

Maximin-Aware Allocations of Indivisible Goods

Hau Chan^{*1}, Jing Chen², Bo Li² and Xiaowei Wu³

¹Department of Computer Science and Engineering, University of Nebraska-Lincoln, USA

²Department of Computer Science, Stony Brook University, USA

³Faculty of Computer Science, University of Vienna, Austria

hchan3@unl.edu, {jingchen, boli2}@cs.stonybrook.edu, xiaowei.wu@univie.ac.at

Abstract

We study envy-free allocations of indivisible goods to agents in settings where each agent is unaware of the goods allocated to other agents. In particular, we propose the maximin aware (MMA) fairness measure, which guarantees that every agent, given the bundle allocated to her, is aware that she does not envy at least one other agent, even if she does not know how the other goods are distributed among other agents. We also introduce two of its relaxations, and discuss their egalitarian guarantee and existence. Finally, we present a polynomial-time algorithm, which computes an allocation that approximately satisfies MMA or its relaxations. Interestingly, the returned allocation is also $\frac{1}{2}$ -approximate EFX when all agents have sub-additive valuations, which improves the algorithm in [Plaut and Roughgarden, 2018].

1 Introduction

In the last few years or so, there has been a tremendous demand for fair division services to provide systematic and *fair* ways of dividing a set of (indivisible) goods such as rooms, courses, and properties among a group of agents so that the agents do not envy each other. Such demand gave rise to, for examples, Spliddit¹, the University of Pennsylvania’s Course Match², and Fair Outcomes, Inc³ which are all based on mathematical and fairness notions. To capture the fairness of an allocation, which is arguably initiated by the work of [Foley, 1967], *envy-freeness (EF)* (and its relaxations, such as EF1 and EFX) is often used to ensure that each agent should not envy or prefer the allocated goods of other agents.

In this paper, we study an envy-free allocation domain where the planner of the division tasks wishes to withhold allocation information of others from the user or the user simply does not know the allocation of others in the system. There are a couple of good reasons why it is desirable for the planner to withhold such information. First, in many private

fair allocations of goods such as tasks or gifts, the planner requires the system to preserve anonymity as not to give away the received bundles of other agents. Second, due to the large number of (unrelated) agents and items that could be potentially be involved in the division tasks (e.g., on the Internet such as MTurks), it is not meaningful for the planner to provide such information due to various reasons. Motivated by this domain, we focus on answering the following questions.

When indivisible goods are to be allocated among unaware agents, what is the appropriate envy-free notion and how efficiently can the allocation be found subject to the notion?

Proportionality (PROP) and maximin share (MMS) [Budish, 2011] are two widely studied and well accepted fair allocation notions, both of which are defined for unaware agents. In PROP, it is required that the value of every agent’s bundle is at least a $\frac{1}{n}$ fraction of her value for the whole goods, where n is the number of agents. It is well known that there may not exist an allocation that has guaranteed approximation of proportionality, thus a weaker and more realistic notion is desired for indivisible goods. MMS is a proper relaxation of PROP, which studies an adversarial situation: when the goods are partitioned into bundles and an agent would always get the least preferred bundle of goods, what is the best way she can partition the goods? The value of such a bundle is the MMS value of the agent. In addition to its non-existence result, MMS allocation only guarantees each agent’s best minimum value, and the value of some agent’s bundle can still be the least compared with others, which may cause significant envy (demonstrated by the below example).

Example 1.1 *Suppose there are 2 agents and 2 goods, $v_1(\{1\}) = H$, $v_1(\{2\}) = 1$ and $v_2(\{1\}) = 1$, $v_2(\{2\}) = H$, where H is a sufficiently large number and v_1 and v_2 are valuation functions. Then each agent’s MMS value is 1 and $A_1 = \{2\}$, $A_2 = \{1\}$ is an MMS allocation, but each agent envies each other. However, if we exchange their allocations, everyone’s value can be improved to H .*

Recently, [Aziz *et al.*, 2018] introduced epistemic envy-free (EEF) notion to study unaware setting. With respect to EEF, each agent is satisfied if there exists one reallocation of the goods that she does not get among the other agents, such that her value for her bundle is at least as good as every bundle in this reallocation. This measure is not robust if the reallocations are restricted by agent’s reasoning (e.g., adversarial

*The authors are ordered alphabetically.

¹<http://www.spliddit.org/>

²<https://mba-inside.wharton.upenn.edu/course-match/>

³<https://www.fairoutcomes.com/>

settings in MMS). Moreover, it can be shown that EEF and PROP allocations (and their relaxations, e.g., removing any item) barely exist and cannot be properly approximated⁴.

Due to space limit, some proofs, examples and discussions are provided in the full version of this paper [Chan *et al.*, 2019].

Our Contributions

In this paper, we focus on modeling the envy-freeness for indivisible goods allocation as well as deriving new algorithms to find (approximately) fair allocations subject to different fairness notions in an unaware environment. Our main contributions include the following.

First, we introduce a novel fairness notion of maximin aware (MMA), which guarantees that the agent’s bundle value is at least as much as her value for some other agent’s bundle, no matter how the remaining goods are distributed, i.e., there is always somebody who gets no more than her. We provide a detailed picture of the relationship between MMA and classic fairness notions for different valuations.

Then we show that MMA is a strong requirement and cannot be guaranteed in general, and provide two relaxations of MMA: MMA1 and MMAX. We also prove that MMA1 (and MMAX) potentially has stronger egalitarian guarantee than EF1 and such an allocation is guaranteed to exist for the following situations: (1) there are at most three agents with submodular valuations and (2) there are any number of agents but all of them have identical submodular valuation. If the above requirement of submodularity is replaced by strictly increasing subadditivity, an MMAX allocation is guaranteed to exist. In contrast, MMS allocation may not exist even for three agents with additive valuations [Kurokawa *et al.*, 2018] and an EFX allocation is only known to exist when there are two agents [Plaut and Roughgarden, 2018].

We present a polynomial-time algorithm that computes an allocation such that every agent is either $\frac{1}{2}$ -approximate MMA or exactly MMAX for additive valuations. It is shown in [Plaut and Roughgarden, 2018] that a $\frac{1}{2}$ -approximate EFX allocation exists for general subadditive valuations, but finding it may need exponential time, which leaves an open question whether such an allocation can be found in polynomial time. We show that the allocation by our algorithm computes is also $\frac{1}{2}$ -approximate EFX when all agents have subadditive valuations, and thus answer this problem affirmatively.

Other Related Works

For indivisible goods allocations, EF or PROP allocations may not exist. [Budish, 2011] introduced the relaxed concept of envy-freeness up to one good (EF1). In an EF1 allocation, it is only required that each agent’s value for a bundle is at least as much as her value for every other agent’s bundle minus a single good (in the bundle). It is shown in [Lipton *et al.*, 2004] that an EF1 allocation always exists, and can be found in polynomial time (with a different name). [Caragiannis *et al.*, 2016] introduced the envy-free up to any good

⁴Consider three agents with two goods such that every agent has value 1 for each good. Then the agent that receives no good has no bounded guarantee for both PROP and EEF (and their relaxations by removing any item from the items she has not obtained).

(EFX) notion, where the comparison is made to “any” single good instead of “a” single good. The state-of-the-art results show that an EFX allocation exists in the following settings: (1) there are 2 agents, or (2) any number of agents but all have identical valuation [Plaut and Roughgarden, 2018]. It is still an open question whether an EFX allocation exists in general, even for additive valuations.

Very recently, we see many new fairness notions, such as Nash social welfare [Caragiannis *et al.*, 2016], epistemic envy-freeness [Aziz *et al.*, 2018], pairwise maximin share [Caragiannis *et al.*, 2016], and groupwise maximin share [Barman *et al.*, 2018], which are adapted for different settings. There is also a line of works making connections (or constraints) of fair allocations to graphs or networks [Abebe *et al.*, 2017; Bei *et al.*, 2017; Aziz *et al.*, 2018]. More works studying the fair allocation problem under different constraints can be found in [Ferraioli *et al.*, 2014; Bouveret *et al.*, 2017; Biswas and Barman, 2018].

Organization

We provide necessary definitions in Section 2. We introduce MMA and show its connections with classic fairness notions in Sections 3. In Sections 4 and 5, we study the two relaxations of MMA, MMA1 and MMAX. and discuss their connections with other fairness notions and their existences in different settings. Finally, in Section 6, we present a polynomial time algorithm to compute a $\frac{1}{2}$ -approximate MMAX allocation for additive valuations; and $\frac{1}{2}$ -approximate EFX allocation for subadditive valuations.

2 Preliminaries and Definitions

In the envy-free division problem on indivisible goods, there is a set $N = \{1, \dots, n\}$ of n agents and a set $M = \{1, \dots, m\}$ of m indivisible goods. Each agent $i \in N$ has a valuation function $v_i : 2^N \rightarrow \mathbb{R}_{\geq 0}$ that maps every subset of goods to a non-negative real number. We assume that v_i is normalized (i.e., $v_i(\emptyset) = 0$) and monotone (i.e., $v_i(S) \leq v_i(T)$ for every $S \subseteq T \subseteq M$) for each agent $i \in N$. For convenience we use $v(j)$ to denote $v(\{j\})$, for every valuation v and every good $j \in M$. Throughout the whole paper, we call a valuation v *additive* if $v(S) = \sum_{j \in S} v(j)$ for each $S \subseteq M$; *binary additive* (BA) if v is additive, and $v(j) \in \{0, 1\}$ for any good j ; *subadditive* (SA) if $v(S \cup T) \leq v(S) + v(T)$ for any $S, T \subseteq M$; *submodular* (SM) if for any $S \subseteq T$ and $e \in M \setminus T$, $v(S \cup \{e\}) - v(S) \geq v(T \cup \{e\}) - v(T)$.

An *allocation* $A = (A_1, \dots, A_n)$ is a partition of M among agents in N , i.e., $\bigcup_{i \in N} A_i = M$ and $A_i \cap A_j = \emptyset$ for any $i \neq j$. In other words, A_i is the set of goods allocated to agent i . Let $A_{-i} = \bigcup_{j \neq i} A_j$ be the goods not assigned to agent i . For ease of analysis, we assume that for every $j \in M$, there exists $i \in N$ such that $v_i(j) > 0$.

Below we state some standard fairness definitions.

Definition 2.1 (EF) For any $\alpha \in [0, 1]$, an allocation A is α -*envy-free* (α -EF) if for any $i, j \in N$, $v_i(A_i) \geq \alpha \cdot v_i(A_j)$. The allocation is EF when $\alpha = 1$.

Definition 2.2 (EF1) An allocation A is *envy-free up to one good* (EF1) if for any $i, j \in N$, there exists $e \in A_j$ such that $v_i(A_i) \geq v_i(A_j \setminus \{e\})$.

Definition 2.3 (EFX) For any $\alpha \in [0, 1]$, an allocation A is α -envy-free up to any good (α -EFX) if for any $i, j \in N$, $v_i(A_i) \geq \alpha \cdot v_i(A_j \setminus \{e\})$ for any $e \in A_j$. The allocation is EFX when $\alpha = 1$.

Next, we introduce another fairness notion called *maximin share* (MMS). Let $\Pi_k(S)$ be the set of all k -partitions of a set S . The maximin share of agent i on S among k players is

$$\text{MMS}_i(S, k) = \max_{\pi \in \Pi_k(S)} \min_{S \in \pi} v_i(S).$$

Definition 2.4 (MMS) For any $\alpha \in [0, 1]$, an allocation A is α -MMS if for any $i \in N$, $v_i(A_i) \geq \alpha \cdot \text{MMS}_i(M, n)$. The allocation is MMS when $\alpha = 1$.

The most relevant fairness notion to this work is the epistemic envy-freeness (EEF) proposed in [Aziz *et al.*, 2018]. An allocation A is EEF if for any i , there exists an allocation $\bar{A}^i = (\bar{A}_j^i)_{j \in N}$ such that $\bar{A}_i^i = A_i$ and $v_i(A_i) \geq v_i(\bar{A}_j^i)$ for any $j \in N$. For the purpose of presenting our results, we provide the definition of proportionality (PROP) fairness. An allocation A is PROP if for any $i \in N$, $v_i(A_i) \geq \frac{1}{n} \cdot v_i(M)$.

Finally, we say that $X \xrightarrow{\text{type}} Y$ if every X allocation is also a Y allocation when agents have *type* valuations, where X and Y are the fairness notions (e.g., MMA1 and MMS) and *type* is the function type (e.g. SA and SM).

3 Maximin-Aware Allocation

In this section, we introduce the fairness notion of maximin-aware (MMA), and study its connection to other existing fairness notions.

Definition 3.1 (MMA) For any $\alpha \in [0, 1]$, an allocation A is α -MMA if for any agent $i \in N$, $v_i(A_i) \geq \alpha \cdot \text{MMS}_i(A_{-i}, n - 1)$. The allocation is MMA when $\alpha = 1$.

There is a modeling and conceptual advantage of MMA: an MMA allocation guarantees for every agent, no matter how the goods that she does not get are distributed among the other agents, there is always an agent who gets no more than her. Compared to EEF, where the happiness of each agent depends on the existence of a reallocation of the other goods, MMA provides a robust way for the agents to reason about their reallocations of the remaining goods.

Observe that an MMS allocation need not be MMA. Recall, Example 1.1 where $A_1 = \{2\}$ and $A_2 = \{1\}$ is an MMS allocation. However, this allocation is far from being MMA since $v_1(A_1) \ll v_1(A_2)$ and $v_2(A_2) \ll v_2(A_1)$. The only MMA allocation in this example is $A_1 = \{1\}$ and $A_2 = \{2\}$.

Seemingly MMA being a stronger definition than MMS, Example 3.1 shows that this is not true.

Example 3.1 Suppose there are four agents ($N = \{1, 2, 3, 4\}$) and eight goods ($M = \{1, 2, \dots, 8\}$). Agent 1 has the additive valuation shown in Table 1.⁵ It is easy to check that $\text{MMS}_1(M, 4) = 0.5$. Let $S = \{5, 6, 7, 8\}$. Then $v_1(S) = 0.4 < \text{MMS}_1(M, 4)$. That is, for agent 1, S does not satisfy the condition of MMS. However, $\text{MMS}_1(M \setminus S, 3) = 0.4 = v_1(S)$, which satisfies the condition of MMA.

⁵It is easy to design other agents' valuations such that S is allocated to agent 1.

Goods	1	2, 3, 4	5, 6, 7, 8
Value	1	0.4	0.1

Table 1: MMA does not imply MMS in general.

However, when all agents have binary additive valuations, MMA actually implies MMS.

Lemma 3.1 $\text{MMA} \xrightarrow{BA} \text{MMS}$.

Proof: Fix any agent $i \in N$, let $m_i = v_i(M)$ be the number of goods agent i is interested in. Then $\text{MMS}_i(M, n) = k_i = \lfloor \frac{m_i}{n} \rfloor$. Fix any MMA allocation A , we show that it is also MMS. Suppose otherwise, i.e., $v_i(A_i) < k_i$. Then we have

$$\text{MMS}_i(A_{-i}, n - 1) = \lfloor \frac{m_i - v_i(A_i)}{n - 1} \rfloor \geq \lfloor \frac{m_i - \frac{m_i}{n}}{n - 1} \rfloor \geq k_i,$$

which contradicts the fact that A is MMA. Hence A is also an MMS allocation. \square

Indeed, similar to MMS, MMA is a more realistic and appropriate definition for indivisible goods allocation, in the sense that MMA is more approachable than EF and PROP for additive valuations. In the following, we show that $\text{PROP} \xrightarrow{A} \text{MMA}$. Combining with Theorem 4 of [Aziz *et al.*, 2018], it implies $\text{EF} \xrightarrow{A} \text{EEF} \xrightarrow{A} \text{PROP} \xrightarrow{A} \text{MMA}$.

Lemma 3.2 $\text{EF} \xrightarrow{A} \text{EEF} \xrightarrow{A} \text{PROP} \xrightarrow{A} \text{MMA}$

Proof: Fix any PROP allocation A and any agent i . It follows that $v_i(A_i) \geq \frac{1}{n} \cdot v_i(M)$. Since the valuations are additive,

$$\text{MMS}_i(M \setminus A_i, n - 1) \leq \frac{1}{n-1} \cdot (1 - \frac{1}{n}) \cdot v_i(M) = \frac{v_i(M)}{n}.$$

Hence, the allocation is also MMA. \square

Note that the other direction is not true: in Example 3.1, agent 1 is satisfied with $S = \{5, 6, 7, 8\}$ w.r.t. MMA, but the portion is only $\frac{0.4}{2.6} < 0.25$. On the other hand, MMA is a strong requirement for indivisible goods allocation, in the sense that it cannot be satisfied even in some simple settings. The proof of Lemma 3.3 is included in the full version.

Lemma 3.3 *MMA allocations may not exist even when the agents have identical binary additive valuation.*

4 Maximin-Aware up to One Good

In this section, we relax the MMA notion to maximin-aware up to one good (MMA1) that is analogous to EF1.

Definition 4.1 (MMA1) For any $\alpha \in [0, 1]$, an allocation A is α -MMA1 if for any i , there exists $e \in A_{-i}$ such that

$$v_i(A_i) \geq \alpha \cdot \text{MMS}_i(A_{-i} \setminus \{e\}, n - 1).$$

The allocation is MMA1 when $\alpha = 1$.

4.1 Connection to EF1 and MMS

We first note that while MMA1 is a relaxed version of MMA, it may have better egalitarian guarantee than EF1 allocations. An MMA1 allocation A guarantees that for each agent i and her favorite item $e \in A_{-i}$ (assume $e \in A_k$, $k \neq i$), $v_i(A_i)$ is

at least as large as the worst bundle in any $(n-1)$ -partition of $A_{-i} \setminus \{e\}$. However, EF1 implies that there exists an $(n-1)$ -partition of $A_{-i} \setminus \{e\}$, (i.e., $A_{i'}$ for $i' \neq k$ and $A_k \setminus \{e\}$) such that $v_i(A_i)$ is at least as large as the worst bundle, i.e. $A_k \setminus \{e\}$. We illustrate this as the following example.

Example 4.1 Consider $n+1$ agents and $2n+1$ goods, where every agent has identical additive valuation v with $v(\{j\}) = n$ for $1 \leq j \leq n$ and $v(\{j\}) = 1$ for $n+1 \leq j \leq 2n+1$. Thus allocation $A_i = \{i, n+i\}$ for $1 \leq i \leq n$ and $A_{n+1} = \{2n+1\}$ is EF1. In such an allocation, agent $n+1$ only gets value 1. However, there is an allocation (for example: $A_i = \{i\}$ for $1 \leq i \leq n$ and $A_{n+1} = \{n+1, \dots, 2n+1\}$) where everyone obtains value at least n .⁶

In Example 4.1, an EF1 allocation can be arbitrarily bad to agent $n+1$. However, we show that any MMA1 allocation A in this example guarantees that each agent's value is at least $\lceil \frac{n}{2} \rceil$. If everyone's value is at least n , then our claim trivially holds. Suppose for some agent i , $v(A_i) = k < n$. Then among all goods in A_{-i} , there are n goods with value n and $n+1-k$ goods with value 1. Thus by removing the most favorite good, the maximin share for the remaining set is $\min\{n, n+1-k\}$. Since A is an MMA1 allocation, $v(A_i) = k \geq n+1-k$, which means $v(A_i) \geq \lceil \frac{n}{2} \rceil$.

In the following, we show that MMS implies MMA1 when the agents's valuations are submodular.

Lemma 4.1 $MMS \xrightarrow{SM} MMA1$.

Proof: Let A be an MMS allocation. It follows that $v_i(A_i) \geq MMS_i(M, n)$ for every agent i . Fix any agent $i \in N$, it suffices to show that there exists $e \in A_{-i}$ such that $v_i(A_i) \geq MMS(A_{-i} \setminus \{e\}, n-1)$.

If for every $e \in A_{-i}$, $v_i(A_i \cup \{e\}) = v_i(A_i)$, then by the submodularity of v , we have

$$v_i(M) \leq v_i(A_i) + \sum_{e \in A_{-i}} (v_i(A_i \cup \{e\}) - v_i(A_i)) = v_i(A_i),$$

where the inequality holds since the marginal contribution of each good does not increase. Thus A is also an MMA1 allocation for agent i .

Otherwise, let $e \in A_{-i}$ be such that $v_i(A_i \cup \{e\}) > v_i(A_i)$. For the sack of contradiction, suppose there exists a partition $B = (B_j)_{j \in N \setminus \{i\}}$ of $A_{-i} \setminus \{e\}$ such that $v_i(B_j) > v_i(A_i)$ for all $j \in N \setminus \{i\}$. Then if we include e into A_i , the resulting partition of M (into n parts) has minimum value strictly larger than $MMS_i(M, n)$ on every agent, which contradicts the definition of $MMS_i(M, n)$. \square

By a similar argument, we can show that every α -MMS allocation is also α -MMA1. Recall Example 3.1 where we show an MMA allocation is not necessarily MMS (thus neither an MMA1 allocation), this implication is strict.

Next, we show that for binary additive valuations, MMA1 and MMS are actually equivalent.

Lemma 4.2 $MMA1 \xleftrightarrow{BA} MMS$.

Proof: Since binary additive functions are submodular, it suffices to prove the " \Rightarrow " direction. Fix any MMA1 allocation

$A = (A_i)_{i \in N}$, we prove that $v_i(A_i) \geq k_i$ for every agent i , where $k_i = MMS_i(M, n) = \lfloor \frac{v_i(M)}{n} \rfloor$.

Fix any agent i and assume the contrary that $v_i(A_i) < k_i$. Let $m_i = v_i(M)$. Then for any $e \in A_{-i}$ we have

$$\begin{aligned} MMS_i(A_{-i} \setminus \{e\}, n-1) &= \lfloor \frac{m_i - 1 - v_i(A_i)}{n-1} \rfloor \\ &\geq \lfloor \frac{m_i - 1 - k_i + 1}{n-1} \rfloor \geq k_i > v_i(A_i), \end{aligned}$$

which contradicts the fact that A is MMA1. \square

It is shown in [Kurokawa *et al.*, 2018] that even for three agents with additive valuations, MMS allocations do not always exist. As we will show later, for a broader class of situations, an MMA1 allocation is guaranteed to exist.

4.2 The Existence of MMA1 Allocations

By Lemma 4.2, we know that an MMA1 allocation always exists when the valuations are binary additive. In this section, we explore the existence of MMA1 allocations for instances with subadditive valuations.

First, for any valuation v over a set M of goods, we define the *leximin n -partition* as follows. A leximin n -partition is a partition which divides M into n subsets and maximizes the lexicographical order when the values of the partitions are sorted in non-decreasing order. In other words, a leximin n -partition maximizes the minimum value over all n -partitions, and if there is a tie it selects the one maximizing the second minimum value, and so on.

In the following, we show that MMA1 allocations exist when all agents have identical submodular valuation, or when there are three agents with (different) submodular valuations.

We first show that when all agents have identical submodular valuation, the leximin n -partition of M is indeed an MMA1 allocation.

Theorem 4.1 *When all agents have the identical submodular valuation, the leximin n -partition of M is MMA1.*

Proof: Let v be the submodular valuation and A be the leximin n -partition of M . We show that allocating each partition A_i to agent i gives an MMA1 allocation. As before, if for all $e \in A_{-i}$, $v(A_i \cup \{e\}) = v(A_i)$, then we have

$$v(A_i) = v(M) \geq MMS(A_{-i}, n-1)$$

by submodularity of v . Otherwise let $e \in A_{-i}$ be any good such that $v(A_i \cup \{e\}) > v(A_i)$.

Assume the contrary that $v(A_i) < MMS(A_{-i} \setminus \{e\}, n-1)$.

Let $B = (B_j)_{j \in N \setminus \{i\}}$ be the partition of $A_{-i} \setminus \{e\}$ such that $v(A_i) < v(B_j)$ for all $j \neq i$. Then by including e into A_i , we obtain a partition of M such that every bundle has value strictly larger than $v(A_i)$. In other words, the resulting partition is strictly better than the leximin n -partition, which is a contradiction. \square

Next we show a divide-and-choose style algorithm (shown in Algorithm 1) which gives an MMA1 allocation for three agents. For space limit, the proof of Theorem 4.2 is deferred to the full version.

Theorem 4.2 *When $n = 3$ and the valuations are submodular, Algorithm 1 computes an MMA1 allocation.*

⁶We note that there is no exact EF allocation for this example.

Algorithm 1 Divide-and-choose algorithm for three agents

1. Agent 3 divides M as a leximin partition to her, which is denoted by $A = (A_1, A_2, A_3)$.
2. If agent 1's and agent 2's favorite bundles are different, each of them takes her favorite bundle and the remaining bundle is allocated to agent 3.
3. Otherwise (agent 1 and agent 2 have the same favorite bundle), we compare their second favorite bundle. If their second favorite bundles are also the same, agent 2 takes her favorite two bundles (w.o.l.g, say A_1 and A_2), and leaves A_3 to agent 3. Next, agent 2 repartitions $A_1 \cup A_2$ by her leximin 2-partitions, denoted by B_1 and B_2 . Agent 1 chooses her favorite one in B_1 and B_2 , and the other one is allocated to agent 2.
4. Finally, suppose their second favorite bundles are different. W.o.l.g, suppose $v_1(A_1) \geq v_1(A_2) \geq v_1(A_3)$ and $v_2(A_1) \geq v_2(A_3) \geq v_2(A_2)$.
 - (1) If $2 \cdot v_1(A_2) > v_1(A_1) + v_1(A_3)$ and $2 \cdot v_2(A_3) > v_2(A_1) + v_2(A_2)$, allocate A_2 to agent 1, A_3 to agent 2 and A_1 to agent 3.
 - (2) Otherwise (assume w.l.o.g. that $2 \cdot v_1(A_2) \leq v_1(A_1) + v_1(A_3)$). Agent 2 takes A_1 and A_3 , which are her favorite two bundles, and leave A_2 to agent 3. Next, agent 2 repartitions $A_1 \cup A_3$ by her leximin 2-partition $C = (C_1, C_2)$, and let Agent 1 choose her favorite one in C_1 and C_2 . The other one is allocated to agent 2.

5 Maximin-Aware up to Any Good

Definition 5.1 (MMAX) For any $\alpha \in [0, 1]$, an allocation A is α -MMAX if for any i ,

$$v_i(A_i) \geq \alpha \cdot \text{MMS}_i(A_{-i} \setminus \{e\}, n - 1)$$

for any $e \in A_{-i}$. The allocation is MMAX when $\alpha = 1$.

We first note that MMAX allocations give an even better egalitarian guarantee than MMA1 and EF1 allocations. Recall Example 4.1, where we show there is an EF1 allocation such that some agent only gets value 1 and any MMA1 allocation guarantees each agent's value is at least $\lceil \frac{n}{2} \rceil$. It is not difficult to show that any MMAX allocation guarantees each agent's value is at least n .

Obviously, an MMAX allocation is also MMA1. Together with Lemma 4.2, we have $\text{MMAX} \xrightarrow{BA} \text{MMA1} \xleftrightarrow{BA} \text{MMS}$. However, beyond BA valuations, the above implication does not hold. We can see this by borrowing an example from [Kurokawa *et al.*, 2018], for which there is no MMS allocation, but MMAX (and MMA1) allocations do exist. For space limit we defer the example and analysis to the full version.

Next we show the connection between MMAX and EFX when the valuations are binary additive.

Lemma 5.1 $\text{EFX} \xrightarrow{BA} \text{MMAX}$.

Proof: Let A be any EFX allocation. Consider any agent $i \in N$ and let $x = v_i(A_i)$ be the value of i 's bundle under the allocation A . Since A is EFX, then for any $j \in N \setminus \{i\}$,

one of the following statements holds: (1) $v_i(A_j) = x + 1$ and $|A_j| = x + 1$ (i.e., agent i has value 1 for every good in A_j); (2) $v_i(A_j) \leq x$ (i.e., A_j may contain more than $x + 1$ goods, but agent i has value 1 for exact x goods). This is because if both of them do not hold, then $v_i(A_j) \geq x + 1$ and $|A_j| > x + 1$, which means agent i has value 0 for at least one good in A_j . Thus agent i still envies agent j after removing the good with value 0 from A_j .

Note that as long as there exists an agent $j \neq i$ such that $v_i(A_j) \leq x$, then $v_i(A_{-i}) \leq (x + 1)(n - 1) - 1$. Hence

$$\text{MMS}_i(A_{-i}, n - 1) \leq x = v_i(A_i).$$

Otherwise we have $v_i(A_{-i}) = |A_{-i}| = (x + 1)(n - 1)$. Then by removing any good $e \in A_{-i}$, we have

$$\text{MMS}_i(A_{-i} \setminus \{e\}, n - 1) \leq x = v_i(A_i).$$

In both case, agent i is satisfied w.r.t. MMAX. \square

However, when agents have general additive valuations, EFX does not imply MMAX. More specifically, there are examples (we will include one in our full version), for which EFX allocations exist but are not MMAX.

The existence of MMAX allocations. First, note that for two agents, MMAX is equivalent to EFX and thus exists for general subadditive valuations [Plaut and Roughgarden, 2018]. Indeed, following exactly the same analysis as in Theorems 4.1 and 4.2, we obtain the following result.

Theorem 5.1 *When there are at most three agents having strictly increasing subadditive valuations, or any number of agents having identical strictly increasing subadditive valuations, an MMAX allocation exists.*

6 Computing Allocations Efficiently

In this section, we provide a polynomial-time algorithm (see Algorithm 2) for computing an allocation such that when all agents have additive valuations every agent is either $\frac{1}{2}$ -MMA or MMAX. Moreover, the allocation our algorithm outputs is $\frac{1}{2}$ -EFX when all agents have subadditive valuations.

In each round (the while loop) of Algorithm 2, we assign at most one good to each agent that is not envied by any other agent. The assignment of goods to agents in each round is given by a maximum weight matching between the unenvied agents L and the unallocated goods R , where the weight of every edge is given by the marginal increase in the value, after allocating the good to the corresponding agent. If all edges between L and R have weight 0, we assign goods to unenvied agents according to some maximum cardinality matching. To make sure that there is always an unenvied agent, we use the approach (procedure \mathcal{P} in the following) introduced in [Lipton *et al.*, 2004], which is now widely used in the resource allocation literature.

Next, we define a *procedure* \mathcal{P} , which is useful in our main algorithm. Given an allocation A , we define its corresponding *envy-graph* \mathcal{G} as follows. Every agent is represented by a node in \mathcal{G} and there is a directed edge from node i to node j iff $v_i(A_i) < v_i(A_j)$, i.e., i is envies j . A directed cycle in \mathcal{G} is called an *envy-cycle*. Let $C = i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_t \rightarrow i_1$ be such a cycle. If we reallocate $A_{i_{k+1}}$ to agent i_k for all

Algorithm 2 Algorithm for (sub-)additive valuations

- 1: Initiate $L = N$ and $R = M$.
 - 2: Initiate $A_i = \emptyset$ for all $i \in N$.
 - 3: **while** $R \neq \emptyset$ **do**
 - 4: Compute a maximum weight matching \mathcal{M} between L and R , where the weight of edge between $i \in L$ and $j \in R$ is given by $v_i(A_i \cup \{j\}) - v_i(A_i)$. If all edges have weight 0, then we compute a maximum cardinality matching \mathcal{M} instead.
 - 5: For every edge $(i, j) \in \mathcal{M}$, allocate j to i : $A_i = A_i \cup \{j\}$ and exclude j from R : $R = R \setminus \{j\}$.
 - 6: As long as there is an envy-cycle w.r.t. $A = (A_i)_{i \in N}$, invoke procedure \mathcal{P} (to be described later).
 - 7: Update $A = (A_i)_{i \in N}$ to be the allocations after \mathcal{P} .
 - 8: Update the set of agents not envied by any other agents: $L = \{i \in N : \forall j \in N, v_j(A_j) \geq v_j(A_i)\}$.
 - 9: **end while**
-

$k \in [t-1]$, and reallocate A_{i_t} to agent i_t , the number of edges of \mathcal{G} will be strictly decreased. Thus, by repeatedly using this procedure, we eventually get another allocation whose envy-graph is acyclic. It is shown in [Lipton *et al.*, 2004] that \mathcal{P} can be done in time $O(mn^3)$.

Next, we prove the following main result of this section.

Theorem 6.1 *Algorithm 2 computes an allocation in polynomial time that is (1) $\frac{1}{2}$ -EFX and EF1 when all agents have subadditive valuations; (2) either $\frac{1}{2}$ -MMA or exact MMAX for each agent when the valuations are additive.*

Proof: We first show that Algorithm 2 is well defined. In each round, there must be at least one agent who is not envied by anybody, since after Step 6, the envy-graph is acyclic. Thus there will be at least one agent whose in-degree is 0, i.e., $L \neq \emptyset$. Moreover, since at least one good will be allocated in each round (each while loop), Algorithm 2 terminates in at most m rounds. Given that each round (including procedure \mathcal{P} , the computation of maximum weight matching, and the updates of edge weights and unenvied agents L) can be done in polynomial time, Algorithm 2 runs in polynomial time.

Fix any agent i . We classify other agents into three sets:

- agents not envied by i : $N_1 = \{j \in N \setminus \{i\} : v_i(A_i) \geq v_i(A_j)\}$;
- agents envied by i that receive only one good: $N_2 = \{j \in N \setminus \{i\} : v_i(A_i) < v_i(A_j) \text{ and } |A_j| = 1\}$;
- agents envied by i that receive more than one good: $N_3 = \{j \in N \setminus \{i\} : v_i(A_i) < v_i(A_j) \text{ and } |A_j| \geq 2\}$.

By definition we have $N_1 \cup N_2 \cup N_3 = N \setminus \{i\}$.

We first show that the final allocation A is $\frac{1}{2}$ -EFX and EF1 when the valuations are subadditive. Note that it suffices to consider N_3 , as the other agents are either not envied by i , or allocated only one good.

Fix any agent $j \in N_3$. Let e_j be the last good allocated to j . As we only allocate goods to unenvied agents, we have $v_i(A_i) \geq v_i(A_j \setminus \{e_j\})$, which implies that agent i is satisfied w.r.t. EF1. On the other hand, we also have $v_i(A_i) \geq v_i(e_j)$, since otherwise in the first round, matching e_j (which is not allocated yet) with i (which is not envied) gives a matching with strictly larger weight. Combing the two inequalities and

by subadditivity of valuation v_i , we have

$$2 \cdot v_i(A_i) \geq v_i(A_j \setminus \{e_j\}) + v_i(e_j) \geq v_i(A_j),$$

which implies that agent i is satisfied w.r.t. $\frac{1}{2}$ -EF, and $\frac{1}{2}$ -EFX.

Finally, we show that when the valuations are additive, agent i is also satisfied w.r.t. $\frac{1}{2}$ -MMA or MMAX.

The case when $N_2 = N \setminus \{i\}$, i.e., all other agents receive one good and is envied by agent i , is trivial: by removing any good e from A_{-i} , there is always some agent with no good. Thus we have

$$\text{MMS}_i(A_{-i} \setminus \{e\}, n-1) = 0 \leq v_i(A_i),$$

which implies that agent i is also satisfied w.r.t. MMAX.

Now suppose that $N_2 \neq N \setminus \{i\}$. In other words, we have $N_1 \cup N_3 \neq \emptyset$, and $n-1 - |N_2| \geq 1$. We prove that in this case we have $\text{MMS}_i(A_{-i}, n-1) \leq 2 \cdot v_i(A_i)$, which implies that agent i is satisfied w.r.t. $\frac{1}{2}$ -MMA.

We first prove that $\text{MMS}_i(A_{-i}, n-1)$ is at most the average value of bundles allocated to agents in $N_1 \cup N_3$:

$$\text{MMS}_i(A_{-i}, n-1) \leq \frac{\sum_{j \in N_1} v_i(A_j) + \sum_{j \in N_3} v_i(A_j)}{n-1 - |N_2|}.$$

Let $\tau = \text{MMS}_i(A_{-i}, n-1)$. Observe that if we have $v_i(e) \leq \tau$ for every $e \in A_{-i}$, then

$$\begin{aligned} \sum_{j \in N_1} v_i(A_j) + \sum_{j \in N_3} v_i(A_j) &= v_i(M) - \sum_{j \in N_2} v_i(A_j) \\ &\geq v_i(M) - |N_2| \cdot \tau \geq (n-1 - |N_2|) \cdot \text{MMS}_i(A_{-i}, n-1). \end{aligned}$$

Indeed, the assumption is w.l.o.g., since otherwise we can reset every $v_i(e)$ to $\min\{\tau, v_i(e)\}$, which does not change $\text{MMS}_i(A_{-i}, n-1)$. The reason is, in any partitioning of A_{-i} into $n-1$ bundles, the bundle with minimum value (whose value is at most $\tau = \text{MMS}_i(A_{-i}, n-1)$) is not touched.

As we have shown before (for subadditive valuations), for every $j \in N_3$, we have $v_i(A_j) \leq 2 \cdot v_i(A_i)$. Since we have $v_i(A_j) \leq v_i(A_i)$ for all $j \in N_1$, we conclude that

$$\text{MMS}_i(A_{-i}, n-1) \leq 2 \cdot v_i(A_i),$$

which implies that agent i is satisfied w.r.t. $\frac{1}{2}$ -MMA. \square

7 Conclusion

In this paper, we introduce the novel fairness notions MMA and its two relaxations MMA1 and MMAX in an unaware environment. We study their connections with other fairness notions, and propose an efficient algorithm for computing allocations that is (approximately) MMA and MMAX.

We leave the existence of MMA1 and MMAX allocations for the broader class of valuations as a future direction. Another promising direction would be to extend our work to other preference representations, including ordinal preferences, or to chores instead of goods. Finally, it will be interesting to obtain analogue concepts for asymmetric agents, where agents may have different entitlements or shares.

Acknowledgments

The authors thank Haris Aziz and David Parkes for reading a draft of this paper and for helpful discussions. This work is partially supported by NSF CAREER Award No. 1553385.

References

- [Abebe *et al.*, 2017] Rediet Abebe, Jon Kleinberg, and David C Parkes. Fair division via social comparison. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 281–289. International Foundation for Autonomous Agents and Multiagent Systems, 2017.
- [Aziz *et al.*, 2018] Haris Aziz, Sylvain Bouveret, Ioannis Caragiannis, Ira Giagkousi, and Jérôme Lang. Knowledge, fairness, and social constraints. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, 2018.
- [Barman *et al.*, 2018] Siddharth Barman, Arpita Biswas, Sanath Kumar Krishnamurthy, and Yadati Narahari. Groupwise maximin fair allocation of indivisible goods. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, 2018.
- [Bei *et al.*, 2017] Xiaohui Bei, Youming Qiao, and Shengyu Zhang. Networked fairness in cake cutting. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 3632–3638. AAAI Press, 2017.
- [Biswas and Barman, 2018] Arpita Biswas and Siddharth Barman. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 91–97, 2018.
- [Bouveret *et al.*, 2017] Sylvain Bouveret, Katarína Cechlárová, Edith Elkind, Ayumi Igarashi, and Dominik Peters. Fair division of a graph. In *26th International Joint Conference on Artificial Intelligence*, 2017.
- [Budish, 2011] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2016] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 305–322. ACM, 2016.
- [Chan *et al.*, 2019] Hau Chan, Jing Chen, Bo Li, and Xiaowei Wu. Maximin-aware allocations of indivisible goods (full version). <https://arxiv.org/abs/1905.09969>, 2019.
- [Ferraioli *et al.*, 2014] Diodato Ferraioli, Laurent Gourvès, and Jérôme Monnot. On regular and approximately fair allocations of indivisible goods. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 997–1004. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [Foley, 1967] Duncan K Foley. An improved envy-free cake cutting protocol for four agents. *Yale Economics Essays*, pages 45–98, 1967.
- [Kurokawa *et al.*, 2018] David Kurokawa, Ariel D Procaccia, and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM (JACM)*, 65(2):8, 2018.
- [Lipton *et al.*, 2004] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 125–131. ACM, 2004.
- [Plaut and Roughgarden, 2018] Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2584–2603. SIAM, 2018.