$ for her, and she achieves this with the partition $(\{a\}, \{b, c\}, \{d, e\})$. Similarly, agent 3 can guarantee a value of 1 (which is best possible as it is equal to $v_3(M)/n$) by the partition $(\{a, b\}, \{c\}, \{d, e\})$. This instance admits a maximin share allocation, e.g., $(\{a\}, \{b, c\}, \{d, e\})$, and in fact this is not a unique such allocation.

Let us look now into an allocation with better fairness properties, namely $S = (\{a, e\}, \{b, d\}, \{c\})$. Obviously this is also a MMS allocation. But furthermore, it is also a proportional and envy-free allocation. It is easy to check that each agent is happy with her bundle and would not envy any of the other two bundles.

12.3 Boundaries of Polynomial Time Tractability

In this section, we first examine which fairness criteria admit efficient algorithms. As we will see, two of the concepts defined in Section 12.2 fall into this class. We then proceed to understand the source of computational difficulty for the remaining criteria, and establish NP-hardness for a variety of problems. The main conclusion of this section is that the majority of the problems being studied in the literature are computationally hard.

12.3.1 Efficient Algorithms

One of the easiest criteria that we can satisfy is envy-freeness up to one good. There are several algorithms that produce such allocations. Perhaps the most intuitive one is a simple round robin algorithm (see Algorithm 1). In describing algorithms, we assume that they take as input the set of players N, the set of goods M, and the valuation matrix V that encodes the valuation functions.

Theorem 12.2. *The allocation produced by Algorithm 1 is envy-free up to one good.*

Proof. This is quite easy to establish. Fix an agent i, and let $j \neq i$. We will upper bound the difference $v_i(S_j) - v_i(S_i)$. If j comes after i in the order chosen by the algorithm, then agent i cannot envy j, since i always picks an item at least as desirable as the one j picks. Suppose that j precedes i in the ordering. The algorithm proceeds in $\ell = \lceil m/n \rceil$ rounds. In each round k, let r_k and r'_k be the items allocated to j and i respectively. Then

$$v_i(S_j) - v_i(S_i) = (v_{i,r_1} - v_{i,r'_1}) + (v_{i,r_2} - v_{i,r'_2}) + \cdots + (v_{i,r_\ell} - v_{i,r'_\ell}) .$$

ALGORITHM 1: Greedy Round-Robin(N, M, V)

Set $S_i = \emptyset$ for each $i \in N$
Fix an arbitrary ordering of the agents
while $\exists$ *unallocated items* **do**
    Let $i \in N$ be the next agent to be examined in the current round (proceeding in a round-robin fashion)
    Let $j \in M$ be the most desired item for i, among the currently unallocated items
    $S_i = S_i \cup \{j\}$
return $(S_1, ..., S_n)$

Note that there may be no item r'_ℓ in the last round if the algorithm runs out of goods but this does not affect the analysis (simply set $v_{i,r'_\ell} = 0$).

Since agent i picks her most desirable item when it is her turn to choose, this means that for two consecutive rounds k and $k+1$ it holds that $v_{i,r'_k} \geqslant v_{i,r_{k+1}}$. This directly implies that $v_i(S_j) - v_i(S_i) \leqslant v_{i,r_1} - v_{i,r'_\ell} \leqslant v_{i,r_1}$. But then, if we remove item r_1 from the bundle of j, agent i will not be envious. $\square$

As we will see, the round-robin algorithm can be useful for other criteria as well. Interestingly, there are more algorithms that also achieve envy-freeness up to one good. For example, see the algorithm by Lipton et al. (2004), based on a graph-theoretic modeling of the problem.

We now move to the notion of V_n-fairness. As already mentioned, the function $V_n(\cdot)$ defined in the previous section, expresses the worst possible guarantee for an agent as a function of the maximum value across all items. Even though it was proved by Hill (1987) that there always exists an allocation providing such guarantees, the proof does not yield an efficient algorithm. Later on, a greedy algorithm was proposed by Markakis and Psomas (2011), which provides an interesting interpretation; the function $V_n(\cdot)$ describes precisely the value guaranteed to an agent by the simple greedy process stated in Algorithm 2.

This algorithm tries to greedily find an agent i who can exceed $V_n(\alpha_i)$ with the least number of items. After doing so, it removes agent i with the goods allocated to her, it normalizes the reduced valuation matrix and continues in the same fashion. The following result about the preformance of the algorithm was proved by Markakis and Psomas (2011).

Theorem 12.3. *Algorithm 2 produces an allocation $S = (S_1, \ldots, S_n)$ such that, for each player i, $v_i(S_i) \geqslant V_n(\alpha_i)$, where $\alpha_i = \max_j v_{ij}$.*

Further improvements have been obtained regarding such greedy outcomes, summarized in the following discussion.

Remark 12.4. *The result of Theorem 12.3 is tight on the decreasing segments of the function $V_n(\cdot)$ (see Figure 12.1). In Gourvès et al. (2015), a refined version of $V_n(\cdot)$ was defined and it was established that one can have a slightly higher guarantee when α_i lies in the segments of $[0, 1]$ where $V_n(\cdot)$ is constant.*

ALGORITHM 2: V_n-FAIR(N, M, V)

Set $S_i = \emptyset$ for each $i \in N$, and let $n = |N|$, $\alpha_i = \max_{j \in M} v_{ij}$
while *for every* $i \in N, v_i(S_i) < V_n(\alpha_i)$ **do**
  for *each* $i \in N$ **do**
    S_i = $S_i \cup$ {next most desired item in M}
  //we add one item at a time to each person's bundle till we find an agent satisfiable with the least number of items
Pick an agent i with $v_i(S_i) \geqslant V_n(\alpha_i)$ //arbitrarily in case of ties
Allocate S_i to agent i
if $n = 2$ **then**
  Allocate all other items to the remaining agent
else
  for $k \in N$, $k \neq i$ **do**
    $v_{kj} = v_{kj}/(1 - v_k(S_i))$
    //normalization before going to next round
  V' = new matrix after removing row i and columns corresponding to S_i
  run V_n-FAIR($V', N \setminus \{i\}, M \setminus S_i$)

12.3.2 Hardness Results

The results in the previous subsection hit essentially the boundaries of what is feasible to achieve by polynomial time algorithms. Most other algorithmic problems regarding fair allocations turn out to be NP-hard. The reason is that a special case of all these problems is the well known PARTITION problem. This has been observed already, among others, by Demko and Hill (1988), and by several other works.

Theorem 12.5. *Even for two agents, it is* NP*-complete to decide if there exists an envy-free allocation. It is also* NP*-complete to decide the existence of a proportional allocation.*

Proof. Deciding the existence of fair allocations clearly belongs to the class NP. To establish NP-hardness, consider an instance I of the PARTITION problem. Such an instance is described by a set of n numbers $A = \{a_1, ..., a_n\}$ and we are asked whether we can split A into two sets of equal value, i.e., whether there exists a set $S \subseteq A$ such that $\sum_{i \in S} a_i = \sum_{i \in A \setminus S} a_i$.

Starting from I, we can define now an instance of our problem, with two agents and n goods. Each number a_i in the PARTITION instance corresponds to a good in our instance with value a_i for both agents. Hence, a proportional or an envy-free allocation exists if and only if both players can receive a bundle of value at least $\sum_{i \in A} a_i/2$. But since the total value of the goods is exactly $\sum_{i \in A} a_i$, we conclude that a proportional allocation exists if and only if there exists a solution to the PARTITION problem. $\square$

One might initially feel more hopeful for the notion of MMS allocations, since this is a relaxation of proportionality. Indeed, for two agents, the existence of MMS allocations is a trivial problem, since they always exist (we note though that the status of existence for higher values of n and m is not completely clear, see Kurokawa et al. (2016)). However, when it comes to computing the actual allocations, even for the case of two agents, we still have an NP-hard problem to solve. The same reduction as in the previous theorem can show the following.

Theorem 12.6. *Even for two agents, finding a MMS allocation is an* NP*-hard problem.*

12.4 Approximation Guarantees: Algorithms and Impossibility Results

Motivated by the hardness results, we will now examine ways to construct approximately fair allocations. As it is not a priori clear what it would mean to be approximately fair, we examine each fairness concept separately and argue about possible approximation versions of each criterion. We summarize in each of the following subsections both positive (in terms of approximation algorithms) and negative (in terms of inapproximability) results that have been established.

12.4.1 Allocations with Maximin Share Guarantees

Apart from the computational difficulty, we know that for $n \geqslant 3$, there exist examples showing that MMS allocations do not always exist (Procaccia and Wang, 2014; Kurokawa et al., 2016). Although these examples are rather extreme, they show that even if we disregard time complexity, we cannot construct algorithms that always compute a MMS allocation.

It is then natural to explore what would be the best guarantee we can give. I.e., how close to a MMS allocation can we come? Can we construct an allocation where every agent receives a bundle of goods with a total value that is "close" to her MMS value? In order to quantify the distance from an actual MMS allocation, we use the notions of *additive* and *multiplicative* approximation. By an additive ρ-approximation, we mean an allocation $(S_1, \ldots, S_n)$, where $v_i(S_i) \geqslant \boldsymbol{\mu}_i - \rho$, for some $\rho \leqslant \boldsymbol{\mu}_i$. As for multiplicative approximations, which is the most common approach used in approximation algorithms, we will mean that we demand an allocation such that $v_i(S_i) \geqslant \rho\boldsymbol{\mu}_i$, for some $\rho \leqslant 1$.

Let us first look for an additive approximation. We claim that the round robin algorithm presented in Section 12.3 provides such a guarantee. In particular, recall that we have defined α_i to be the maximum value of any good for agent i. Let also $\alpha = \max_{i,j} v_{ij} = \max_i \alpha_i$.

Theorem 12.7. *If $(S_1, ..., S_n)$ is the output of Algorithm 1, then for every $i \in N$,*

$$v_i(S_i) \geqslant \frac{v_i(M)}{n} - \alpha_i \geqslant \boldsymbol{\mu}_i - \alpha_i \geqslant \boldsymbol{\mu}_i - \alpha\,.$$

Proof. By Theorem 12.2, we know that the round robin algorithm is envy-free up to one good. This implies that for any $i, j \in N$, we have $v_i(S_i) \geqslant v_i(S_j) - \alpha_i$.

If we now sum up these inequalities over every $j \in N$, we get: $nv_i(S_i) \geqslant \sum_j v_i(S_j) - n\alpha_i$, which implies

$$v_i(S_i) \geqslant \frac{\sum_j v_i(S_j)}{n} - \alpha_i = \frac{v_i(M)}{n} - \alpha_i \geqslant \boldsymbol{\mu}_i - \alpha_i\,,$$

where the last inequality holds since the maximin share guarantee is a relaxation to proportionality. $\square$

This is already a positive result, as it reveals that when no item has a very high value for any agent, we can get an approximately fair allocation.

If we now try to obtain a multiplicative rather than an additive guarantee, we see that we would need somehow to handle carefully the goods that are valued highly, even by one agent. Obviously, such goods should end up at one of the agents who have a high value for them. What is interesting here is that if we do take care of highly valuable goods in this manner, we end up with an instance, where running the round robin algorithm provides a multiplicative guarantee since the maximum value now is quite small. Quantifying "highly valuable" to mean at least half of the total value of an agent for the set of available goods, we obtain Algorithm 3, for which we can establish a $1/2$-approximation.

ALGORITHM 3: APX-MMS$_{1/2}(N, M, V)$

Set $S = M$
for $i = 1$ *to* $|N|$ **do**
  Let $\beta_i = \frac{\sum_{j \in S} v_{ij}}{|N|}$
while $\exists i, j$ *s.t.* $v_{ij} \geqslant \beta_i/2$ **do**
  Allocate j to i.
  $S = S \setminus \{j\}$
  $N = N \setminus \{i\}$
  Recompute the β_is
Run Algorithm 1 on the remaining instance

Theorem 12.8. *Let N be a set of n agents, and let M be a set of goods. Algorithm 3 produces an allocation $(S_1, ..., S_n)$ such that*

$$v_i(S_i) \geqslant \frac{1}{2}\boldsymbol{\mu}_i\,,\ \forall i \in N$$

Proof. The important ingredient of the proof is a simple monotonicity property, which says that we can allocate a single good to an agent without decreasing the maximin share of other agents. Recall that $\boldsymbol{\mu}_i$ is defined with respect to the number of agents, the set of goods and the valuation function of i, i.e., $\boldsymbol{\mu}_i := \boldsymbol{\mu}_i(n, M, v_i)$. This means that an analogous quantity can be defined for any sub-instance of an initial instance, hence, for any $S \subseteq M$ or for a smaller number of agents. The following property is then very easy to establish.

Claim 12.9 (Monotonicity property). *For any agent i and any good j, it holds that*

$$\boldsymbol{\mu}_i(n-1, M \setminus \{j\}, v_i) \geqslant \boldsymbol{\mu}_i(n, M, v_i)\,.$$

We will distinguish two cases. Consider an agent i who was allocated a single item during the first phase of the algorithm (lines 4 - 8). Suppose that at the time when i was given her item, there were n_1 active agents, $n_1 \leqslant n$, and that S was the set of currently unallocated items. By the design of the algorithm, this means that the value of what i received is at least

$$\frac{\sum_{j\in S} v_{ij}}{2n_1} \geqslant \frac{1}{2}\boldsymbol{\mu}_i(n_1, S, v_i)\,.$$

But now if we apply the monotonicity property $n - n_1$ times, we obtain that $\boldsymbol{\mu}_i(n_1, S, v_i) \geqslant \boldsymbol{\mu}_i(n, M, v_i) = \boldsymbol{\mu}_i$, and we are done.

Consider now an agent i, who gets a bundle of goods according to Greedy Round-Robin, in the second phase of Algorithm 3. Let n_2 be the number of active agents at that point, and S be the set of goods that are unallocated before Greedy Round-Robin is executed. We know that the maximum value of any remaining good at that point is less than half the current value of β_i for agent i. Hence by the additive guarantee of Greedy Round-Robin, i.e., the first inequality in Theorem 12.7, we have that the bundle received by agent i has value at least

$$\frac{\sum_{j\in S} v_{ij}}{n_2} - \frac{\beta_i}{2} = \frac{\sum_{j\in S} v_{ij}}{2n_2} \geqslant \frac{1}{2}\boldsymbol{\mu}_i(n_2, S, v_i)\,.$$

Again, after applying Claim 12.9 repeatedly, we get that $\boldsymbol{\mu}_i(n_2, S, v_i) \geqslant \boldsymbol{\mu}_i(n, M, v_i) = \boldsymbol{\mu}_i$, which completes the proof. □

In trying to obtain a better approximation for this problem, one needs to take a different approach. This direction was investigated first by Procaccia and Wang (2014), and followed up by Amanatidis et al. (2015). By using matching arguments in a bipartite graph representation of the problem, and exploiting a more involved monotonicity property underlying the maximin shares, a better approximation factor was obtained, initially in exponential time by Procaccia and Wang (2014), and later in polynomial time by Amanatidis et al. (2015). The following statement summarizes the results of these two works.

Theorem 12.10. *For any n and m, and any constant $\epsilon > 0$, there exists a polynomial time algorithm producing an allocation $(S_1, ..., S_n)$ such that $v_i(S_i) \geqslant (\frac{2}{3} - \epsilon)\boldsymbol{\mu}_i$.*

It is interesting to note that up to now it is still an open question whether a better approximation guarantee is possible for this notion. The impossibility results that have been established by Kurokawa et al. (2016), and Procaccia and Wang (2014) do not rule out even an approximation very close to 1.

Finally, the problem is also nontrivial to understand even with a small number of agents. For $n = 2$, it is pointed out by Bouveret and Lemaître (2014) that maximin share allocations always exist via an analog of the cut-and-choose protocol. Using the result of Woeginger (1997), we can then have a Polynomial Time Approximation Scheme (PTAS), i.e., a $(1 - \epsilon)$-approximation in polynomial time, for any constant $\epsilon > 0$. Hence, despite the NP-hardness, we can still have the best possible positive result we could hope for, when $n = 2$. In contrast, as soon as we move to $n = 3$, it has been proved that there exist instances where no maximin share allocation exists (Procaccia and Wang, 2014). The best known approximation guarantee is roughly $8/9$ and was very recently obtained by Gourvès and Monnot (2017). Surprisingly, it is still unclear what is the best we can achieve for 3 agents. All these findings are summarized in the following statement.

Theorem 12.11. *For any number of items m, and*

- *for $n = 2$, there exists a PTAS, i.e, we can produce in polynomial time an allocation (S_1, S_2) such that $v_i(S_i) \geqslant (1-\epsilon)\boldsymbol{\mu}_i$, for any constant $\epsilon > 0$.*
- *for $n = 3$, there exists a polynomial time algorithm producing an allocation (S_1, S_2, S_3), such that $v_i(S_i) \geqslant (8/9-\epsilon)\boldsymbol{\mu}_i$, for any constant $\epsilon > 0$.*

12.4.2 Proportionality

As soon as we move away from maximin shares towards proportionality, the problems become computationally much harder.

For the concept of proportionality, let us start again with additive approximations. One way to define here such a notion of approximation would be as follows: we can say that an algorithm achieves an additive ρ-approximation if it produces an allocation where every agent i receives a bundle worth at least $\frac{1}{n}v_i(M) - \rho$.

By Theorem 12.7, we directly have the following corollary.

Corollary 12.12. *Algorithm 1 achieves an additive α-approximation for proportionality, where $\alpha = \max_{i,j} v_{ij}$.*

It can be easily seen that this guarantee is tight and cannot be further improved.

In analogy to the approximate versions of maximin share allocations, we will say that an algorithm achieves a multiplicative ρ-approximation for proportionality if it produces a partition $(S_1, S_2, \ldots, S_n)$ such that $v_i(S_i) \geqslant \rho \cdot \frac{1}{n} \cdot v_i(M)$.

In the previous subsection we saw that combining the round-robin algorithm with a careful handling of the most valuable goods, we obtained an approximation ratio of $1/2$ for the maximin share guarantees. Unfortunately, this does not carry over to approximate proportionality. In fact, it has been established that we cannot hope to have any decent approximation. The following result, proved by Markakis and Psomas (2011) shows that we cannot even decide in polynomial time, if an instance admits an allocation that achieves a constant approximation, hence we cannot possibly compute one for every instance. The proof is based on a reduction from the 3D-Matching problem.

Theorem 12.13. *For any constant $\rho \leqslant 1$, it is NP-complete to decide if there exists an allocation where every player receives a bundle worth at least $\rho \cdot \frac{1}{n} \cdot v_i(M)$.*

Theorem 12.13 reveals that proportionality is a much stronger concept to satisfy (even approximately), than maximin share fairness under indivisible goods.

12.4.3 Envy-freeness

Coming now to the strongest among the concepts we have considered so far, one might not be so optimistic about obtaining algorithmic results. Nevertheless, we will see that in certain cases, some positive results are feasible.

In order to define approximate versions, let us start with some notation. Given an allocation $S = (S_1, S_2, \ldots, S_n)$, let $e_{ij}(S) = \max\{0, v_i(S_j) - v_i(S_i)\}$ be the envy experienced by agent i towards agent j. When i does not envy j, we have $e_{ij}(S) = 0$. Let also $envy(S) = \max_{i,j\in N} e_{ij}(S)$ be the envy of the allocation S, i.e., we care for the maximum envy between any pair of agents.

Given an instance I, let $OPT(I)$ denote the minimum possible envy that can be achieved over all possible partitions,

$$OPT(I) = \min_{S\in\Pi_n(M)} envy(S)$$

If we look at the notion of envy from an optimization viewpoint, we can define the minimum envy problem as the problem of finding $OPT(I)$ for every instance I. We can now define approximate versions of our problem. Namely, we will say that an algorithm achieves an additive ρ-approximation for the minimum envy problem, if it produces a partition S such that $envy(S) \leqslant OPT(I) + \rho$. We can now easily have the following observation.

Corollary 12.14. *Any algorithm that is envy-free up to one good achieves an additive α-approximation to the minimum envy problem.*

This is true since any algorithm that is envy-free up to one good produces envy at most equal to α. Hence Algorithm 1 in particular obtains this guarantee,

Suppose now that we try to get a multiplicative guarantee. This would mean that we want an allocation such that the envy within this allocation should be at most $\rho OPT(I)$, at a given instance I. We claim that this particular type of multiplicative approximation is problematic for the minimum envy problem. To illustrate this, consider an instance I that admits an anvy-free solution. Hence $OPT(I) = 0$. An algorithm with a multiplicative ρ-approximation in this case means that it would produce an allocation S such that $envy(S) \leqslant \rho OPT(I) = 0$. Hence the algorithm is forced to compute an envy-free allocation whenever there exists one, i.e., this algorithm can be used to decide if an envy-free allocation exists. But since we showed in Theorem 12.5 that deciding existence of envy-free allocations is NP-complete, we then immediately have the following corrolary:

Theorem 12.15. *Unless* P = NP, *there is no finite multiplicative approximation algorithm for the minimum envy problem.*

The discussion above is an artifact of the nature of the problem. In particular, it is a consequence of the fact that the objective function in the optimization problem we try to solve can take the value zero. This is what prevents us from having a multiplicative approximation. We claim that a more suitable objective function for approximating envy is the envy-ratio, defined as taking the ratio rather than the difference between a pair of bundles. More formally, given an allocation $S = (S_1, S_2, \ldots, S_n)$, the envy-ratio of S is defined as the quantity $\max_{i,j \in N}\{1, \frac{v_i(S_j)}{v_i(S_i)}\}$. The minimum envy-ratio then asks for a partition that achieves the best such ratio.

Obviously this problem is still NP-complete. It is however more amenable to multiplicative approximations. Even though we are not aware of any good positive result for general additive valuation functions, it has been established that when the agents have identical preferences, any constant approximation is achievable. The following was established by Lipton et al. (2004).

Theorem 12.16. *Under identical additive valuation functions, there exists a PTAS for the minimum envy-ratio problem.*

It is still an open problem whether we can have positive results for a richer class of additive valuation functions.

12.5 Conclusions and Future Research

In this chapter, we have provided an overview of algorithmic and hardness results, with an emphasis on approximation algorithms. We believe this will continue to be a very active research area over the coming years, as there are still several unexplored problems that are worth further investigation.

To begin with, there are still some challenging open problems regarding the solution concepts we presented. We find the approximability status for MMS allocations one of the most intriguing questions. Is there a better than a $2/3$-approximation to the maximin share guarantee of each agent? This question seems to require radically new ideas. Another interesting question concerns a stronger version of envy-freeness up to one good. Suppose that we require an allocation that would be envy-free up to any single good, meaning that an agent i does not envy any other agent j, after throwing away *any* item from the bundle of j. Recall that the standard definition of envy-freeness up to one good only requires that there exists a good that we can throw away from the bundle of agent j, so that i does not envy her. Can we always guarantee the existence of this stronger notion or can we disprove that?

Coming to other notions, there has recently been a surge of interest in the Maximum Nash Welfare (MNW) solution, which denotes an allocation that maximizes the product of the agents' valuations. It has been argued in Caragiannis et al. (2016) that the MNW solution provides certain fairness guarantees, i.e., it yields an allocation that is envy-free up to one good and is also an approximate MMS allocation. Computing this solution is a hard problem and approximation algorithms have also been suggested (Cole and Gkatzelis, 2015). Unfortunately, as soon as we move to approximations, we lose the fairness guarantees. Hence, the question of interest is whether one can have new approximation algorithms that can retain some of the fairness properties of the MNW solution.

Another approach on fair division is to combine fairness with incentive compatibility. Typically in the fair division literature, monetary payments are not allowed when studying strategic settings. Hence, this becomes a mechanism design problem without money. It has been quite challenging to understand which are the truthful mechanisms in such a setting. Up until recently, the only characterization result that was known concerned two agents and two items, due to Caragiannis et al. (2009). Lately, some further results have been obtained by Amanatidis et al. (2016), and a complete characterization of truthful mechanisms for two agents and any number of items was provided by Amanatidis et al. (2017). These results imply several consequences on the interplay between fairness and incentive compatibility. Undoubtedly, the main open question here is to understand and characterize truthful mechanisms for more than two agents.

Finally, over the years, the vast majority of the related literature has focused on additive valuation functions. In many cases, this assumption is crucial for obtaining positive results. It is true, however, that in some scenarios additivity is not the right assumption, since the goods may exhibit complementarities or substitutabilities. Very recently, new results for finding MMS allocations under submodular valuation functions have been obtained by Barman and Murthy (2017). We believe that further extensions to non-additive valuation functions would be very valuable for making this field even more credible and applicable.

Bibliography

G. Amanatidis, E. Markakis, A. Nikzad, and A. Saberi. Approximation algorithms for computing maximin share allocations. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP), Track A*, pages 39–51, 2015.

G. Amanatidis, G. Birmpas, and E. Markakis. On truthful mechanisms for maximin share allocations. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 31–37, 2016.

G. Amanatidis, G. Birmpas, G. Christodoulou, and E. Markakis. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 545–562, 2017.

S. Barman and S. K. K. Murthy. Approximation algorithms for maximin fair division. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 647–664, 2017.

S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1321–1328, 2014.

S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair allocation of indivisible goods. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12. Cambridge University Press, 2016.

S. J. Brams and A. D. Taylor. *Fair Division: From Cake Cutting to Dispute Resolution*. Cambridge University Press, 1996.

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

I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. On low-envy truthful allocations. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT)*, pages 111–119, 2009.

I. Caragiannis, D. Kurokawa, H. C. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 305–322, 2016.

R. Cole and V. Gkatzelis. Approximating the Nash social welfare with indivisible items. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, pages 371–380, 2015.

S. Demko and T. Hill. Equitable distribution of indivisible items. *Mathematical Social Sciences*, 16:145–158, 1988.

S. Even and A. Paz. A note on cake cutting. *Discrete Applied Mathematics*, 7: 285–296, 1984.

D. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7: 45–98, 1967.

G. Gamow and M. Stern. *Puzzle-Math*. Viking Press, 1958.

L. Gourvès and J. Monnot. Approximate maximin share allocations in matroids. In *Proceedings of the 10th International Conference on Algorithms and Complexity (CIAC)*, pages 310–321, 2017.

L. Gourvès, J. Monnot, and L. Tlilane. Worst case compromises in matroids with applications to the allocation of indivisible goods. *Theoretical Computer Science*, 589:121–140, 2015.

T. Hill. Partitioning general probability measures. *The Annals of Probability*, 15 (2):804–813, 1987.

D. Kurokawa, A. D. Procaccia, and J. Wang. When can the maximin share guarantee be guaranteed? In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 523–529, 2016.

R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 125–131, 2004.

E. Markakis and C. A. Psomas. On worst-case allocations in the presence of indivisible goods. In *Proceedings of the 7th Workshop on Internet and Network Economics (WINE)*, pages 278–289, 2011.

H. Moulin. Uniform externalities: Two axioms for fair allocation. *Journal of Public Economics*, 43(3):305–326, 1990.

H. Moulin. *Fair Division and Collective Welfare*. MIT Press, Cambridge, Massachusetts, 2003.

A. D. Procaccia. Cake cutting algorithms. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13. Cambridge University Press, 2016.

A. D. Procaccia and J. Wang. Fair enough: guaranteeing approximate maximin shares. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, pages 675–692, 2014.

J. M. Robertson and W. A. Webb. *Cake Cutting Algorithms: Be fair if you can*. AK Peters, 1998.

H. Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.

H. Varian. Equity, envy and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.

G. Woeginger. A polynomial time approximation scheme for maximizing the minimum machine completion time. *Operations Research Letters*, 20:149–154, 1997.