Fair Allocation of Indivisible Goods to Asymmetric Agents

Alireza Farhadi
University of Maryland
farhadi@cs.umd.edu

Mohammad Ghodsi
Sharif University of Technology
IPM - Institute for Research in Fundamental Sciences
ghodsi@sharif.edu

MohammadTaghi Hajiaghayi
University of Maryland
hajiagha@cs.umd.edu

Sebastien Lahaie
Google Research
slahaie@microsoft.com

David Pennock
Microsoft Research
dpennock@microsoft.com

Masoud Seddighin
Sharif University of Technology
mseddighin@ce.sharif.edu

Saeed Seddighin
University of Maryland
sseddigh@cs.umd.edu

Hadi Yami
University of Maryland
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 [17] 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 w.l.o.g, 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*.

CCS Concepts

•Computing methodologies → Multi-agent systems;

Keywords

fairness, indivisible, entitlements, proportionality, approximation, stochastic

Appears in: *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

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. This problem dates back to 1948 when Steinhaus [20] 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 ramifications of the cake cutting problem have been studied since then, many of which have not been settled after decades [5, 18, 9, 16, 12, 10, 21, 8, 1]. 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 to the cake such that $\sum e_i = 1$, and the goal is to cut the cake into disproportional 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 *et al.* [4] showed that at least two cuts are necessary to divide the cake between the agents. Furthermore, Robertson *et al.* [19] 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 [15] 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 setting[1], the previous guarantee, that is

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

every agent should obtain $1/n$ of his valuation for all items from his allocated set, is impossible to deliver. To alleviate this problem, Budish [6] proposed a concept of fairness for the allocation of indivisible goods namely *the maxmin 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 [17] (EC'14) 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 $2/3\mathsf{MMS}_i$. Such an allocation is called a $2/3$-MMS allocation. Amanatidis, Markakis, Nikzad, and Saberi [2] 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 items is drawn independently from a uniform distribution. A generalized form of this result was later proposed by Kurokawa *et al.* [14] for arbitrary distributions. In a recent work, Caragiannis *et al.* [7] 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.

For fairly allocating indivisible items to agents with different entitlements, two procedures are proposed in [5]. 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 mentioned in [5] is based on *method of markers* developed by William F. Lucas which is spiritually similar 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, in the indivisible case, 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 $p \cdot e_j/e_i$ for any other agent a_j, so that his own profit per entitlement is a lower bound for all agents. Therefore, a fair answer to this question is 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 the 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 main result is in contrast to that of Procaccia and Wang. We settle the above question by giving a $1/n$ hardness 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 3.1). 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[2]. Note that since those auctions are *truthful*[3], it is the users' best interest to bid their actual valuations for the items and thus the market is transparent. 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 [6]. 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

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

One can interpret the maxmin share of an agent as his outcome as a divider in a divide-and-choose procedure against adversaries [6]. 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^i_A for any allocation $A = \langle A_1, A_2, \ldots, A_n \rangle$ as

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

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

$$\mathsf{MMS}_i = F^i_{A^*}(V_i(\mathcal{M})/n). \tag{2}$$

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:

$$\begin{aligned}\mathsf{MMS}_i &= F^i_{A^*}(V_i(\mathcal{M})/n) \\ &= \frac{\min_j V_i(A^*_j)}{V_i(\mathcal{M})/n}(V_i(\mathcal{M})/n) = \min_j V_i(A^*_j)\end{aligned}$$

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 $A = \langle A_1, A_2, \ldots, A_n \rangle$ as

$$F^i_A = \min_j \frac{V_i(A_j)}{V_i(\mathcal{M})e_j} \tag{3}$$

Let $A^* = \langle A^*_1, A^*_2, \ldots, A^*_n \rangle$ be an allocation by a_i that maximizes $F^i_{A^*}$. The weighted maxmin share of agent a_i is defined in the same way as MMS_i, that is:

$$\mathsf{WMMS}_i = F^i_{A^*} V_i(\mathcal{M}) e_i = e_i \min_j \frac{V_i(A^*_j)}{e_j}$$

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

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

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 $F_A = \min(\frac{9}{24\cdot(1/3)}, \frac{15}{24\cdot(2/3)})$ which means $F_A = 15/16$. Moreover, for allocation $A' = \langle \{b_1, b_2\}, \{b_3, b_4, b_5\} \rangle$, we have $F_{A'} = \min(\frac{8}{24\cdot(1/3)}, \frac{16}{24\cdot(2/3)})$ which means $F_{A'} = 1$. Thus, A' is a fairer allocation than A. In addition, A' is the fairest possible allocation and hence,* $\mathsf{WMMS}_1 = 1\cdot 24\cdot 1/3 = 8$.

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 $\alpha - \mathsf{WMMS}$, if the total value of the share allocated to each agent a_i is worth at least $\alpha\mathsf{WMMS}_i$ to him.

[2]http://cims.nyu.edu/ munoz/data/

[3]An action is called truthful, if no bidder has any incentive to misrepresent his valuation

REFERENCES

[1] R. Alijani, M. Farhadi, M. Ghodsi, M. Seddighin, and A. S. Tajik. Envy-free mechanisms with minimum number of cuts. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 312–318, 2017.

[2] G. Amanatidis, E. Markakis, A. Nikzad, and A. Saberi. Approximation algorithms for computing maximin share allocations. In *ICALP 2015*.

[3] I. Bezáková and V. Dani. Allocating indivisible goods. *ACM SIGecom Exchanges*, 5(3):11–18, 2005.

[4] S. J. Brams, M. A. Jones, and C. Klamler. Proportional pie-cutting. *International Journal of Game Theory*, 36(3-4):353–367, 2008.

[5] S. J. Brams and A. D. Taylor. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press, 1996.

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

[7] I. Caragiannis, D. Kurokawa, H. C. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum nash welfare. In *ACM EC 2016*.

[8] I. Caragiannis, J. K. Lai, and A. D. Procaccia. Towards more expressive cake cutting. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22, page 127, 2011.

[9] 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.

[10] Y. J. Cohler, J. K. Lai, D. C. Parkes, and A. Procaccia. Optimal envy-free cake cutting. 2011.

[11] J. P. Dickerson, J. R. Goldman, J. Karp, A. D. Procaccia, and T. Sandholm. The computational rise and fall of fairness. In *AAAI*, pages 1405–1411. Citeseer, 2014.

[12] J. Edmonds and K. Pruhs. Cake cutting really is not a piece of cake. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 271–278. Society for Industrial and Applied Mathematics, 2006.

[13] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.

[14] D. Kurokawa, A. D. Procaccia, and J. Wang. When can the maximin share guarantee be guaranteed? 2015.

[15] K. McAvaney, J. Robertson, and W. Webb. Ramsey partitions of integers and pair divisions. *Combinatorica*, 12(2):193–201, 1992.

[16] A. D. Procaccia. Cake cutting: not just child's play. *Communications of the ACM*, 56(7):78–87, 2013.

[17] A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. In *ACM EC 2014*.

[18] J. Robertson and W. Webb. Cake-cutting algorithms: Be fair if you can. 1998.

[19] J. M. Robertson and W. A. Webb. Extensions of cut-and-choose fair division. *Elemente der Mathematik*, 52(1):23–30, 1997.

[20] H. Steinhaus. The problem of fair division. *Econometrica*, 16(1), 1948.

[21] G. J. Woeginger and J. Sgall. On the complexity of cake cutting. *Discrete Optimization*, 4(2):213–220, 2007.