

Fair Allocation of Indivisible Goods to Asymmetric Agents

Alireza Farhadi

*University of Maryland, College Park
Maryland, USA*

FARHADI@CS.UMD.EDU

Mohammad Ghodsi

*Sharif University of Technology
IPM - Institute for Research in Fundamental Sciences
Tehran, Iran*

GHODSI@SHARIF.EDU

MohammadTaghi Hajiaghayi

*University of Maryland, College Park
Maryland, USA*

HAIAGHA@CS.UMD.EDU

Sébastien Lahaie

*Google Research
New York City, USA*

SLAHAIE@GOOGLE.COM

David Pennock

*Microsoft Research
New York City, USA*

DPENNOCK@MICROSOFT.COM

Masoud Seddighin

*Sharif University of Technology
Tehran, Iran*

MSEDDIGHIN@CE.SHARIF.EDU

Saeed Seddighin

Hadi Yami
*University of Maryland, College Park
Maryland, USA*

SAEEDREZA.SEDDIGHIN@GMAIL.COM

HADIYAMI@CS.UMD.EDU

Abstract

We study fair allocation of indivisible goods to agents with *unequal entitlements*. Fair allocation has been the subject of many studies in both divisible and indivisible settings. Our emphasis is on the case where the goods are indivisible and agents have unequal entitlements. This problem is a generalization of the work by Procaccia and Wang (Procaccia & Wang, 2014) wherein the agents are assumed to be symmetric with respect to their entitlements. Although Procaccia and Wang show an almost fair (constant approximation) allocation exists in their setting, our main result is in sharp contrast to their observation. We show that, in some cases with n agents, no allocation can guarantee better than $1/n$ approximation of

a fair allocation when the entitlements are not necessarily equal. Furthermore, we devise a simple algorithm that ensures a $1/n$ approximation guarantee.

Our second result is for a restricted version of the problem where the valuation of every agent for each good is bounded by the total value he wishes to receive in a fair allocation. Although this assumption might seem without loss of generality, we show it enables us to find a $1/2$ approximation fair allocation via a greedy algorithm. Finally, we run some experiments on real-world data and show that, in practice, a fair allocation is likely to exist. We also support our experiments by showing positive results for two stochastic variants of the problem, namely *stochastic agents* and *stochastic items*.

1. Introduction

In this work, we conduct a study of *fairly* allocating *indivisible goods* among n agents with unequal claims on the goods. Fair allocation is a very fundamental problem that has received attention in both Computer Science and Economics. In the last decade this subject has attracted the attention of the researchers in Artificial Intelligence, Multiagent Systems, and Computational Social Choice (Caragiannis, Lai, & Procaccia, 2011; Cohler, Lai, Parkes, & Procaccia, 2011; Dickerson, Goldman, Karp, Procaccia, & Sandholm, 2014; Kurokawa, Procaccia, & Wang, 2016; Alijani, Farhadi, Ghodsi, Seddighin, & Tajik, 2017; Chen, Lai, Parkes, & Procaccia, 2013; Kash, Procaccia, & Shah, 2014).

This problem dates back to 1948 when Steinhaus (1948) introduced the *cake cutting* problem as follows: given n agents with different valuation functions for a cake, is it possible to divide the cake between them in such a way that every agent receives a piece whose value to him is at least $1/n$ of the whole cake? Steinhaus answered this question in the affirmative by proposing a simple and elegant algorithm which is called *moving knife*. Although this problem admits a straightforward solution, several versions of the cake cutting problem have been studied since then, many of which have not been settled after decades (Brams & Taylor, 1996; Robertson & Webb, 1998; Chen et al., 2013; Procaccia, 2013; Edmonds & Pruhs, 2006; Cohler et al., 2011; Woeginger & Sgall, 2007; Caragiannis et al., 2011; Alijani et al., 2017). For instance, a natural generalization of the problem in which we discriminate the agents based on their entitlements is still open. In this problem, every agent claims an entitlement e_i of the cake such that $\sum e_i = 1$, and the goal is to cut the cake into pieces and allocate them to the agents such that every agent a_i 's valuation for his piece is at least e_i fraction of his valuation for the entire cake. For two agents, Brams, Jones, and Klamler (2008) showed that at least two cuts are necessary to divide the cake between the agents. Furthermore, Robertson and Web (1997) proposed a modified version of cut and choose method to divide the cake between two agents with portions e_1, e_2 , where e_1 and e_2 are real numbers. McAvaney, Robertson, and Web (1992) considered the case when the entitlements are rational numbers. They used Ramsey partitions to show that when the entitlements are rational, one can make a proper division via $O(n^3)$ cuts.

Recently, a new line of research is focused on the fair allocation of indivisible goods. In contrast to the conventional cake cutting problem, in this problem instead of a heterogeneous cake, we have a set \mathcal{M} of indivisible goods and we wish to distribute them among n agents. Indeed, due to trivial counterexamples in this setting¹, the previous guarantee, that is every agent should obtain $1/n$ of his valuation for all items from his allocated set, is impossible

1. For instance if there is only one item, at most one agent has a non-zero profit in any allocation.

to deliver. To alleviate this problem, Budish (2011) proposed a concept of fairness for the allocation of indivisible goods namely *the maximin share*. Suppose we ask an agent a_i to divide the items between the agents in a way that *he thinks* is fair to everybody. Of course, agent a_i does not take into account other agents' valuations and only incorporates his valuation function in the allocation. Based on this, we define MMS_i equal to the minimum profit that any agent receives in this allocation, according to agent a_i 's valuation function. Obviously, in order to maximize MMS_i , agent a_i chooses an allocation that maximizes the minimum profit of the agents. We call an allocation fair (approximately fair), if every agent a_i receives a set of items that is worth at least MMS_i (a fraction of MMS_i) to him.

It is easy to see that MMS_i is the best possible guarantee that one can hope to obtain in this setting. If all agents have the same valuation function, then at least one of the agents receives a collection of items that are worth no more than MMS_i to him. A natural question that emerges here is whether a fair allocation with respect to MMS_i 's is always possible. Although the experiments are in favor of this conjecture, Procaccia and Wang (2014) refuted this by an elegant and delicate counterexample. They show such a fair allocation is impossible in some cases, even when the number of agents is limited to 3. On the positive side however, they show an approximately fair allocation can be guaranteed. More precisely, they show that there always exists an allocation in which every agent's profit is at least $(2/3)MMS_i$. Such an allocation is called a $2/3$ -MMS allocation. Amanitidis, Markakis, Nikzad, and Saberi (2015) later provided a proof for the existence of an MMS allocation for the case, when there are large enough items and the value of each agent for every item is drawn independently from a uniform distribution. A generalized form of this result was later proposed by Kurokawa *et al.* (2016) for arbitrary distributions. In a recent work, Caragiannis *et al.* (2016) proved that the maximum Nash welfare (MNW) solution, which selects an allocation that maximizes the product of utilities, for each agent guarantees a $2/(1 + \sqrt{4n - 3})$ fraction of her MMS.

Although it is natural to assume the agents have equal entitlements on the items, in most real-world applications, agents have unequal entitlements on the goods. For instance, in various religions, cultures, and regulations, the distribution of the inherited wealth is often unequal. Furthermore, the division of mineral resources of a land or international waters between the neighboring countries is often made unequally based on the geographic, economic, and political status of the countries. Another example is the allocation of resources (e.g., CPU, memory, and network bandwidth) in data centers, clusters, and cloud computing to the agents with heterogeneous demand profiles (Wang, Liang, & Li, 2015; Kash *et al.*, 2014).

For fairly allocating indivisible items to agents with different entitlements, two procedures are proposed by Brams and Taylor (1996). The first one is based on *Knaster's procedure of sealed bids*. In this method, we have an auction for selling each item. Therefore, for using it all the agents should have an adequate reserve of money which is the main issue of the procedure. The second procedure is based on the *method of markers* developed by William F. Lucas which is similar in spirit to the moving knife procedure. In this method, first we line up the items, and then the agents place some markers for dividing the items. This method suffers from high dependency of its final allocation to the order of the items in the line.

Agent duplication is another idea to deal with unequal entitlements. More precisely, when all of the entitlements are fractional numbers, we can duplicate each agent a_i to some

agents with similar valuation functions to a_i . The goal of this duplication is to reduce the problem to the case of equal entitlements. After the allocation, every agent a_i owns all of the allocated items to her duplicated agents. For instance, assume that we have three agents with entitlements $1/2$, $2/5$, and $1/10$, respectively. In this case, we duplicate the first agent to five agents and the second agent to four agents each having an entitlement of $1/10$. This way, we can reduce our problem to the case of equal entitlements. Although agent duplication may be practical when the items are divisible, this method does not apply to the indivisible setting. For instance, if the number of the agents is higher than the number of available items, we cannot allocate anything to some agents. Another issue with this method is that it works only for fractional entitlements.

In this paper, we study fair allocation of indivisible items with different entitlements using a model which resolves the mentioned issues. Our fairness criterion mimics the general idea of Budish for defining maxmin shares. Similar to Budish’s proposal, in order to define a maxmin share for an agent a_i , we ask the following question: how much benefit does agent a_i expect to receive from a fair allocation, if we were to divide the goods *only based on his valuation function*? If agent a_i expects to receive a profit of p from the allocation, then he should also recognize a minimum profit of $p \cdot e_j/e_i$ (e_i is the entitlement of agent a_i) for any other agent a_j , so that his own profit per entitlement is a lower bound for all agents. Therefore, one answer to this question would be the maximum value of p for which there exists an allocation such that agent a_i ’s profit-per-entitlement can be guaranteed to all other agents (according to his own valuation function). We define the maxmin shares of the agents based on this intuition.

Recall that we denote the number of agents with n and the entitlement of every agent a_i with e_i . We assume the entitlements always add up to 1. For every agent a_i , we define the weighted maxmin share denote by WMMS_i , to be the highest value of p for which there exists an allocation of the goods to the agents in which every agent a_j receives a profit of at least $p \cdot e_j/e_i$ based on agent a_i ’s valuation function. Similarly, we call an allocation α -WMMS, if every agent a_i obtains an α fraction of WMMS_i from his allocated goods. Notice that in case $e_i = 1/n$ for all agents, this definition is identical to Budish’s definition. Since our model is a generalization of Budish’s model, it is known that a fair allocation is not guaranteed to exist for every scenario. However, whether a $2/3$ approximation or in general a constant approximation WMMS allocation exists remains an open question.

Our first result is in contrast to that of Procaccia and Wang. We settle the above question by giving a $1/n$ hardness of approximation result for this problem. In other words, we show no algorithm can guarantee any allocation which is better than $1/n$ -WMMS in general. We further complement this result by providing a simple algorithm that guarantees a $1/n$ -WMMS allocation to all agents. As we show in Section 2, this hardness is a direct consequence of unreasonably high valuation of agents with low entitlements for some items. Moreover, in Section 3 we discuss that not only are such valuation functions unrealistic, but also an agent with such a valuation function has an incentive to misrepresent his valuations (Observation 1). Therefore, a natural limitation that one can add to the setting is to assume no item is worth more than WMMS_i for any agent a_i . We also study the problem in this mildly restricted setting and show in this case a $1/2$ -WMMS guarantee can be delivered via a greedy algorithm.

In contrast to our theoretical results, we show in practice a fair allocation is likely to exist by providing experimental results on real-world data. The source of our experiments is a publicly available collection of bids for eBay goods and services (Medina, 2014). More details about the experiments can be found in Section 4. We also support our claim by presenting theoretical analysis for the stochastic variants of the problem in which the valuation of every agent for a good is drawn from a given distribution.

1.1 Our Model

Let \mathcal{N} be a set of n agents, and \mathcal{M} be a set of m items. Each agent a_i has an *additive* valuation function V_i for the items. In addition, every agent a_i has an entitlement to the items, namely e_i . The entitlements add up to 1, i.e., $\sum e_i = 1$.

Since our model is a generalization of maxmin share, we begin with a formal definition of the maxmin shares for equal entitlements, proposed by Budish (2011). In this case, we assume all of the entitlements are equal to $1/n$. Let $\Pi(\mathcal{M})$ be the set of n -partitionings of the items. Define the maxmin share of agent a_i (MMS_i) of player i as

$$\text{MMS}_i = \max_{\langle A_1, A_2, \dots, A_n \rangle \in \Pi(\mathcal{M})} \min_{j \in [n]} V_i(A_j). \quad (1)$$

One can interpret the maxmin share of an agent as his worst case outcome as a divider in a divide-and-choose procedure against adversaries. Consider a situation that a cautious agent knows his own valuation on the items, but the valuations of other agents are unknown to him. If we ask the agent to run a divide-and-choose procedure, he tries to split the items in a way that the least valuable bundle is as attractive as possible.

When the agents have different entitlements, the above interpretation is no longer valid. The problem is that the agents have different entitlements and this discrepancy must somehow be considered in the divide-and-choose procedure. Thus, we need an interpretation of the maxmin share that takes the entitlements into account.

Let us get back to the case with the equal entitlements. Another way to interpret maxmin share is this: suppose that we ask agent a_i to fairly distribute the items in \mathcal{M} between n agents of \mathcal{N} , based on his own valuation function. In an ideal situation (e.g., if the goods are completely divisible), we expect a_i to allocate a share with value $V_i(\mathcal{M})/n$ to every agent. However, since the goods are indivisible, some sort of unfairness is inevitable. For this case, we wish that a_i does his best to retain fairness. MMS_i is in fact, a parameter that reveals how much fairness a_i can guarantee, regarding his valuation function.

Formally, to measure the fairness of an allocation by a_i , define a value F_A^i for any allocation $A = \langle A_1, A_2, \dots, A_n \rangle$ as

$$F_A^i = \frac{\min_j V_i(A_j)}{V_i(\mathcal{M})/n}.$$

In fact, we wish to make sure a_i reports an allocation A^* such that $F_{A^*}^i$ is as close to 1 as possible. The maxmin share of a_i is therefore defined as

$$\text{MMS}_i = F_{A^*}^i(V_i(\mathcal{M})/n). \quad (2)$$

It is easy to observe that Equations (1) and (2) are equivalent, since the fairest allocation in the absence of different entitlements is an allocation that maximizes value of the minimum bundle:

$$\begin{aligned} \text{MMS}_i &= F_{A^*}^i(V_i(\mathcal{M})/n) \\ &= \frac{\min_j V_i(A_j^*)}{V_i(\mathcal{M})/n}(V_i(\mathcal{M})/n) = \min_j V_i(A_j^*). \end{aligned}$$

Now, consider the case with different entitlements. Let e_i be the entitlement of agent a_i . Similar to the second interpretation for MMS_i , ask agent a_i to fairly distribute the items between the agents, but this time, considers the entitlements. In an ideal situation (e.g., a completely divisible resource), we expect the allocation to be proportional to the entitlements, i.e. a_i allocates a share to agent a_j with value exactly $V_i(\mathcal{M})e_j$ (note that when the entitlements are equal, this value equals to $V_i(\mathcal{M})/n$ for every agent). But again, such an ideal situation is very rare to happen and thus we allow some unfairness. In the same way, define the fairness of an allocation $A = \langle A_1, A_2, \dots, A_n \rangle$ as

$$F_A^i = \min_j \frac{V_i(A_j)}{V_i(\mathcal{M})e_j}. \quad (3)$$

Let $A^* = \langle A_1^*, A_2^*, \dots, A_n^* \rangle$ be an allocation by a_i that maximizes $F_{A^*}^i$. The weighted maxmin share of agent a_i is defined in the same way as MMS_i , that is:

$$\text{WMMS}_i = F_{A^*}^i V_i(\mathcal{M})e_i = e_i \min_j \frac{V_i(A_j^*)}{e_j}.$$

In summary, the value WMMS_i for every agent a_i is defined as follows:

$$\text{WMMS}_i = \max_{\langle A_1, A_2, \dots, A_n \rangle \in \Pi(\mathcal{M})} \min_{j \in [n]} V_i(A_j) \frac{e_i}{e_j}.$$

For more intuition, consider the following example:

Example 1. Assume that we have two agents a_1, a_2 with $e_1 = 1/3$ and $e_2 = 2/3$. Furthermore, suppose that there are 5 items b_1, b_2, b_3, b_4, b_5 with the following valuations for a_1 : $V_1(\{b_1\}) = V_1(\{b_2\}) = V_1(\{b_3\}) = 4, V_1(\{b_4\}) = 3$ and $V_1(\{b_5\}) = 9$. For the allocation $A = \langle \{b_5\}, \{b_1, b_2, b_3, b_4\} \rangle$, we have $F_A^1 = \min(\frac{9}{24 \cdot (1/3)}, \frac{15}{24 \cdot (2/3)})$ which means $F_A^1 = 15/16$. Moreover, for allocation $A' = \langle \{b_1, b_2\}, \{b_3, b_4, b_5\} \rangle$, we have $F_{A'}^1 = \min(\frac{8}{24 \cdot (1/3)}, \frac{16}{24 \cdot (2/3)})$ which means $F_{A'}^1 = 1$. Thus, A' is a fairer allocation than A . In addition, A' is the fairest possible allocation and hence, $\text{WMMS}_1 = 1 \cdot 24 \cdot 1/3 = 8$.

Example 1 also gives an insight about why agent duplication (as introduced in the Introduction) is not a good idea. For this example, if we duplicate agent a_2 , we have three agents with the same entitlements. But any partitioning of the items into three bundles, results in a bundle with value at most 7 to a_1 .

Finally, an allocation of the items in \mathcal{M} to the agents in \mathcal{N} is said to be α -WMMS, if the total value of the share allocated to each agent a_i is worth at least αWMMS_i to him.

2. A Tight $1/n$ Bound on the Optimal Allocation

In spite of the fact that there exists a $2/3$ -WMMS guarantee when all the entitlements are equal, in our general setting surprisingly we provide a counterexample which proves that there is no guarantee better than $1/n$ -WMMS. We complement this result by showing that a $1/n$ -WMMS always exists. Thus, these two theorems make a tight bound for the problem.

We provide a counterexample in Theorem 2.1 which shows no approximation better than $1/n$ -WMMS is possible. The main property of our counterexample is a large gap between the value of items for different agents.

Theorem 2.1. *There exists no guarantee better than $1/n$ -WMMS when the entitlements to the items may differ.*

Proof. We propose an example that admits no allocation better than $1/n$ -WMMS. To this end, consider an instance with n agents and $2n - 1$ items and let $e_i = \epsilon$ for all $i < n$ and $e_n = 1 - (n - 1)\epsilon$. The valuation functions of the first $n - 1$ agents are the same. For every agent a_i with $1 \leq i < n$, $V_i(\{b_j\})$ is as follows:

$$V_i(\{b_j\}) = \begin{cases} \epsilon & \text{if } j \leq n - 1 \\ 1 - (n - 1)\epsilon & \text{if } j = n \\ 0 & \text{if } j > n. \end{cases}$$

Also, for agent a_n we have:

$$V_n(\{b_j\}) = \begin{cases} \frac{1 - (n - 1)\epsilon}{n} & \text{if } j \leq n \\ \epsilon & \text{if } j > n. \end{cases}$$

First, note that $\text{WMMS}_i = \epsilon$ for the first $n - 1$ agents and $\text{WMMS}_n = 1 - (n - 1)\epsilon$. For the first $n - 1$ agents, the optimal partitioning is to allocate b_i to a_i for all $i \leq n$. Furthermore, the optimal partitioning for a_n is to allocate items $b_{n+1}, b_{n+2}, \dots, b_{2n-1}$ to the first $n - 1$ agents and keep the first n items for himself. In this case, $V_n(\{b_1, b_2, \dots, b_n\}) = 1 - (n - 1)\epsilon$, and $V_n(\{b_{n+i}\}) = \epsilon$ for $1 \leq i < n$. Therefore, $\text{WMMS}_n = 1 - (n - 1)\epsilon$.

On the other hand, in any allocation that guarantees a non-zero fraction of WMMS for every agent, at most one of the items b_1, b_2, \dots, b_n is allocated to a_n , since the rest of the items have value 0 for the first $n - 1$ agents. Therefore, the items allocated to a_n are worth at most

$$\frac{1 - (n - 1)\epsilon}{n} + (n - 1)\epsilon = 1/n + \epsilon(n - 1 - \frac{n - 1}{n}) \leq 1/n + n\epsilon.$$

to him. Thus, the best fraction of WMMS that can be guaranteed is

$$\frac{1/n + n\epsilon}{1 - (n - 1)\epsilon}. \quad (4)$$

Equation (4) can be made arbitrarily close to $1/n$, by choosing sufficiently small ϵ . Thus, no allocation can guarantee an approximation better than $1/n$ for this example. \square

Algorithm 1: $1/n$ -WMMS allocation

Input: \mathcal{N} , \mathcal{M} , valuation functions V_1, \dots, V_n , and entitlements e_1, \dots, e_n (without loss of generality, sorted in descending order).

Output: Allocation $A = A_1, \dots, A_n$.

1
 1: **for** i from one to $|\mathcal{M}|$ **do**
 2: $r \leftarrow (i \bmod |\mathcal{N}|) + 1$
 3: $b_j = \arg \max_{b_k \in \mathcal{M}} V_r(b_k)$
 4: assign item b_j to a_r and remove b_j from \mathcal{M}

Theorem 2.1 gives a $1/n$ -WMMS upper-bound. In Theorem 2.2, we show that the provided upper-bound is tight. Algorithm 1 uses a simple greedy procedure, which is similar in spirit to an algorithm of Amanatidis *et al.* (2015), guaranteeing $1/n$ -WMMS allocation as follow: sort the agents in descending order of entitlements. Starting from the first agent, ask every person to collect the most valuable item from the remaining set, one by one. Repeat the process until no more item is left.

Theorem 2.2. *Algorithm 1 guarantees a $1/n$ -WMMS allocation.*

Proof. If the number of items is smaller than the number of agents then the proof is trivial. Therefore, from this point on, we assume $m \geq n$. Without loss of generality we assume agents are sorted in descending order of their entitlements, that is $e_1 \geq e_2 \dots \geq e_n$. The goal is to prove that for each agent a_i he receives at least $1/n$ -WMMS $_i$ using Algorithm 1. Suppose that $B^* = \langle B_1^*, \dots, B_n^* \rangle$ is the optimal allocation for a_i (i.e., the allocation that maximizes $F_{A^*}^i$). Without loss of generality suppose items are sorted according to their value for a_i in descending order that is:

$$V_i(\{b_1\}) \geq V_i(\{b_2\}) \geq \dots \geq V_i(\{b_{|\mathcal{M}|}\}).$$

We call items b_1, b_2, \dots , and b_{i-1} *heavy items* for a_i . Let H be the set of heavy items for a_i . Since the entitlements of agents are sorted from a_1 to a_n , for agents a_j and $a_{j'}$ when $e_j > e_{j'}$ we have:

$$V_i(B_j^*) \geq V_i(B_{j'}^*). \quad (5)$$

Now, the goal is to prove $V_i(B_i^*) \leq V_i(\mathcal{M} \setminus H)$ and use it to the guarantee of the algorithm. Since $B_i^* \subseteq \mathcal{M}$, if $B_i^* \cap H = \emptyset$, clearly $V_i(B_i^*) \leq V_i(\mathcal{M} \setminus H)$ holds. In case $B_i^* \cap H \neq \emptyset$, consider agent $a_{k'}$ where $1 \leq k' < i$ and $B_{k'}^* \cap H = \emptyset$. According to Inequality (5), $a_{k'}$ has a greater entitlement than a_i . Therefore, we have: $V_i(B_{k'}^*) \geq V_i(B_i^*)$ which yields that $B_{k'}^*$ is worth at least $V_i(\{b_k\})$ for a_i where all the items of $B_{k'}^*$ are in $\mathcal{M} - H$. Therefore, we can imply that $V_i(B_i^*) \leq V_i(\mathcal{M} \setminus H)$. The items are sorted in the descending order of their value to a_i from b_1 to $b_{|\mathcal{M}|}$ which implies

$$V_i(\{b_i\}) + V_i(\{b_{n+i}\}) + V_i(\{b_{2n+i}\}) + \dots \geq V_i(\mathcal{M} \setminus H)/n.$$

The l -th assigned item to a_i by the algorithm is not worth less than $V_i(\{b_{(l-1)n+i}\})$. Hence, the assigned items to a_i is worth at least $V_i(\{b_i\}) + V_i(\{b_{n+i}\}) + V_i(\{b_{2n+i}\}) + \dots$ for him which is not less than $V_i(\mathcal{M} \setminus H)/n \geq V_i(B_i^*)/n \geq 1/n$ -WMMS $_i$. □

Algorithm 2: 1/2-WMMS allocation for the restricted case

Input: \mathcal{N} , \mathcal{M} , valuation functions V_1, \dots, V_n , and entitlements e_1, \dots, e_n (without loss of generality, sorted in descending order).

Output: Allocation $A = A_1, \dots, A_n$.

-
- 1: **while** $\mathcal{M} \neq \emptyset$ and $\exists A_i$ where $V_i(A_i) < 1/2\text{-WMMS}_i$ **do**
 - 2: define candidate set $C_i = \emptyset$ for each agent a_i .
 - 3: **for** $b_j \in \mathcal{M}$ **do**
 - 4: add b_j to C_i where $V_i(\{b_j\})/V_i(\mathcal{M})F_{A^*}^i$ is maximum among all of the unsatisfied agents.
 - 5: Choose unsatisfied agent a_i and item b_j in its candidate set which maximize $V_i(\{b_j\})/V_i(\mathcal{M})F_{A^*}^i$ among all of the unsatisfied agents and items in their candidate sets.
 - 6: Assign item b_j to agent a_i .
-

3. Allocation for the Restricted Case

In Section 2, we gave a tight $1/n$ -WMMS guarantee for the fair allocation problem with unequal entitlements. In this section, we consider a reasonable restriction of the problem which gives a 1/2-WMMS guarantee. In this restricted setting the value of each item b_j to each agent a_i is no more than WMMS_i . Observation 1 shows that how this assumption can be invaluable.

Observation 1. *If $V_i(b_j) \geq \text{WMMS}_i$, it is in the best interest of a_i to report his valuation for b_j equal to infinity. Because his WMMS may increase in this way, and he achieves more items after the allocation of the algorithm, because his WMMS may increase in this way, and he will be satisfied even if he receives only this item.*

In this section, we provide an algorithm with 1/2-WMMS guarantee for the restricted case of the problem. For a case where the entitlements are equal, an algorithm, namely *bag filling* guarantees 1/2-WMMS for all the agents. In the bag filling algorithm, we start with an empty bag. In each step, we add a remaining item to the bag. After each addition, if the total value of items in the bag becomes more than $1/2\text{-WMMS}_i$ for an unsatisfied agent a_i , we allocate all the items in the bag to him and repeat the procedure with an empty bag. After running this simple procedure, each agent a_i receives at least $1/2\text{-WMMS}_i$.

Our algorithm allocates the items to agents in a more clever way. We allocate an item to an agent in each step of the algorithm until each agent a_i receives at least $1/2\text{-WMMS}_i$ by the allocation. To this end, in each step, first the algorithm for each agent a_i creates a candidate set of items where b_j is in the candidate set of a_i if $V_i(\{b_j\})/V_i(\mathcal{M})F_{A^*}^i$ be the maximum number among all of the unsatisfied agents. Then, the algorithm chooses unsatisfied agent a_i and item b_j in its candidate set which maximize $V_i(\{b_j\})/V_i(\mathcal{M})F_{A^*}^i$ among all of the unsatisfied agents and items in their candidate sets, and assigns this item to the agent.

Before proving Algorithm 2 guarantees a 1/2-WMMS allocation, we prove an auxiliary lemma which argues that using Algorithm 2 no agent receives more than his WMMS.

Lemma 3.1. *Algorithm 2 does not allocate more than $WMMS_i$ for any agent a_i .*

Proof. The algorithm does not assign any more items to a satisfied agent. Hence, any satisfied agent a_i has less than $1/2$ - $WMMS_i$ value of items in A_i before he receives the last item. For the sake of contradiction, suppose that the algorithm allocates more than $WMMS_i$ to a_i . Without loss of generality, suppose that the last item allocated to a_i is b_j . Since before allocating b_j to a_i he had less than $1/2$ - $WMMS_i$ value of items in A_i , the value of b_j to a_i is more than $1/2$ - $WMMS_i$. Since $V_i(b_j) \leq WMMS_i$, before allocating b_j to a_i we have $A_i \neq \emptyset$. Since $V_i(b_j)$ is more than the value of all the other allocated items to a_i , b_j was not in the candidate set of a_i when we were assigning the other items to a_i . Hence, b_j was in the candidate set of an other agent $a_{i'}$. Therefore, $V_{i'}(\{b_j\})/V_{i'}(\mathcal{M})F_{A^*}^{i'} \geq V_i(\{b_j\})/V_i(\mathcal{M})F_{A^*}^i$. Since this value is greater than the values of all other items in A_i , the algorithm first allocate b_j to unsatisfied agent $a_{i'}$. \square

Now, using Lemma 3.1, we prove the approximation guarantee of the algorithm.

Theorem 3.2. *Algorithm 2 ensures a $1/2$ - $WMMS$ guarantee when for each agent a_i and item b_j , $V_i(b_j) \leq WMMS_i$.*

Proof. It is clear that if the algorithm satisfies all the agents, it ensures the approximation guarantee. Now, for the sake of contradiction assume that there exists an unsatisfied agent a_i at the end of the algorithm. For each item b_j we define

$$v'_{b_j} = V_{i'}(b_j)/V_{i'}(\mathcal{M})F_{A^*}^{i'},$$

where $a_{i'}$ is the recipient of b_j in the allocation. Since v'_{b_j} is maximal according to the algorithm, and a_i is not a satisfied agent, we have:

$$V_i(b_j)/V_i(\mathcal{M})F_{A^*}^i \leq V_{i'}(b_j)/V_{i'}(\mathcal{M})F_{A^*}^{i'}. \quad (6)$$

Lemma 3.1 implies that $\sum_{b_j} v'_{b_j} \leq 1$, and since we have at least one unsatisfied agent we can write:

$$\sum_{b_j} v'_{b_j} < 1. \quad (7)$$

Inequality (6) along with Inequality (7) implies

$$\sum_{b_j} V_i(b_j)/V_i(\mathcal{M})F_{A^*}^i < 1. \quad (8)$$

Finally Inequality (8) yields that $\sum_{b_j} V_i(b_j) < V_i(\mathcal{M})F_{A^*}^i$ which is a contradiction. \square

4. Empirical Results

As we discussed in Section 2, in extreme cases, making a $WMMS$ allocation or even an approximately $WMMS$ allocation is theoretically impossible. However, our counter-example is extremely delicate and thus very unlikely to happen in real-world. Here, we show in practice fair allocations w.h.p exist, especially when the number of items is large.

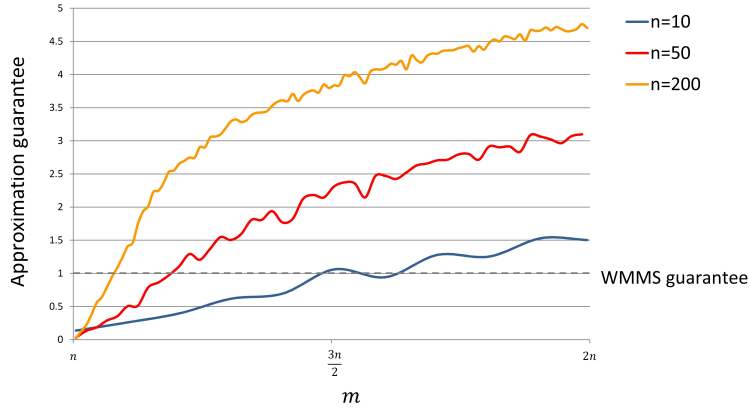


Figure 1: The vertical line denotes the ratio of the valuation of the allocated set to the maxmin guarantee of the agents. The horizontal line shows the number of items varying from n to $2n$. Blue, red, and yellow poly lines illustrate the performance of our algorithm for $n = 10$, $n = 50$, and $n = 200$ respectively.

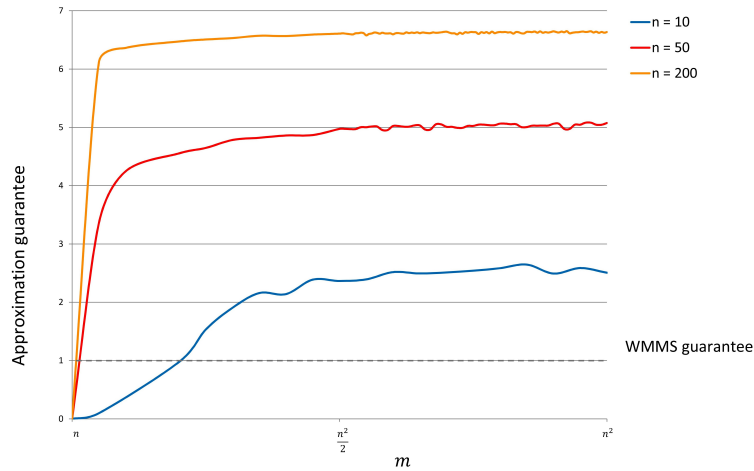


Figure 2: The vertical line denotes the ratio of the valuation of the allocated set to the maxmin guarantee of the agents. The horizontal line shows the number of items varying from n to n^2 . Blue, red, and yellow poly lines illustrate the performance of our algorithm for $n = 10$, $n = 50$, and $n = 200$ respectively.

We draw the valuation of the agents for the goods based on a collection of publicly available bids for eBay items (Medina, 2014). More precisely, for m items, we randomly choose m different categories of goods from the dataset. Moreover, for every agent a_i and item b_j , we set $V_i(\{b_j\})$ to a submitted bid for the corresponding category of item b_j chosen uniformly at random.

Despite many real-world examples of allocations with different entitlements,² we could not find any appropriate publicly available data for our setting.³ That’s why we ran the experiments on randomly generated vectors for entitlements. In this sense, our data is not exactly real-world, but a mix of real-world data and random entitlements. For an instance of the problem with n agents and m items, we run the experiments with 1000 different vector of entitlements drawn from the uniform distribution (and scaled up to satisfy $\sum e_i = 1$). For every n and m , we take **the minimum WMMS guarantee obtained by all 1000 runs**, and show it in Figures 1 and 2.

We used heuristic methods to compute the maxmin guarantees and estimate maximin-share values. These heuristic methods include greedy algorithms such as local search and Algorithm 2 as well as random allocations. In our local search approach, we start by an arbitrary allocation and iteratively improve the guarantee by local exchanges of items. Another greedy algorithm we use is as follows: let $\gamma_i = V_i(S_i)/\text{WMMS}_i$, where S_i is the set of currently allocated items to agent a_i . At each step, we select an agent a_i which minimizes γ_i , and allocate him the most valuable remaining good.

Since the obtained allocations are not necessarily optimal, our results are only *lower bounds* to the actual WMMS guarantees. Nonetheless, the optimal guarantees are very close to the estimated ones.

Figures 1 and 2 illustrate the result of the runs for $n = 10$, $n = 50$, and $n = 200$ respectively. Figure 1 only depicts the WMMS guarantees for $m \in [n, 2n]$ whereas in Figure 2 the number of items varies from n to n^2 .

As shown in Figures 1 and 2, the approximation guarantee improves as we increase the number of items. Moreover, unless m is very close to n , a WMMS allocation exists in our experiments (notice that the guarantee is above 1 when m is considerably larger than n).

5. Stochastic Setting

In Section 2 we presented a counterexample to show that no allocation better than $1/n$ -WMMS can be guaranteed. However, the construction described in the counterexample is very unlikely to happen in the real settings. Here, we show that WMMS allocation exists with high probability when a small randomness is allowed in the setting.

Considering stochastic settings is common in the fair allocation problems since many real-world instances can be modeled with random distributions (Budish, 2011; Kurokawa et al., 2016; Amanatidis et al., 2015; Dickerson et al., 2014). The general probabilistic model used in previous works is as follows: every agent a_i has a probability distribution \mathcal{D}_i over $[0, 1]$ and for every item b_j , the value for $V_i(\{b_j\})$ is randomly sampled from \mathcal{D}_i . Amanatidis *et al.* (2015) proved the existence of an MMS allocation for the special case of $\mathcal{D}_i = U(0, 1)$, where $U(0, 1)$ is the standard uniform distribution with minimum 0 and maximum 1. Kurokawa *et al.* (2016) considered the problem for arbitrary random distribution \mathcal{D}_i with the condition that $\mathbb{V}[\mathcal{D}_i] \geq c$ for a positive constant c .⁴ A considerable part of their proof for the existence of an MMS allocation is referred to the work of Dickerson *et*

2. We mentioned some of these scenarios in the Introduction

3. There are some online services such as “Assign Credit” (Procaccia & Goldman, 2014), where their data can help generate entitlement vectors. However, their data are not publicly available.

4. $\mathbb{V}[\mathcal{D}_i]$ is the variance of \mathcal{D}_i .

al. (2014), where the authors prove the existence of an envy-free allocation in the stochastic settings with arbitrary random distributions.

In this section, we consider two different probabilistic models. Our first model is the same as the model of Krukawa *et al.* (2016), with the exception that we omit the restriction $\mathbb{V}[\mathcal{D}_i] \geq c$. We name this model as *Stochastic Agents* model. In the second model, every item b_i has a probability distribution \mathcal{D}_i and for every agent a_j , value of $V_j(\{b_i\})$ is randomly drawn from \mathcal{D}_i . We choose the name *Stochastic Items* for the second model. We believe that *Stochastic Items* model is more realistic since the first model does not make any distinguish between the items. None of the previous works mentioned above considered this model.

We leverage *Hoeffding's* inequality to prove the existence of WMMS allocation. Theorem 5.1 states the general form of this inequality (Hoeffding, 1963).

Theorem 5.1 (General Form of Hoeffding (1963)). *Let X_1, X_2, \dots, X_n be random variables bounded by the interval $[0, 1] : 0 \leq X_i \leq 1$. We define the empirical mean of these variables by $\bar{X} = \frac{1}{n}(X_1 + X_2 + \dots + X_n)$. Then, the following inequality holds:*

$$\mathbb{P}(|\bar{X} - \mathbb{E}(\bar{X})| \geq t) \leq 2e^{-2nt^2}. \quad (9)$$

Regarding Theorem 5.1, let $X = n \cdot \bar{X} = \sum_i X_i$ and let $\mu = n \cdot \mathbb{E}(\bar{X})$. By Inequality (9), we have:

$$\mathbb{P}(|X - \mu| \geq nt) \leq 2e^{-2nt^2}. \quad (10)$$

By setting $nt = \delta\mu$, we rewrite Equation (10) as:

$$\mathbb{P}(|X - \mu| \geq \delta\mu) \leq 2e^{-\frac{2(\delta\mu)^2}{n}}. \quad (11)$$

5.1 Model I: Stochastic Agents

As mentioned before, in the first model we assume that every agent has a probability distribution \mathcal{D}_i and for every item b_j , the value of $V_i(\{b_j\})$ is randomly sampled from \mathcal{D}_i . Furthermore, we suppose $\mu_i = \mathbb{E}(\mathcal{D}_i)$. Throughout this section, we assume that $m \geq n$. This is without loss of generality, because for the case $m < n$, WMMS $_i$ for every agent a_i equals to zero. For the Stochastic Agents model, we state Theorem 5.2.

Theorem 5.2. *Consider an instance of the fair allocation problem with unequal entitlements, such that the value of every item for every agent a_i is randomly drawn from distribution \mathcal{D}_i . Furthermore, let $s = \min_i s_i$ and $\mu = \min_i \mu_i$. Then, for every $0 < \epsilon < 1$, there exists a value $m' = m'(\frac{1}{\epsilon}, \frac{1}{\mu}, \frac{1}{\epsilon})$ (which means m' is a function of $\frac{1}{\epsilon}, \frac{1}{\mu}, \frac{1}{\epsilon}$) such that if $m \geq m'$, then almost surely a $(1 - \epsilon)$ -WMMS allocation exists.*

In the rest of this section, we prove Theorem 5.2. Consider the algorithm that allocates $t_i = \lfloor me_i \rfloor$ items to every agent a_i . We know that the value of item b_i for agent a_j is randomly sampled from \mathcal{D}_i . From the point of view of the algorithm, it trivially does not matter whether the value of items are sampled after the allocation or before the allocation. Thus, we can suppose that the value of every item is sampled after the allocation of items.

For now, we know that $t_i = \lfloor me_i \rfloor$ number of items are assigned to a_i . Argue that for every $\epsilon > 0$, there exists a value m' , such that for every $m \geq m'$, $t_i \geq me_i(1 - \epsilon)$. In Lemma 5.3 we bound the value of m' in terms of e_i and ϵ .

Lemma 5.3. *For every $m > \frac{1}{\epsilon e_i}$, $\lfloor me_i \rfloor \geq me_i(1 - \epsilon)$.*

Proof. We know $\lfloor me_i \rfloor \geq me_i - 1 = me_i(1 - \frac{1}{me_i})$. If $m > \frac{1}{\epsilon e_i}$, then $\lfloor me_i \rfloor \geq me_i(1 - \frac{1}{\frac{1}{\epsilon e_i} e_i}) = me_i(1 - \epsilon)$. This, completes the proof. \square

For the rest of the proof, suppose that $m > \frac{1}{\epsilon e_i}$. Let X_i be the variable indicating total value of items allocated to a_i . Note that $\mathbb{E}(X_i) = t_i \mu_i$. Regarding Equation (11), we have:

$$\mathbb{P}(|X_i - t_i \mu_i| \geq \delta t_i \mu_i) \leq 2e^{-\frac{2(\delta t_i \mu_i)^2}{t_i}}.$$

We want to choose δ such that $\mathbb{P}(|X_i - t_i \mu_i| \geq \delta t_i \mu_i) \leq \frac{1}{2mn}$. We have:

$$2e^{-\frac{2(\delta t_i \mu_i)^2}{t_i}} \leq \frac{1}{2mn} \Rightarrow -2(\delta t_i \mu_i)^2 \leq t_i \ln \frac{1}{4mn} \Rightarrow \delta \geq \sqrt{\frac{t_i \ln 4mn}{2(t_i \mu_i)^2}}.$$

Regarding the facts that $t_i \geq me_i(1 - \epsilon)$ and $m \geq n$, we have:

$$\sqrt{\frac{t_i \ln 4mn}{2(t_i \mu_i)^2}} \leq \sqrt{\frac{\ln 2 + \ln m}{e_i m \mu_i^2 (1 - \epsilon)}}.$$

Therefore, it's enough to choose δ such that

$$\delta \geq \sqrt{\frac{\ln 2 + \ln m}{e_i m \mu_i^2 (1 - \epsilon)}}. \quad (12)$$

Now, let t'_i be the number of items that are not assigned to a_i . Since $t_i + t'_i = m$, regarding the fact that $t_i \geq me_i(1 - \epsilon)$, we have $t'_i \leq m - me_i(1 - \epsilon)$, which means $t'_i \leq m + me_i(\epsilon - 1)$. On the other hand, $t'_i \geq m(1 - e_i)$. Also, let X'_i be the variable indicating total value of the items that are not allocated to a_i . By the same deduction as t_i for t'_i we have:

$$\mathbb{P}(|X'_i - t'_i \mu_i| \geq \delta' t'_i \mu_i) \leq 2e^{-\frac{2(\delta' t'_i \mu_i)^2}{t'_i}}.$$

Let δ' be the value that $\mathbb{P}(|X_i - t'_i \mu_i| \geq \delta' t'_i \mu_i) < \frac{1}{2mn}$. We have:

$$2e^{-\frac{2(\delta' t'_i \mu_i)^2}{t'_i}} \leq \frac{1}{2mn} \Rightarrow (\delta' t'_i \mu_i)^2 \geq \frac{t'_i \ln 4mn}{2}$$

$$\delta' \geq \sqrt{\frac{\ln 4mn}{2t'_i \mu_i^2}}. \quad (13)$$

Thus, it's enough to choose δ in a way that Inequality (13) holds. Regarding the facts that $m > n$ and $t'_i \geq m(1 - e_i)$,

$$\sqrt{\frac{\ln 4mn}{2t'_i \mu_i^2}} \leq \sqrt{\frac{\ln 2 + \ln m}{m(1 - e_i) \mu_i^2}}.$$

Therefore, it's enough to choose δ' in a way that

$$\delta' \geq \sqrt{\frac{\ln 2 + \ln m}{m(1 - e_i)\mu_i^2}}. \quad (14)$$

Now, suppose that both Inequalities (12) and (14) are held. Considering S_i as the set of items assigned to a_i , with the probability of at least $1 - (\frac{1}{2mn} + \frac{1}{2mn}) = 1 - \frac{1}{mn}$ we have:

$$\frac{V_i(S_i)}{V_i(\mathcal{M})} = \frac{(1 - \delta)t_i\mu_i}{(1 + \delta)t_i\mu_i + (1 + \delta')t'_i\mu_i}.$$

It is easy to show that, there always exist an m'_i such that for all $m \geq m'_i$ both Inequalities (12) and (14) hold for $\delta = \epsilon$ and $\delta' = \epsilon$. Regarding this, we have:

$$\begin{aligned} \frac{V_i(S_i)}{V_i(\mathcal{M})} &\geq \frac{(1 - \epsilon)t_i\mu_i}{(1 + \epsilon)t_i\mu_i + (1 + \epsilon)t'_i\mu_i} \\ &= \frac{(1 - \epsilon)t_i}{(1 + \epsilon)m} \\ &\geq \frac{(1 - \epsilon)m e_i (1 - \epsilon)}{(1 + \epsilon)m} \\ &= \frac{(1 - \epsilon)^2}{1 + \epsilon} e_i = \frac{(1 + \epsilon)^2 - 4\epsilon}{1 + \epsilon} e_i \\ &= (1 + \epsilon)e_i - \frac{4\epsilon}{(1 + \epsilon)} e_i \\ &\geq (1 + \epsilon)e_i - 4\epsilon e_i = (1 - 3\epsilon)e_i. \end{aligned}$$

Therefore, with the probability at least $1 - \frac{1}{mn}$ we have:

$$\forall a_i \in \mathcal{N} \quad \frac{V_i(S_i)}{V_i(\mathcal{M})} \geq (1 - 3\epsilon)e_i. \quad (15)$$

Now, suppose that the Inequality 15 holds for every agent a_i , with probability at least $1 - \frac{1}{mn}$. Considering all the agents, with the probability at least $(1 - \frac{1}{mn})^n \geq 1 - \frac{n}{nm} = 1 - \frac{1}{m}$, value of $\frac{V_i(S_i)}{V_i(\mathcal{M})}$ for every agent a_i is at least $e_i(1 - 3\epsilon)$. Regarding the fact that $\text{WMMS}_i \leq e_i V_i(\mathcal{M})$, we have

$$\forall a_i \in \mathcal{N} \quad V_i(S_i) \geq \text{WMMS}_i(1 - 3\epsilon).$$

This completes the proof.

5.2 Model II: Stochastic Items

As mentioned, in the Stochastic Items model, every item b_i has a probability distribution \mathcal{D}_i and the value of every agent a_j for item b_i is randomly chosen from \mathcal{D}_i . For this model, we prove Theorem 5.4. The theorem states that for large enough m , almost surely a $(1 - \epsilon)$ -WMMS allocation exists.

Theorem 5.4. *Suppose that for every agent a_i , $\mathbb{E}(\mathcal{D}_i) > c$ for a non-negative constant c . Then for all $0 < \epsilon < 1$, there exists $m' = m'(c, \epsilon, n, e_1, \dots, e_n)$ such that if $m \geq m'$, then, almost surely, $(1 - \epsilon)$ -WMMS allocation exists.*

In the rest of the section, we prove Theorem 5.4. First, in Lemma 5.6 we prove the existence of an allocation that assigns to every agent a_i , a set of items with value at least $\text{WMMS}_i - M_i$, where $M_i = \max_j (V_i(\{b_j\}))$. The idea to prove this fact is inspired by Bezakova and Dani (2005). Argue that we can formulate the allocation problem with unequal entitlements as the following integer program:

$$\begin{aligned}
 \sum_{b_j \in \mathcal{M}} V_i(\{b_j\}) \cdot f_{i,j} &\geq V(\mathcal{M}) \cdot e_i \quad \forall a_i \in \mathcal{N} \\
 \sum_{a_i \in \mathcal{N}} f_{i,j} &= 1 \quad \forall b_j \in \mathcal{M} \\
 f_{i,j} &\in \{0, 1\}
 \end{aligned} \tag{16}$$

In IP(16), variable $f_{i,j}$ determines whether b_j is assigned to agent a_i or not. Considering the fact that $V_i(\mathcal{M}) \cdot e_i$ is a trivial upper bound on WMMS_i , any solution to IP16 is a feasible solution to the assignment problem with unequal entitlements. By relaxing the second and the third condition, we can convert IP16 to LP17:

$$\begin{aligned}
 \sum_{b_j \in \mathcal{M}} V_i(\{b_j\}) \cdot f_{i,j} &\geq V(\mathcal{M}) \cdot e_i \quad \forall a_i \in \mathcal{N} \\
 \sum_{a_i \in \mathcal{N}} f_{i,j} &\leq 1 \quad \forall b_j \in \mathcal{M} \\
 f_{i,j} &\geq 0
 \end{aligned} \tag{17}$$

For every feasible solution \mathcal{A} to LP17, we construct the bipartite graph $G_{\mathcal{A}} \langle I, J, E \rangle$ where $I = \{1, 2, \dots, n\}$ and $J = \{1, 2, \dots, m\}$ correspond to the set of players and items, respectively. An edge i, j is included if $f_{i,j} > 0$. Then by the same way used by Bezakova and Dani, we will prove the following theorem.

Lemma 5.5. *There exists a solution \mathcal{A}' to LP17, such that $G_{\mathcal{A}'}$ is a pseudoforest. (each component of the graph is either a tree or a tree with an extra edge)*

Proof. We have $mn + m + n$ inequalities defining the polytope of feasible solutions of LP17. We have mn variables $f_{i,j}$, therefore every solution which is located in the corner of polytope satisfies at least mn inequalities as equalities and there will be at most $m + n$ non-zero variables in these solutions. By the same method used by Bezakova and Dani, it is clear to show that if \mathcal{A}' is corresponding solution to a corner of polytope, then $G_{\mathcal{A}'}$ is a pseudoforest. \square

We call the solution with the property defined in Lemma 5.5 as constrained solution. It is shown that every constrained solution for LP17, can be converted to a solution for IP16, such that every agent a_i loses at most one item b_j where $f_{i,j} > 0$ (Bezakova & Dani, 2005).

Lemma 5.6. *There exists an allocation in which every agent a_i gets at least $\text{WMMS}_i - \max_j V_i(\{b_j\})$.*

Proof. The polytope of of feasible solutions of LP17 is non-empty, because it has at least one solution which is, $f_{i,j} = \frac{1}{e_i}$ for every agent a_i and item b_j . Therefore there exists a

constrained solution for LP17, and using the same method of Bezakova and Dani, this solution can be converted to a solution for IP16, such that every agent loses at most one item. \square

Proofs of Lemmas 5.5 and 5.6 are omitted and included in the full version.

Now we show that there exists $m'_i = m'_i(c, \epsilon, n, e_1, \dots, e_n)$, such that if $m \geq m'_i$, then, $\text{WMMS}_i \geq \frac{1}{\epsilon}$ with the probability at least $1 - \frac{1}{mn}$. Therefore, whenever $m \geq \max(m'_1, m'_2, \dots, m'_n)$, with the probability at least $(1 - \frac{1}{mn})^n \geq 1 - \frac{n}{nm} = 1 - \frac{1}{m}$, for every agent a_i , $\text{WMMS}_i \geq \frac{1}{\epsilon}$. According to Lemma 5.6, there is an allocation in which every agent gets at least

$$\begin{aligned} \text{WMMS}_i - \max_j V_j(\{b_j\}) &\geq \text{WMMS}_i - 1 \\ &\geq \text{WMMS}_i \left(1 - \frac{1}{\text{WMMS}_i}\right) \\ &\geq \text{WMMS}_i(1 - \epsilon). \end{aligned}$$

Therefore, a $(1 - \epsilon)$ -WMMS allocation is guaranteed to exist with the probability at least $1 - \frac{1}{m}$.

Lemma 5.7. *Suppose that for every agent a_j , $\mathbb{E}(\mathcal{D}_j) > c$ for a non-negative constant c . Then for all $0 < \epsilon < 1$, there exists $m'_i = m'_i(c, \epsilon, n, e_1, \dots, e_n)$ such that if $m \geq m'_i$, $\text{WMMS}_i \geq \frac{1}{\epsilon}$ with the probability at least $1 - \frac{1}{mn}$.*

Proof. If there exists a partition of \mathcal{M} , $\pi = B_1, B_2, \dots, B_n$ in which $V_i(B_j) \geq \frac{1}{\epsilon e_i}$ for every agent a_j , then:

$$\text{WMMS}_i \geq e_i \cdot \min\left(\frac{1}{\epsilon e_i e_1}, \dots, \frac{1}{\epsilon e_i e_n}\right) \geq \min\left(\frac{1}{\epsilon e_1}, \dots, \frac{1}{\epsilon e_n}\right) \geq \frac{1}{\epsilon}. \quad (18)$$

Suppose that $m = \frac{\alpha}{\epsilon}$, then for every item b_k and every agent a_j , we assign this item to this agent with the probability e_j . Let $X_{j,k}$ be a random variable that takes the value $V_i(\{b_k\})$ with the probability e_j and 0 otherwise. We have $\mathbb{E}[X_{j,k}] > \frac{c}{e_j}$. Let $X_j = \sum_{k=1}^m X_{j,k}$. By setting $nt = \gamma$, we rewrite Equation (10) as:

$$\mathbb{P}(|X - \mu| \geq \gamma) \leq 2e^{-\frac{2\gamma^2}{n}}. \quad (19)$$

$\mathbb{E}(X_j)$ will be at least $\frac{\alpha c}{\epsilon e_j}$. Regarding Equation (19), we have:

$$\begin{aligned} P(|X_j - \mathbb{E}(X_j)| \geq \frac{\alpha c}{\epsilon e_j} - \frac{1}{\epsilon e_i}) &\leq 2e^{-\frac{2\epsilon}{\alpha} \left(\frac{\alpha c}{\epsilon e_j} - \frac{1}{\epsilon e_i}\right)^2} \\ \Rightarrow P(|X_j - \mathbb{E}(X_j)| \geq \frac{\alpha c}{\epsilon e_j} - \frac{1}{\epsilon e_i}) &\leq 2e^{-\frac{2}{\epsilon \alpha} \left(\frac{\alpha c}{e_j} - \frac{1}{e_i}\right)^2}. \end{aligned}$$

Let α_j be the value that $P(|X_j - \mathbb{E}(X_j)| \geq \frac{\alpha_j c}{\epsilon e_j} - \frac{1}{\epsilon e_i}) \leq \frac{1}{mn^2}$, then:

$$\begin{aligned}
 & \frac{-2}{2e^{\epsilon \alpha_j}} \left(\frac{\alpha_j c}{e_j} - \frac{1}{e_i} \right)^2 \leq \frac{1}{mn^2} \\
 \Rightarrow & \frac{-2}{2e^{\epsilon \alpha_j}} \left(\frac{\alpha_j c}{e_j} - \frac{1}{e_i} \right)^2 \leq \frac{\epsilon}{\alpha_j n^2} \\
 \Rightarrow & \frac{2}{\epsilon \alpha_j} \left(\frac{\alpha_j c}{e_j} - \frac{1}{e_i} \right)^2 \geq \ln\left(\frac{2\alpha_j n^2}{\epsilon}\right) \\
 \Rightarrow & \frac{2\alpha_j c^2}{\epsilon e_j^2} + \frac{2}{\epsilon \alpha_j e_i^2} - \frac{4c}{\epsilon e_i e_j} \geq \ln\left(\frac{2\alpha_j n^2}{\epsilon}\right) \\
 \Rightarrow & \frac{2\alpha_j c^2}{\epsilon e_j^2} + \frac{2}{\epsilon \alpha_j e_i^2} - \ln(\alpha_j) \geq \frac{4c}{\epsilon e_i e_j} + \ln\left(\frac{2n^2}{\epsilon}\right).
 \end{aligned}$$

Since the right hand side of the inequality is a constant, and the left hand side is an increasing function on its domain, we can find an α'_j such that whenever $\alpha_j \geq \alpha'_j$, this inequality holds. Therefore, $V_i(B_j) \geq \frac{1}{\epsilon e_i}$ with the probability at least $1 - \frac{1}{mn^2}$ whenever $\alpha_j \geq \alpha'_j$. Thus, by choosing proper α such that for all $a_j \in \mathcal{N}$, $\alpha \geq \alpha'_j$, WMMS_i would not be less than $\frac{1}{\epsilon}$ with the probability at least $(1 - \frac{1}{mn^2})^n \geq 1 - \frac{1}{mn}$. \square

6. Discussion

In this work we conduct a study of fair allocation when the agents have different entitlements. The original notion of maxmin share is proposed for the case where all agents are the same in terms of the entitlements. We extend this notion to the case of different entitlement and show that unlike the symmetric case, when the entitlements are different, finding an almost WMMS allocation is not always possible. More precisely, we show that the best allocation that one can hope for is a $1/n$ -WMMS allocation and that such an allocation can be obtained via a somewhat round-robin procedure.

Our experimental results suggest that in reality WMMS allocation is very likely to exist. Of course this is in contrast to our theoretical results. Therefore, it is important to study the problem under reasonable restrictions that rule out the unlikely worst-case scenarios. We initiate this study by considering two limitations to the problem and show substantially better results for these restricted cases. First, we show that if the valuations of the agents for the items are limited by their maxmin share, one can devise a greedy algorithm to achieve a $1/2$ -WMMS allocation. We then proceed by studying the case where the valuations of the agents for the items are drawn from known distributions. We show that in these cases, a WMMS allocation exists with high probability.

Although these observations partially justify our empirical results, it seems that the problem is not yet well understood when it comes to real-world settings. Therefore, we believe that future work can investigate this problem under other reasonable and realistic assumptions and further explain why *almost fair* allocations can be guaranteed in practice.

References

- Alijani, R., Farhadi, M., Ghodsi, M., Seddighin, M., & Tajik, A. S. (2017). Envy-free mechanisms with minimum number of cuts. In *AAAI*, pp. 312–318.
- Amanatidis, G., Markakis, E., Nikzad, A., & Saberi, A. (2015). Approximation algorithms for computing maximin share allocations. In *ICALP*.
- Bezakova, I., & Dani, V. (2005). Allocating indivisible goods. *ACM SIGecom Exchanges*, 5(3), 11–18.
- Brams, S. J., Jones, M. A., & Klamler, C. (2008). Proportional pie-cutting. *International Journal of Game Theory*, 36(3-4), 353–367.
- Brams, S. J., & Taylor, A. D. (1996). *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press.
- Budish, E. (2011). The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6), 1061–1103.
- Caragiannis, I., Kurokawa, D., Moulin, H. C., Procaccia, A. D., Shah, N., & Wang, J. (2016). The unreasonable fairness of maximum nash welfare. In *ACM EC*.
- Caragiannis, I., Lai, J. K., & Procaccia, A. D. (2011). Towards more expressive cake cutting. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, Vol. 22, p. 127.
- Chen, Y., Lai, J. K., Parkes, D. C., & Procaccia, A. D. (2013). Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1), 284–297.
- Cohler, Y. J., Lai, J. K., Parkes, D. C., & Procaccia, A. D. (2011). Optimal envy-free cake cutting. In *AAAI*.
- Dickerson, J. P., Goldman, J. R., Karp, J., Procaccia, A. D., & Sandholm, T. (2014). The computational rise and fall of fairness. In *AAAI*, pp. 1405–1411. Citeseer.
- Edmonds, J., & Pruhs, K. (2006). Cake cutting really is not a piece of cake. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pp. 271–278. Society for Industrial and Applied Mathematics.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301), 13–30.
- Kash, I., Procaccia, A. D., & Shah, N. (2014). No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research*, 51, 579–603.
- Kurokawa, D., Procaccia, A. D., & Wang, J. (2016). When can the maximin share guarantee be guaranteed?. In *AAAI*, Vol. 16, pp. 523–529.
- McAvaney, K., Robertson, J., & Webb, W. (1992). Ramsey partitions of integers and pair divisions. *Combinatorica*, 12(2), 193–201.
- Medina, M. (2014). *Ebay Data Set*. <http://cims.nyu.edu/~munoz/data/>.
- Procaccia, A. D. (2013). Cake cutting: not just child’s play. *Communications of the ACM*, 56(7), 78–87.
- Procaccia, A. D., & Goldman, J. (2014). *Spliddit*. <http://www.spliddit.org/apps/credit>.

- Procaccia, A. D., & Wang, J. (2014). Fair enough: Guaranteeing approximate maximin shares. In *ACM EC*.
- Robertson, J., & Webb, W. (1998). *Cake-cutting algorithms: Be fair if you can*. AK Peters/CRC Press.
- Robertson, J. M., & Webb, W. A. (1997). Extensions of cut-and-choose fair division. *Elemente der Mathematik*, 52(1), 23–30.
- Steinhaus, H. (1948). The problem of fair division. *Econometrica*, 16(1).
- Wang, W., Liang, B., & Li, B. (2015). Multi-resource fair allocation in heterogeneous cloud computing systems. *IEEE Transactions on Parallel and Distributed Systems*, 26(10), 2822–2835.
- Woeginger, G. J., & Sgall, J. (2007). On the complexity of cake cutting. *Discrete Optimization*, 4(2), 213–220.