On Truthful Mechanisms for Maximin Share Allocations

Georgios Amanatidis, Georgios Birmpas, Evangelos Markakis
Department of Informatics, Athens University of Economics and Business, Athens, Greece
{gamana, gebirbas, markakis}@aueb.gr

Abstract

We study a fair division problem with indivisible items, namely the computation of maximin share allocations. Given a set of n players, the maximin share of a single player is the best she can guarantee to herself, if she would partition the items in any way she prefers, into n bundles, and then receive her least desirable bundle. The objective then is to find an allocation, so that each player is guaranteed her maximin share.

Previous works have studied this problem purely algorithmically, providing constant factor approximation algorithms. In this work, we embark on a mechanism design approach and investigate the existence of truthful mechanisms. We propose three models regarding the information that the mechanism attempts to elicit from the players, based on the cardinal and ordinal representation of preferences. We establish positive and negative (impossibility) results for each model and highlight the limitations imposed by truthfulness on the approximability of the problem. Finally, we pay particular attention to the case of two players, which already leads to challenging questions.

1 Introduction

We study the design of mechanisms for a fair division problem with indivisible items. The objective in fair division is to allocate a set of resources to a set of players in a way that leaves everyone satisfied, according to their own preferences. Over the past decades, several fairness concepts have been proposed and the area gradually gained popularity in computer science as well, since most of the questions are inherently algorithmic. We refer the reader to the upcoming survey [Procaccia, 2015] for more recent results and to the classic textbooks [Brams and Taylor, 1996; Robertson and Webb, 1998] for an overview of the area.

Our focus here is on the concept of maximin share allocations, which has already attracted a lot of attention, ever since it was introduced by [Budish, 2011]. The rationale for this notion is as follows: suppose that a player, say player i, is asked to partition the items into n bundles and then the rest of the players select a bundle before i. In the worst case, player i will be left with her least valuable subset. Hence, a risk-averse player would choose a partition that maximizes the minimum value of a bundle. This value is called the maximin share of agent i and the goal then is to find an allocation where every person receives at least her maximin share.

The existence of maximin share allocations is not always guaranteed under indivisible items. This has led to a series of works that have either established approximation algorithms (i.e., every player receives an approximation of her own maximin share) or have resolved special cases of the problem; see our Related Work section. Currently, the best algorithms we are aware of, achieve an approximation ratio of $2/3$ [Procaccia and Wang, 2014; Amanatidis *et al.*, 2015], and it is still a challenging open problem if one can do better.

These previous works have studied the problem purely from an algorithmic point of view, and one aspect that has not been addressed so far is incentive compatibility. Players may have incentives to misreport their valuation functions and in fact, the proposed approximation algorithms are not truthful. Is it possible then to have truthful algorithms with the same approximation guarantee? Truthfulness is a demanding constraint especially in settings without monetary transfers, as is the case here, and our goal is to explore the effects on the approximability of the problem if we impose such a constraint.

Contribution. We investigate the existence of truthful mechanisms for constructing approximate or exact maximin share allocations. In doing so, we consider three models regarding the information that the mechanism attempts to elicit from the players. The first one is the more straightforward approach where players have to submit their entire additive valuation function to the mechanism. We then move to mechanisms where the manipulating power of the players is restricted by the type of information that they are allowed to submit. Namely in our second model, players only submit their ranking over the items, motivated by the fact that many mechanisms in the fair division literature fall within this class. Finally, in our third model, we assume the mechanism designer knows the ranking of each player over the items and asks for a valuation function consistent with the ranking. This can be appropriate for settings where the items are distinct enough to be able to extract a ranking, or when the players are known to belong to specific behavioral types. For each of these models, we establish positive and negative (impossibility) results and highlight the differences and similarities between them. Our results provide a clear separation between the guarantees achievable by truthful and non-truthful mechanisms. As an example, it is known that for two players there is a non-truthful PTAS, whereas we establish that this is not the case for truthful algorithms in any of our models. We also note that all our positive results yield polynomial time algorithms, whereas the impossibility results are independent of the running time of an algorithm. Finally, we pay particular attention to the case of two players, which already gives rise to non-trivial questions, even with a small number of items.

Related Work. The notion of maximin share allocations was introduced by [Budish, 2011] (building on concepts of [Moulin, 1990]), and later on defined by [Bouveret and Lemaître, 2014] in the setting that we study here. Both experimental and theoretical evidence, see [Bouveret and Lemaître, 2014; Kurokawa *et al.*, 2016; Amanatidis *et al.*, 2015], indicate that such allocations do exist almost always. As for computation, a $2/3$ approximation algorithm has been established in [Procaccia and Wang, 2014] and later on, a polynomial time algorithm with the same guarantee was provided in [Amanatidis *et al.*, 2015].

Regarding incentive compatibility, we are not aware of any prior work that addresses the design of truthful mechanisms for maximin share allocations. There have been quite a few works on mechanisms for other fairness notions, see among others [Caragiannis *et al.*, 2009; Chen *et al.*, 2013; Lipton *et al.*, 2004]. Parts of our work are motivated by the question of what is the power of cardinal information vs ordinal information. We note that exploring what can be done using only ordinal information has been recently studied for other optimization problems too, (see [Anshelevich and Sekar, 2016]). A popular class of mechanisms based only on ordinal preferences is the class induced by "picking sequences", introduced by [Kohler and Chandrasekaran, 1971]. We make use of such algorithms to establish some of our positive results.

2 Preliminaries

For any $k \in \mathbb{N}$, we denote by $[k]$ the set $\{1, ..., k\}$. Let $N = [n]$ be a set of n players and $M = [m]$ be a set of indivisible items. We assume each player has an additive valuation function $v_i(\cdot)$ over the items, and we will write v_{ij} instead of $v_i(\{j\})$. For $S \subseteq M$, we let $v_i(S) = \sum_{j \in S} v_{ij}$. An allocation of M to the n players is a partition, $T = (T_1, ..., T_n)$, where $T_i \cap T_j = \emptyset$ and $\bigcup_i T_i = M$. Let $\Pi_n(M)$ be the set of all partitions of a set M into n bundles.

Definition 2.1. *Given n players, and a set M of items, the n-maximin share of a player i with respect to M is:*

$$\boldsymbol{\mu}_i(n, M) = \max_{T \in \Pi_n(M)} \min_{T_j \in T} v_i(T_j) \,.$$

When it is clear from context what n, M are, we will simply write $\boldsymbol{\mu}_i$ instead of $\boldsymbol{\mu}_i(n, M)$. The solution concept defined in [Budish, 2011] asks for a partition that gives each player her maximin share.

Definition 2.2. *Given n players and a set of items M, a partition $T = (T_1, ..., T_n) \in \Pi_n(M)$ is called a maximin share allocation if $v_i(T_i) \geq \boldsymbol{\mu}_i$, for every $i \in [n]$. If $v_i(T_i) \geq \rho \cdot \boldsymbol{\mu}_i$, $\forall i \in [n]$, with $\rho \leq 1$, then T is called a ρ-approximate maximin share allocation.*

It can be easily seen that this is a relaxation of the classic notion of proportionality.

Example 1. Consider an instance with 3 players and 5 items:

	a	b	c	d	e
Player 1	$1/2$	$1/2$	$1/3$	$1/3$	$1/3$
Player 2	$1/2$	$1/4$	$1/4$	$1/4$	0
Player 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 $\boldsymbol{\mu}_1(3, M) = 1/2$, $\boldsymbol{\mu}_2(3, M) = 1/4$, $\boldsymbol{\mu}_3(3, M) = 1$. For player 1, no matter how she partitions the items into three bundles, the worst bundle will be worth at most $1/2$ for her. Similarly, player 3 can guarantee a value of 1 (which is best possible as it is equal to $v_3(M)/n$) by the partition $(\{a, b\}, \{c\}, \{d, e\})$. Note that this instance admits a maximin share allocation, e.g., $(\{a\}, \{b, c\}, \{d, e\})$, and in fact this is not unique.

Note also that if we remove some player, say player 2, the maximin values for the other two players increase. E.g., $\boldsymbol{\mu}_1(2, M) = 1$, achieved by the partition $(\{a, b\}, \{c, d, e\})$. Similarly, $\boldsymbol{\mu}_3(2, M) = 3/2$.

2.1 Mechanism Design Aspects

Following most of the fair division literature, our focus is on mechanism design without money, i.e., we do not allow side payments to the players. The standard way to define truthfulness then, is as follows: an instance of the problem can be described as an $n \times m$ valuation matrix $V = [v_{ij}]$, as in Example 1 above. For any mechanism A, we denote by $A(V) = (A_1(V), \ldots, A_n(V))$ the allocation output by A on input V. Also, let V_i denote the ith row of V, and V_{-i} denote the remaining matrix. Finally, let (V_i', V_{-i}) be the matrix we get by changing the ith row of V from V_i to V_i'.

Definition 2.3. *A mechanism A is truthful if for any instance V, any player i, and any possible declaration V_i' of i: $v_i(A_i(V)) \geq v_i(A_i(V_i', V_{-i}))$.*

Obtaining a good understanding of truthful mechanisms and their performance for other fairness notions has been a difficult problem; see among others [Lipton *et al.*, 2004; Caragiannis *et al.*, 2009] for approximating minimum envy with truthful mechanisms. The difficulty is that an algorithm that uses an m-dimensional vector of values for each player, can create many subtle ways for players to benefit by misreporting. One can try to alleviate this by restricting the type of information that is requested from the players. As a first instantiation of this, we note that many mechanisms in the literature end up utilizing only the ranking of each player for the items, and not the entire valuation function, (see our discussion in Section 4 and references therein). This yields simpler, intuitive mechanisms, at the expense of possibly sacrificing performance, since the mechanism uses less information. As a second instantiation, one can exploit information that could be available to the mechanism so as to restrict the allowed valuations. For example, in some scenarios, it is realistic to assume that the ranking of each player over the items is public knowledge. If the items are distinct enough, it is possible that one could extract such information (a special case is that of full correlation, considered in [Brams and Fishburn, 2002; Bouveret and Lang, 2011], where all players have the same ranking). Therefore, the players in such cases can only submit values that agree with their (known) ranking.

Motivated by the above considerations, we study the following three models:

- The Cardinal or Standard Model. Every player submits a valuation function, without any restrictions. To represent the input of player i, we fix an ordering of the items and write the corresponding vector of values as $\mathbf{v}_i = [v_{i1}, v_{i2}, \ldots, v_{im}]$.
- The Ordinal Model. Here, an instance is again determined by a matrix V, however a mechanism only asks players to submit a ranking on the items. Note that Definition 2.3 of truthfulness needs to be modified accordingly. That is, let $\succeq_i$ be any total order consistent with $\mathbf{v}_i$ (there may be many in case of ties). A mechanism is truthful if for any tuple of rankings $\succeq_{-i}$, for the other players, and any ranking $\succeq_i'$: $v_i(A_i(\succeq_1, ..., \succeq_n)) \geq v_i(A_i(\succeq_i', \succeq_{-i}))$.
- The Public Rankings Model. Now, the ranking of each player is known to the mechanism, say it is $\succeq_i$. Hence, each player is asked to submit a valuation function consistent with $\succeq_i$.

It is not hard to see how the different scenarios we investigate are related to each other; this is summarized in the following lemma. The proof is omitted due to space constraints.

Lemma 2.4. (i) *Assume there exists a truthful ρ-approximation mechanism A in the cardinal model. Then, A can be efficiently turned into a truthful ρ-approximation mechanism for the public rankings model.*
(ii) *Assume there exists a truthful ρ-approximation mechanism A for the ordinal model. Then, A can be efficiently turned into a truthful ρ-approximation mechanism for the cardinal model.*

3 The Cardinal Model

As already alluded to, designing mechanisms that utilize the values submitted by each player, so as to achieve a good approximation and at the same time induce truthful behavior, is a very challenging problem. This is true even in the case of $n = 2$ players. Therefore, we start first with a rather weak result for general n and m, and then move on to discuss the case of two players. The main message from this section (Theorem 3.3) is that there is a clear separation, regarding the approximation guarantees of truthful and non-truthful algorithms.

Theorem 3.1. *For any $n \geq 2, m \geq 1$, there is a truthful $1/\left\lfloor \frac{\max\{2, m-n+2\}}{2} \right\rfloor$-approximation mechanism for the cardinal model.*

The proof follows from results in the next section (see the discussion before and after Theorem 4.1). For the case of two players, the mechanism of Theorem 3.1 has the following form:

Mechanism M: Given the reported valuations of the players, allocate to player 1 her best item, and to player 2 the remaining items.

Although the approximation ratio achieved by Theorem 3.1 is quite small, it is still an open question whether there exist better mechanisms for general n, m. We note also that the mechanism of Theorem 3.1 only utilizes the preference rankings of the players. Hence it is not even clear if there exist truthful mechanisms that can exploit more information from the valuation functions to achieve a better approximation.

For the remainder of this section, we discuss the case of $n = 2$. We recall that for two players, the discretized cut and choose procedure is a non-truthful algorithm that produces an exact maximin share allocation; one player partitions the goods into two bundles that are as equal as possible, and the other player chooses her best bundle. To implement this in polynomial time, we can produce an approximate partitioning using the result of [Woeginger, 1997] and then we can guarantee at least $(1 - \varepsilon)\boldsymbol{\mu}_i$ to each player, $\forall \varepsilon > 0$. This is not a truthful algorithm however, since player 1 can manipulate the partitioning; in fact, she can compute her optimal strategy if she knows the valuations of player 2 by solving a Knapsack instance. Thus, the question we would like to resolve is to find the best truthful approximation that we can guarantee for two players.

Notice that for $n = 2, m < 4$, the mechanism M does output an exact maximin share allocation. Further, when $m \in \{4, 5\}$, M outputs a $1/2$-approximation, according to Theorem 3.1.

On the other hand, we can deduce an impossibility result, using Theorem 5 of [Markakis and Psomas, 2011], which yields[1]:

Corollary 3.2 (implied by [Markakis and Psomas, 2011])**.** *For $n = 2, m \geq 4$, and any $\varepsilon \in (0, 1/3]$, there is no truthful $(2/3 + \varepsilon)$-approximation mechanism for the cardinal model.*

The above corollary leaves open whether there exist better mechanisms than M with approximation guarantees in $(1/2, 2/3]$. Our main result in this section is that we close this gap, by providing a stronger negative result, which shows that the mechanism M above is optimal for $n = 2, m = 4$.

Theorem 3.3. *For $n = 2, m \geq 4$, and for any $\varepsilon \in (0, 1/2]$, there is no truthful $(1/2 + \varepsilon)$-approximation mechanism for the cardinal model.*

We prove the theorem for $m = 4$, since by adding dummy items of no value, we can trivially extend it to any number of items. The proof follows from Lemmas 3.4 and 3.5 below. Notice that the theorem is valid even if we drastically restrict the possible values of the items.

Lemma 3.4. *For $n = 2, m = 4$, and for any $\varepsilon \in (0, 1/2]$, there is no truthful $(1/2 + \varepsilon)$-approximation mechanism for the cardinal model that allocates two items to each player at every instance where the profiles are permutations of $\{2 + \varepsilon, 1 + \varepsilon, 1 - \varepsilon, \varepsilon/2\}$ or $\{2 - \varepsilon, 1 + \varepsilon, 1 - \varepsilon, \varepsilon/2\}$.*

[1] The work of [Markakis and Psomas, 2011] concerns a different problem however the arguments for their impossibility result can be employed here as well.

The proof of Lemma 3.4 is omitted due to lack of space. It is based on constructing a series of preference profiles and then reaching a contradiction by exploiting the constraints imposed by truthfulness. Using Lemma 3.4, we deduce that there must exist some instance where the mechanism must allocate one item to one player and three items to the other. We prove below that this is not possible either.

Lemma 3.5. *For $n = 2, m = 4$, and for any $\varepsilon \in (0, 1/2]$, there is no truthful $(1/2 + \varepsilon)$-approximation mechanism for the cardinal model, which at some instance where the profiles are permutations of $\{2 + \varepsilon, 1 + \varepsilon, 1 - \varepsilon, \varepsilon/2\}$ or $\{2 - \varepsilon, 1 + \varepsilon, 1 - \varepsilon, \varepsilon/2\}$, allocates exactly one item to one of the players.*

Proof. Let us first fix an ordering of the four items, say a, b, c, d. For the sake of readability we write $2^+, 2^-, 1^+, 1^-$, 0^+ instead of $2 + \varepsilon, 2 - \varepsilon, 1 + \varepsilon, 1 - \varepsilon$ and $\varepsilon/2$.

Suppose that there is such a truthful mechanism, and an instance $\{[v_{1a}, v_{1b}, v_{1c}, v_{1d}], [v_{2a}, v_{2b}, v_{2c}, v_{2d}]\}$ (that we refer to as the *initial profile*), where the mechanism gives one item to p_1 and three items to p_2 (the symmetric case can be handled accordingly). Since we want allocations that give to each player items of value at least $1/2 + \varepsilon$ of their maximin share, it is trivial to check that allocating $\{1^+, 0^+\}$ or $\{1^-, 0^+\}$ to a player is not feasible (we use this repeatedly below).

Recall that the values of each player are a permutation of either $\{2^+, 1^+, 1^-, 0^+\}$ or $\{2^-, 1^+, 1^-, 0^+\}$. Since p_1 gets only one item, its value must be 2^+ or 2^-. W.l.o.g. we may assume that this item is a, so the produced allocation is $(\{a\}, \{b, c, d\})$. We will now construct a chain of profiles (Profiles 1–4) which will help us establish a contradiction.

Profile 1: $\{[v_{1a}, v_{1b}, v_{1c}, v_{1d}], [2^-, v_{1b}, v_{1c}, v_{1d}]\}$. It is easy to see that p_2 can not get just item a, or item a and the item that has value 0^+, or any proper subset of $\{b, c, d\}$, since she could then play $[v_{2a}, v_{2b}, v_{2c}, v_{2d}]$ as in the initial profile, and end up strictly better. Moreover, p_2 cannot get a bundle that contains a and (at least) one item with value 1^- or 1^+, because then there is not enough value left for p_1. Thus, the only feasible allocation here is $(\{a\}, \{b, c, d\})$. W.l.o.g., by possibly renaming items b, c, d, we take Profile 1 to be $\{[v_{1a}, 1^+, 1^-, 0^+], [2^-, 1^+, 1^-, 0^+]\}$.

Profile 2: $\{[v_{1a}, 1^+, 1^-, 0^+], [0^+, 2^-, 1^+, 1^-]\}$. It is easy to notice that in any feasible allocation other than $(\{a\}, \{b, c, d\})$, p_2 could play $\mathbf{v}'_2 = [2^-, 1^+, 1^-, 0^+]$ as in Profile 1, and end up with a better value. Thus, the mechanism has to output $(\{a\}, \{b, c, d\})$ at Profile 2.

Profile 3: $\{[1^-, v_{1a}, 0^+, 1^-], [0^+, 2^-, 1^+, 1^-]\}$. Here, p_1 cannot get a proper superset of $\{a\}$, since then at Profile 1, she could have played $\mathbf{v}'_1 = [1^-, v_{1a}, 0^+, 1^-]$ like here, and end up with strictly more items. The only other feasible allocation here is $(\{b\}, \{a, c, d\})$.

Profile 4: $\{[1^-, v_{1a}, 0^+, 1^-], [2^-, 1^+, 1^-, 0^+]\}$. Here, p_2 cannot get $\{b, c\}$ or any proper subset of $\{a, c, d\}$, since she could then play $\mathbf{v}'_2 = [0^+, 2^-, 1^+, 1^-]$ like in Profile 3, and end up with a total value of $3 - 3\varepsilon/2$, which is strictly better. The only other feasible allocation here is $(\{b\}, \{a, c, d\})$.

By starting now at Profile 2 and repeating the arguments for Profiles 1, 2, and 3 –shifted one position to the right– we have that for *Profile 5:* $\{[1^-, v_{1a}, 0^+, 1^+], [1^-, 0^+, 2^-, 1^+]\}$ the only possible allocation is $(\{b\}, \{a, c, d\})$, and for *Profile 6:* $\{[1^+, 1^-, v_{1a}, 0^+], [1^-, 0^+, 2^-, 1^+]\}$ the only possible allocation is $(\{c\}, \{a, b, d\})$.

Profile 7: $\{[1^+, 1^-, v_{1a}, 0^+], [2^-, 1^+, 1^-, 0^+]\}$. Here, p_2 cannot receive $\{b, c\}$ or any proper subset of $\{a, c, d\}$, since she could then play $\mathbf{v}'_2 = [1^-, 0^+, 2^-, 1^+]$ as in Profile 6, and be better off. The only other feasible allocation is $(\{c\}, \{a, b, d\})$.

Final profile: $\{[1, 1, 1, 1], [2^-, 1^+, 1^-, 0^+]\}$. Here, any feasible allocation has to give p_1 at least two items, otherwise it is not a $(1/2 + \varepsilon)$-approximation. However, one can check that for any such allocation, there is a profile among Profiles 1, 4 and 7, where p_1 could play $\mathbf{v}'_1 = [1, 1, 1, 1]$ and end up strictly better. Thus, we conclude that there are no possible allocations here, arriving at a contradiction. □

This concludes the proof of Theorem 3.3.

4 The Ordinal Model

Several works in the fair division literature have proposed mechanisms that only ask for the ordinal preferences of the players. There are various reasons for such assumptions; apart from their simplicity in implementing them, the players themselves may feel more at ease as they may be reluctant to fully reveal their valuation. Here, one extra motive is to restrict the players' ability to manipulate the outcome.

A class of such simple and intuitive mechanisms is the class of *picking sequence mechanisms*, see, e.g., [Kohler and Chandrasekaran, 1971; Brams and King, 2005; Brams and Taylor, 2000; Bouveret and Lang, 2011; 2014; Kalinowski *et al.*, 2013a; 2013b] and references therein. A picking sequence $\pi = p_{i_1} p_{i_2} \dots p_{i_k}$ is simply a sequence of players (possibly with repetitions). Each picking sequence, naturally induces a deterministic allocation mechanism for the ordinal model as follows: first give to player p_{i_1} her favorite item, then give to p_{i_2} her favorite among the remaining items, and so on, and keep cycling through π until all the items are allocated. Sometimes, periodicity is absent, and then the length of the given sequence is at least m. Notice that these mechanisms can be implemented by asking each player for her ranking over the items. And note also that these mechanisms are not generally truthful, unless they are sequential dictatorships, i.e., they are induced by picking sequences of the form $p_{i_1}^{m_1} p_{i_2}^{m_2} \dots p_{i_k}^{m_k}$, where $p_{i_1}, p_{i_2}, \dots, p_{i_k}$ are all different players and $\sum_i m_i = m$ (see [Bouveret and Lang, 2014]).

Given a set of n players $p_1, \dots, p_n$, we now define the mechanism $M(n, m)$ induced by the picking sequence $\pi = p_1 p_2 p_3 \dots p_{n-2} p_{n-1} p_n p_n \dots p_n$. Thus, the first $n - 1$ players receive exactly one item, and the last player receives the rest $m - n + 1$ items. This is a truthful mechanism, given the observation above. It is easy to see that if $m \leq n + 1$ then $M(n, m)$ constructs an exact maximin share allocation. For large values of m, however, the approximation deteriorates fast and we also have a strong impossibility result.

Theorem 4.1. *The mechanism $M(n, m)$ defined above, is a truthful $1/\left\lfloor \frac{m-n+2}{2} \right\rfloor$-approximation for the ordinal model, for any $n \geq 2, m \geq n + 2$. Moreover, there is no truthful* *mechanism for the ordinal model, induced by some picking sequence, that achieves a better approximation factor.*

The proof is omitted due to lack of space. Notice that Theorem 4.1 combined with Lemma 2.4(ii) imply Theorem 3.1.

For $n = 2$, the mechanism $M(2, m)$ is identical to mechanism M defined in Section 3. Hence, as already pointed out there, this mechanism achieves a $1/2$-approximation for $m \in \{4, 5\}$. We can now combine the impossibility result of Theorem 3.3 and Lemma 2.4(ii) to conclude that $M(2, m)$ is optimal for the ordinal model for $m \in \{4, 5\}$.

Corollary 4.2. *For $n = 2, m \geq 4$, and for any $\varepsilon \in (0, 1/2]$, there is no truthful $(1/2 + \varepsilon)$-approximation mechanism for the ordinal model.*

The impossibility results of Theorem 3.3 and Corollary 4.2 have a surprising consequence. The mechanism $M(2, m)$ achieves the best possible approximation both for the cardinal and the ordinal model, for $m \in \{4, 5\}$. Thus, providing access to more information for the mechanism does not improve the approximation factor at all, when truthfulness is required!

We conclude this section with a general result on the limitations of the ordinal model. Apart from the truthfulness requirement, an additional issue here is the lack of information itself. Below, we prove an inapproximability result for any mechanism in the ordinal model, independent of whether it is truthful or not.

Theorem 4.3. *For $n \geq 2$, and for any $\varepsilon > 0$, there is no $(1/H_n + \varepsilon)$-approximation algorithm, be it truthful or not, for the ordinal model, where H_n is the n^{th} harmonic number, with $H_n = \Theta(\ln n)$. Moreover, for $n = 3$, there is no $(1/2 + \varepsilon)$-approximation algorithm for the ordinal model.*

Proof. Let A be an α-approximation algorithm for the ordinal model, where $\alpha > 0$. Consider an instance with large enough m, where all the players agree on the ranking $1 \succeq 2 \succeq \ldots \succeq m$. Let g_i be the best item that player i receives by A. We renumber the players, if needed, so that if $i < j$ then $g_i < g_j$. We claim that $g_i = i$. To see that, consider player n. Clearly, by the definition of g_n and the renumbering of the players, we have $g_n \geq n$. If $g_n > n$, let $v_{n1} = \ldots = v_{nn} = 1$ and $v_{n,n+1} = \ldots = v_{nm} = 0$. Then, in such an instance, algorithm A will fail to give an α-approximation of $\boldsymbol{\mu}_n$ to player n. It follows that $g_n = n$, and therefore $1 = g_1 < g_2 < \ldots < g_{n-1} < n$, which implies $g_i = i$ for every $i \in [n]$.

Now, for $i \geq 1$, suppose that $v_{i1} = \ldots = v_{i,i-1} = 1$ and $v_{ii} = \ldots = v_{im} = \frac{1}{m-i+1}$. Observe that $\boldsymbol{\mu}_i = \lfloor \frac{m-i+1}{n-i+1} \rfloor \frac{1}{m-i+1}$, and A must give at least $\lceil \alpha \lfloor \frac{m-i+1}{n-i+1} \rfloor \rceil$ items to player i. Since there are m items in total, we must have $\sum_{i=1}^{n} \lceil \alpha \lfloor \frac{m-i+1}{n-i+1} \rfloor \rceil \leq m$. It follows that for any $\varepsilon > 0$ and large enough m

$$\alpha \leq \frac{m}{\sum_{i=1}^{n} \lfloor \frac{m-i+1}{n-i+1} \rfloor} \leq \frac{m}{\sum_{i=1}^{n} \left(\frac{m-i+1}{n-i+1} - 1 \right)}$$
$$= \frac{1}{\left(1 - \frac{n}{m}\right) \sum_{i=1}^{n} \frac{1}{n-i+1}} < \frac{1}{H_n} + \varepsilon \, .$$

Especially for $n = 3$, assume that $\alpha > 1/2$ and consider the same analysis as above with $m = 6$. We get the contradiction

$$6 \geq \sum_{i=1}^{3} \lceil \alpha \lfloor \tfrac{7-i}{4-i} \rfloor \rceil = \lceil \alpha \lfloor \tfrac{6}{3} \rfloor \rceil + \lceil \alpha \lfloor \tfrac{5}{2} \rfloor \rceil + \lceil \alpha \lfloor \tfrac{4}{1} \rfloor \rceil$$
$$\geq 2 + 2 + 3 = 7 \, .$$ □

5 The Public Rankings Model

When the players' rankings are publicly known, one would expect to achieve better approximation ratios, while still maintaining truthfulness. Indeed, the mechanism now has more information, while the options for manipulation are greatly reduced. In particular, note that any picking sequence induces a truthful mechanism for the public rankings model.

We show that indeed this is the case; the impossibility results we obtain are less severe and we have improvements for the case of more than two players as well.

We focus first on two players. For $m < 4$, the mechanism $M(2, m)$ from Section 4 gives an exact solution, like before. However, unlike what happens in the other two scenarios, for $m = 4$ we now have a truthful exact mechanism. Before we describe the mechanism, we introduce some useful notation. For a player i, we will denote by $B_i(k_1, k_2, \ldots k_\ell)$ the set of items that are in the positions $k_1, k_2, \ldots k_\ell$, of her ranking. E.g., $B_2(2, 4)$ denotes the bundle that contains the second and the fourth items in the ranking of player 2.

Mechanism M_1: Given the reported valuations of the two players p_1, p_2, and their actual rankings, consider two cases:
–If their most valuable items are different, allocate the items according to the picking sequence $p_1 p_2 p_2 p_1$.
–Otherwise, give to player 1 her most valuable bundle among $B_1(1)$ and $B_1(2, 3)$, and to player 2 the remaining items.

Theorem 5.1. *Mechanism M_1, is truthful and produces an exact maximin share allocation for the public rankings model, for $n = 2, m = 4$.*

Proof. To see why M_1 is truthful, note that the players cannot affect which of the two cases of M_1 will be employed, since this is defined by the publicly known rankings. In addition, only p_1 could strategize, in the case where she agrees with p_2 on the most valuable item. However, in that case M_1 gives her the best bundle between two choices defined by her ranking, thus there is no incentive to lie about her true values.

To prove now the guarantee for the maximin share, observe that when the two players disagree on their most valuable item, p_1 receives one of $B_1(1, 2)$, $B_1(1, 3)$, or $B_1(1, 4)$, and p_2 receives either $B_2(1, 2)$, or $B_2(1, 3)$. Similarly, when they agree on their most valuable item, p_1 receives her best bundle among $B_1(1)$ and $B_1(2, 3)$, and p_2 receives either a bundle of three items, or one of $B_2(1, 2)$, $B_2(1, 3)$, or $B_2(1, 4)$.

Consider the seven possible ways p_i can split the four items into non-empty bundles: $(B_i(1,2), B_i(3,4))$, $(B_i(1,3), B_i(2,4))$, $(B_i(1,4), B_i(2,3))$, $(B_i(1), B_i(2,3,4))$, $(B_i(2), B_i(1,3,4))$, $(B_i(3), B_i(1,2,4))$ and $(B_i(4), B_i(1,2,3))$. From the definition of maximin share, in at least one of those, both bundles have value at least $\boldsymbol{\mu}_i$.

It is easy to see that the total value of $B_i(1, 3)$ (and thus of $B_i(1, 2)$), is always at least $\boldsymbol{\mu}_i$, and the same holds for any bundle that contains three items. Moreover, we claim that both $v_i(B_i(1, 4))$ and $\max\{v_i(B_i(1)), v_i(B_i(2, 3))\}$ are at least $\boldsymbol{\mu}_i$, which suffice to prove the theorem. Indeed, if $v_i(B_i(1, 4)) < \boldsymbol{\mu}_i$ or $\max\{v_i(B_i(1)), v_i(B_i(2, 3))\} < \boldsymbol{\mu}_i$, this implies that each one of $B_i(1)$, $B_i(2)$, $B_i(3)$, $B_i(4)$, $B_i(2,3)$, $B_i(2,4)$, and $B_i(3,4)$ also has value less than $\boldsymbol{\mu}_i$. Thus, none of the possible partitions has both bundles worth at least $\boldsymbol{\mu}_i$, contradicting the definition of maximin share. □

An interesting question is whether the above can be extended for any number of items. We exhibit below that the answer is no, hence non-truthful algorithms have a strictly better performance under this model as well. However for general m, we provide later on an improved approximation in comparison to the other two settings.

Theorem 5.2. *For $n = 2$, and $m = 5$, there is no truthful $(5/6 + \varepsilon)$-approximation mechanism for any $\varepsilon \in (0, 1/6]$, while for $m \geq 6$, there is no truthful $(4/5+\varepsilon)$-approximation mechanism for any $\varepsilon \in (0, 1/5]$.*

Proof. We give the proof for $m = 6$, which can be extended to $m \geq 6$, by adding dummy items of no value. The proof for $m = 5$ is of similar flavor, albeit more complicated.

Suppose that there is a deterministic truthful mechanism for the public rankings model, that achieves a $(4/5 + \varepsilon)$-approximation for some $\varepsilon > 0$. We study five profiles where the ranking of the six items is $a \succeq_i b \succeq_i c \succeq_i d \succeq_i e \succeq_i f$ for $i \in \{1, 2\}$, thus it is feasible for both players to move between these profiles in order to increase the value they get. Recall that in our current model, a player can strategize using the values of the items, but without changing their publicly known ranking.

Profile 1: $\{[1, 1, 1, 1, 1, 1], [1, 1, 1, 1, 1, 1]\}$. Here, $\boldsymbol{\mu}_i = 3$ for $i \in \{1, 2\}$, so in order to achieve a better than a 0.8-approximation, the mechanism must give to each player items of value greater than $0.8 \cdot \boldsymbol{\mu}_i = 2.4$. Thus each player has to receive three items. W.l.o.g. we may assume that p_1 gets item a (the symmetric case is handled in an analogous manner).

Profile 2: $\{[1, 0.2, 0.2, 0.2, 0.2, 0.2], [1, 1, 1, 1, 1, 1]\}$. Here, $\boldsymbol{\mu}_1 = 1$ and $\boldsymbol{\mu}_2 = 3$. The mechanism must give to p_1 a total value greater than $0.8 \cdot 1 = 0.8$ and to p_2 a total value greater than $0.8 \cdot 3 = 2.4$. Notice, now, that p_2 has to get at least three items, and therefore p_1 has to get a superset of $\{a\}$. In fact, p_1 gets a superset of $\{a\}$ of size three, otherwise she could play $\mathbf{v}'_1 = [1, 1, 1, 1, 1, 1]$ like in Profile 1, and end up strictly better. So, we conclude that both players get three items each, and p_1 gets item a.

Profile 3: $\{[1, 0.2, 0.2, 0.2, 0.2, 0.2], [1, 0.2, 0.2, 0.2, 0.2, 0.2]\}$. Here, $\boldsymbol{\mu}_i = 1$ for $i \in \{1, 2\}$, so in order to achieve something strictly greater than $0.8 \cdot 1 = 0.8$, there are only two feasible allocations: i) $(\{b, c, d, e, f\}, \{a\})$, and ii) $(\{a\}, \{b, c, d, e, f\})$. Now, notice that allocation ii) is not possible, since then at Profile 2, p_2 could play $\mathbf{v}'_2 = [1, 0.2, 0.2, 0.2, 0.2, 0.2]$ like here, and end up strictly better. Thus, the mechanism outputs $(\{b, c, d, e, f\}, \{a\})$.

Profile 4: $\{[1, 1, 1, 1, 1, 1], [1, 0.2, 0.2, 0.2, 0.2, 0.2]\}$. Here, $\boldsymbol{\mu}_1 = 3$ and $\boldsymbol{\mu}_2 = 1$. The mechanism must give to p_1 a total value greater than $0.8 \cdot 3 = 2.4$ and to p_2 a total value greater than $0.8 \cdot 1 = 0.8$. Notice now that p_1 has to get five items, since otherwise she could play $\mathbf{v}'_1 = [1, 0.2, 0.2, 0.2, 0.2, 0.2]$ like in Profile 2, and end up strictly better. Thus p_2 has to get $\{a\}$ to achieve the desired ratio.

Profile 5: $\{[1, 1, 1, 1, 1, 1], [0.7, 0.3, 0.25, 0.25, 0.25, 0.25]\}$. Here, $\boldsymbol{\mu}_1 = 3$ and $\boldsymbol{\mu}_2 = 1$. The mechanism must give to p_1 a total value greater than $0.8 \cdot 3 = 2.4$ and to p_2 a total value greater than $0.8 \cdot 1 = 0.8$. First, notice that p_1 must get at least three items. Moreover, if the mechanism does not give item a to p_2, then there is no way for p_2 to get total value strictly greater than 0.8 with at most three items. Therefore, p_2 has to get a strict superset of $\{a\}$. However, this is not feasible either, since at Profile 4, p_2 could play $\mathbf{v}'_2 = [0.7, 0.3, 0.25, 0.25, 0.25, 0.25]$ like here, and end up strictly better. Thus, we conclude that there are no possible allocations here, arriving at a contradiction. □

Exploiting the fact that picking sequences induce truthful mechanisms for the public rankings model, we can get more positive results for two players and any m. Let M_2 be the mechanism for two players induced by the picking sequence $p_1p_2p_2$. We have the following result for M_2.

Theorem 5.3. *For $n = 2$ and any $m \geq 1$, M_2 is a truthful $2/3$-approximation mechanism for the public rankings model.*

Hence, for $n = 2$, we have a pretty clear picture on what we can achieve for any m, leaving only a small gap, i.e., $[2/3, 4/5]$, between the impossibility result and Theorem 5.3.

We can also obtain constant factor approximations for more than two players, which has been elusive in the other two models. E.g., for $n = 3$, we can achieve a $1/2$-approximation. In particular, for any $n \geq 2$, and $m \geq 1$, let M_n be the mechanism induced by the picking sequence $p_1p_2p_3 \ldots p_{n-1}p_np_n$.

Theorem 5.4. *For any $n \geq 2$, and any $m \geq 1$, the mechanism M_n is a truthful $\frac{2}{n+1}$-approximation mechanism for the public rankings model. E.g., for $n = 3$, this yields a $1/2$-approximation.*

Note that Theorem 5.3 is a corollary of Theorem 5.4. Also, observe that $\frac{2}{n+1}$ is better than the guarantee of Theorem 4.1, however, it remains an open problem to design truthful mechanisms achieving better, or even no, dependence on n.

6 Conclusions

We embarked on the existence of truthful mechanisms for approximate maximin share allocations. In doing so, we considered three models regarding the information that a mechanism elicits from the players, and studied their power and limitations. Quite surprisingly, we have exhibited cases with two players, where the best possible truthful approximation is achieved by using only ordinal information.

Our work leaves several interesting questions for future research. A great open problem is whether there exist better truthful mechanisms in the cardinal model, that explicitly take into account the players' valuation functions rather than just ordinal information. Another more general question is to tighten the upper and lower bounds obtained here; especially for a large number of players, these bounds are quite loose.

Acknowledgments

Research supported by an internal research funding program of the Athens University of Economics and Business.

References

[Amanatidis *et al.*, 2015] G. Amanatidis, E. Markakis, A. Nikzad, and A. Saberi. Approximation algorithms for computing maximin share allocations. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Proceedings, Part I*, pages 39–51, 2015.

[Anshelevich and Sekar, 2016] E. Anshelevich and S. Sekar. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In *30th AAAI Conference on Artificial Intelligence, AAAI 2016,*, 2016.

[Bouveret and Lang, 2011] S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI 2011*, pages 73–78, 2011.

[Bouveret and Lang, 2014] S. Bouveret and J. Lang. Manipulating picking sequences. In *ECAI 2014 - 21st European Conference on Artificial Intelligence*, pages 141–146, 2014.

[Bouveret and Lemaître, 2014] S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '14*, pages 1321–1328, 2014.

[Brams and Fishburn, 2002] S. J. Brams and P. Fishburn. Fair division of indivisible items between two people with identical preferences. *Social Choice and Welfare*, 17:247–267, 2002.

[Brams and King, 2005] S. J. Brams and D. King. Efficient fair division - help the worst off or avoid envy. *Rationality and Society*, 17(4):387–421, 2005.

[Brams and Taylor, 1996] S. J. Brams and A. D. Taylor. *Fair Division: from Cake Cutting to Dispute Resolution*. Cambridge University press, 1996.

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

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

[Caragiannis *et al.*, 2009] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. On low-envy truthful allocations. In *First International Conference on Algorithmic Decision Theory, ADT 2009*, pages 111–119, 2009.

[Chen *et al.*, 2013] Y. Chen, J. K. Lai, D. C. Parkes, and A. D. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1):284–297, 2013.

[Kalinowski *et al.*, 2013a] T. Kalinowski, N. Narodytska, and T. Walsh. A social welfare optimal sequential allocation procedure. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, IJCAI '13, pages 227–233, 2013.

[Kalinowski *et al.*, 2013b] T. Kalinowski, N. Narodytska, T. Walsh, and L. Xia. Strategic behavior when allocating indivisible goods sequentially. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, July 14-18, 2013, Bellevue, Washington, USA.*, 2013.

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

[Kurokawa *et al.*, 2016] D. Kurokawa, A. D. Procaccia, and J. Wang. When can the maximin share guarantee be guaranteed? In *30th AAAI Conference on Artificial Intelligence, AAAI 2016,*, 2016.

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

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

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

[Procaccia and Wang, 2014] A. D. Procaccia and J. Wang. Fair enough: guaranteeing approximate maximin shares. In *ACM Conference on Economics and Computation, EC '14*, pages 675–692, 2014.

[Procaccia, 2015] 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, 2015.

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

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