

Fair Enough: Guaranteeing Approximate Maximin Shares

ARIEL D. PROCACCIA, Carnegie Mellon University
JUNXING WANG, Tsinghua University

We consider the problem of *fairly* allocating indivisible goods, focusing on a recently-introduced notion of fairness called *maximin share guarantee*: Each player's value for his allocation should be at least as high as what he can guarantee by dividing the items into as many bundles as there are players and receiving his least desirable bundle. Assuming additive valuation functions, we show that such allocations may not exist, but allocations guaranteeing each player $2/3$ of the above value always exist, and can be computed in polynomial time when the number of players is constant. These theoretical results have direct practical implications.

Categories and Subject Descriptors: F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Fair division, Computational social choice

1. INTRODUCTION

We are interested the *fair* allocation of *indivisible* goods, but to explain the intricacies of this problem we start from discussing the easier case of *divisible* goods. In the latter setting, known as *cake cutting*, we need to divide a heterogeneous cake between players with different valuation functions (that is, different players may have different values for the same piece of cake).

When there are only two players, the *Cut and Choose* protocol provides a compelling method for dividing a cake — and will play an important conceptual role later on. Under this protocol, player 1 cuts the cake into two pieces that he values equally, and player 2 subsequently chooses the piece that he prefers, giving the other piece to player 1. The resulting allocation is fair in a precise, formal sense known as *envy-freeness*: Each player prefers his own allocation to the allocation of the other player. Envy-free cake divisions exist for any number of players; today we know exactly how many cuts are needed to achieve such allocations in the worst case [Alon 1987], and how to constructively find them [Brams and Taylor 1995] (although subtle complexity questions remain open [Procaccia 2009; 2013]). It is interesting to note that — in the standard cake-cutting setting — envy-freeness implies another natural fairness property called *proportionality*: Each player in the set \mathcal{N} receives a piece of cake whose value is at least $1/|\mathcal{N}|$ of the player's value for the entire cake.

Cake cutting is a nice metaphor for real-world problems like land division; the study of cake cutting distills insights about fairness that are useful in related settings, such as the allocation of computational resources [Ghodsi et al. 2011; Parkes et al. 2012;

Authors' addresses: A. D. Procaccia, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA, email: arielpro@cs.cmu.edu; J. Wang, Institute for Theoretical Computer Science (ITCS), Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China, email: thuwjx@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'14, June 8–12, 2014, Stanford, CA, USA. Copyright © 2014 ACM 978-1-4503-2565-3/12/06 ...\$15.00.
<http://dx.doi.org/10.1145/10.1145/2600057.2602835>

Procaccia 2013]. However, typical real-world situations where fairness is a chief concern, notably divorce settlements and the division of an estate between heirs, involve *indivisible goods* (e.g., houses, cars, and works of art) — which in general preclude envy-free, or even proportional, allocations. As a simple example, if there are several players and only one indivisible item to be allocated, the allocation cannot possibly be proportional or envy free. Foreshadowing the approach we take below, we note that no allocation can be even *approximately* (in a multiplicative sense) fair according to these notions, because some players receive an empty allocation of zero value.

So how can we divide an estate without lawyers? Potentially using an intriguing alternative to classical fairness notions, recently presented by Budish [2011] (building on concepts introduced by Moulin [1990]). Imagine that player 1 partitioned the items into $|\mathcal{N}|$ bundles, and each player in $\mathcal{N} \setminus \{1\}$ *adversarially* chose a bundle before player 1. A smart player would partition the bundles to maximize his minimum value for any bundle. For the same reason we intuitively view the Cut-and-Choose protocol as fair to player 1, even before specifying fairness axioms, the allocation that leaves player 1 with his least desired bundle seems fair to player 1 — as he is the one who divided the items in the first place. Budish calls the value player 1 can guarantee in this way his *maximin share (MMS) guarantee*.¹ But an allocation based on the division of player 1 may make another player regret the fact that he was not the one to divide the items. The question is: Can we allocate the items in a way that *all* players receive a bundle worth at least as much as their MMS guarantee? This question was recently addressed by Bouveret and Lemaître [2014], and while they were able to answer it for special cases (which we list in §1.3), they left the general question open.

1.1. Model, Conceptual Contribution, and Technical Results

Denote the set of players by \mathcal{N} and the set of indivisible items to be allocated by \mathcal{M} . Each player i is endowed with a valuation function $v_i : 2^{\mathcal{M}} \rightarrow \mathbb{R}^+$. We simplify notation by writing $v_i(j)$ instead of $v_i(\{j\})$ for an item $j \in \mathcal{M}$. We assume that the valuation functions are *additive*:

$$\forall S \subseteq \mathcal{M}, v_i(S) = \sum_{j \in S} v_i(j).$$

This assumption is also made in most of the related work on fair division of indivisible goods (see §1.3), including the paper of Bouveret and Lemaître [2014] that studies the maximin share guarantee in the same setting. And, more importantly, people find it difficult to specify combinatorial preferences, which is why some deployed implementations of fair division methods (see §1.2) rely on additive valuation functions. Finally, our positive result does not hold under larger classes of valuation functions, e.g., sub-additive and superadditive functions.

For a set $S \subseteq \mathcal{M}$, let $\Pi_n(S)$ be the set of n -partitions of S . Define the *n -maximin share (n -MMS) guarantee* of player $i \in \mathcal{N}$ as

$$\text{MMS}_i^{(n)}(S) = \max_{(T_1, \dots, T_n) \in \Pi_n(S)} \min_{j \in [n]} v_i(T_j),$$

where $[n] = \{1, \dots, n\}$; we call a partition that realizes this value player i 's *n -maximin partition of S* . The valuation function used to determine a player's MMS guarantee will be clear from the context. An *allocation* $(A_1, \dots, A_{|\mathcal{N}|}) \in \Pi_{|\mathcal{N}|}(\mathcal{M})$ allocates the subset of items A_i to each player i . We say that $(A_1, \dots, A_{|\mathcal{N}|})$ is a *maximin share (MMS)*

¹This term should not be confused with the terminology of the systems literature, where max-min fairness simply refers to maximizing the value any player receives [Demers et al. 1989] rather than an axiomatic notion of fairness.

allocation if and only if

$$\forall i \in \mathcal{N}, v_i(A_i) \geq \text{MMS}_i^{(|\mathcal{N}|)}(\mathcal{M}).$$

Our first result is negative:

Theorem 2.1. *For any set of players \mathcal{N} such that $|\mathcal{N}| \geq 3$ there exist \mathcal{M} and (additive) valuation functions that do not admit an MMS allocation.*

We find this theorem surprising because extensive automated experiments by several groups of researchers (including us) have failed to find a counterexample. Indeed, the counterexample relies on a very precise construction. In §2 we first provide an explicit counterexample for the case of three players (relying on a Sudoku-like construction) that illustrates the key ideas, and then give the full proof.

While it seems that MMS allocations almost always exist, we wish to relax this fairness notion in order to guarantee existence. Fortunately, unlike other fairness notions such as envy-freeness and proportionality, the MMS guarantee supports a multiplicative notion of approximation. Our main question is:

Is there a constant $c > 0$ such that we can always find an allocation A_1, \dots, A_n that satisfies $v_i(A_i) \geq c \cdot \text{MMS}_i^{(|\mathcal{N}|)}(\mathcal{M})$?

We answer this question in the positive for $c = 2/3$:

Theorem 3.1. *Let there be $|\mathcal{N}| \geq 1$ players and a set of items \mathcal{M} . Then there exists an allocation $A_1, \dots, A_{|\mathcal{N}|}$ such that for all $i \in \mathcal{N}$, $v_i(A_i) \geq \frac{2}{3} \text{MMS}_i^{(|\mathcal{N}|)}(\mathcal{M})$. Moreover, such an allocation can be found in polynomial time if $|\mathcal{N}|$ is constant.*

In fact, we prove a slightly stronger approximation ratio that converges to $2/3$ as $|\mathcal{N}|$ goes to infinity; for the important cases of three and four players, the ratio is $3/4$. Our technical approach relies on an intricate lemma — the *Density Balance Lemma* — which relates the value of a subset of items to the value of its complement. We use the lemma to design an exponential-time recursive algorithm (in §3.1), which lets one of the players compute a maximin partition (an NP-hard problem), and creates a bipartite graph between players and bundles where an edge exists if a bundle satisfies the player. We show (using Hall’s Theorem) that a perfect matching exists between a nonempty subset of the players and some of the bundles. The algorithm then recurses on the remaining players and items. In §3.2 we show that the algorithm can be adapted to run in polynomial time if the number of players is constant, by leveraging a result of Woeginger [1997].

1.2. Practical Applications of Our Results

The theory of fair division has been extensively studied, as shown, e.g., by the books by Moulin [2003] and Brams and Taylor [1996]. Despite the abundance of extremely clever fair division algorithms, very few have been implemented. Budish’s [2011] work is a rare example; his method is currently used for MBA course allocation at the Wharton School of the University of Pennsylvania. Another example is the *adjusted winner* method [Brams and Taylor 1996], which assumes that there are exactly two players (with additive valuation functions). Adjusted winner has been patented by NYU and licensed to *Fair Outcomes, Inc.* It can be used at a free NYU website.²

One of us (Procaccia) has been leading an effort to change this situation by building a not-for-profit fair-division website, tentatively called *Spliddit*; its purpose is to educate the public, gather data on fair division, and hopefully make the world just a bit

²<http://www.nyu.edu/projects/adjustedwinner/>

fairer. The website — spliddit.org — contains implementations of mechanisms for rent division [Abdulkadiroğlu et al. 2004] and assignment of scientific credit [de Clippeel et al. 2008]. However, for the third application — dividing indivisible goods — we were unable to find satisfactory methods for more than two players, despite discussions with leading experts on fair division such as Steven Brams and Hervé Moulin (we survey some existing methods in §1.3). This provided strong motivation for the theoretical work reported here.

The approach we ultimately implemented relies heavily on Theorem 3.1. We consider three “levels” of fairness: envy-freeness, proportionality, and (approximate) MMS guarantee. It is easy to verify that each of these fairness notions implies the ones following it. Users specify their valuation functions by distributing a fixed pool of points between the items. We then find an allocation that maximizes social welfare — $\sum_{i \in \mathcal{N}} v_i(A_i)$ — subject to the strongest feasible fairness constraint (using an integer linear programming formulation, which is solved via CPLEX). For MMS, we maximize the value of c for which the c -MMS guarantee is feasible. By Theorem 3.1, achieving $2/3$ of the MMS guarantee is always feasible, so the theorem ensures an outcome that is, well, fair enough. In our view, this method is (arguably) the most compelling method to date for fair division settings involving indivisible goods and more than two players.

1.3. Related Work

Motivated by the problem of allocating courses to students, Budish [2011] studies a solution concept that he calls *approximate competitive equilibrium from equal incomes (CEEI)*. Budish shows the existence of an approximate CEEI (with certain approximation parameters), even when the preferences of players are unrestricted (so they may correspond to any combinatorial valuation functions). Roughly speaking, an approximate CEEI guarantees that $v_i(A_i) \geq \text{MMS}_i^{(|\mathcal{N}|+1)}(\mathcal{M})$, that is, each of the $|\mathcal{N}|$ players receives its $(|\mathcal{N}| + 1)$ -MMS guarantee. However, this result takes advantage of an approximation error in the items that are allocated (some items might be in excess demand or excess supply). The approximation error grows with the overall number of items, and with the number of items demanded by each player, but not with the number of players or the number of copies of each item. Therefore, as the two latter parameters go to infinity, the error goes to zero. A large economy, in this sense, is plausible in the context of MBA course allocation, because there are many MBA students, many seats in each course, but relatively few courses that are offered, and even fewer courses a single student can take. But Budish’s results do not provide practical guarantees when there are, say, three or four players, and (very possibly) only one copy of each item — which is the setting we are interested in.

Like us, Bouveret and Lemaître [2014] focus on the division of indivisible goods between players with additive valuations. They study a hierarchy of fairness properties, of which the maximin share guarantee is the weakest (it is easy to see that allocations satisfying the other properties may not exist). Among other results, they show that MMS allocations exist in the following cases: (i) valuations for items are 0 or 1; (ii) the values different players assign to items form identical multisets; and (iii) $|\mathcal{M}| \leq |\mathcal{N}| + 3$. They also present results from extensive simulations using different distributions over item values; MMS allocations exist in each and every trial.

Also related is the work of Lipton et al. [2004]. Among other results, they give a polynomial-time algorithm that computes approximately envy-free allocations, where the approximation is *additive*. Specifically, they let α be the largest possible increase in value a player can have from adding one item to his bundle, and produce an allocation such that $v_i(A_i) \geq v_i(A_j) - \alpha$ for all $i, j \in \mathcal{N}$. This interesting result may not be very practical; for example, if one of the items is extremely valuable, the players would not

be guaranteed anything. In contrast, assuming items have positive values, an MMS allocation (or any multiplicative approximation thereof) gives some player a bundle worth zero (if and) only if *any* allocation gives some player a bundle worth zero.

Hill [1987] shows that when valuations are additive, indivisible items can be allocated in a way that a certain value is guaranteed to each player; and Markakis and Psomas [2011] refine this guarantee and construct a polynomial time algorithm that achieves it. However, the guaranteed value is defined using an unwieldy function that depends on the number of players as well as on the value of the most valuable item, and even for three players the function’s value quickly goes down to zero as the most valuable item becomes more valuable.

When there are exactly two players, practical methods for dividing indivisible goods are available. For example, recent work by Brams et al. [2014] gives a method satisfying several desirable properties, including envy-freeness; its main shortcoming is that it may not allocate all items (it generates a “contested pile” of unallocated items). The *adjusted winner* method [Brams and Taylor 1996], mentioned above, is another practical method (which is routinely being used, as discussed in §1.2) — but it implicitly assumes that the items are divisible and would typically require splitting one of the items. In any case, for more than two players, one encounters a great many paradoxes when contemplating standard fairness notions [Brams et al. 2003]. Moreover, generalizing these practical 2-player protocols is impossible; for example, adjusted winner can be interpreted as a special case of the *egalitarian equivalent* [Pazner and Schmeidler 1978] rule (for two players and additive valuation functions), but the latter method strongly relies on divisibility and may end up splitting all goods.

From an algorithmic viewpoint, our work is related to papers on the problem of allocating indivisible goods to maximize the minimum value any player has for his bundle (under additive valuation functions) — also known as the *Santa Claus* problem [Bezáková and Dani 2005; Bansal and Sviridenko 2006; Asadpour and Saberi 2007]. Woeginger [1997] studies the special case of players with identical valuations, and presents a polynomial time approximation scheme that we leverage below.

Somewhat further afield, recent years have seen quite a bit of computational work on cake cutting; see [Procaccia 2013] for an overview. One question that received some attention from the theoretical computer science community is the complexity of proportional and envy-free cake cutting in a concrete complexity model [Magdon-Ismail et al. 2003; Edmonds and Pruhs 2006b; 2006a; Woeginger and Sgall 2007; Procaccia 2009]

1.4. Open Problems

One obvious question remains open. Theorem 2.1 does not provide an upper bound on the the constant $c > 0$ such that c -MMS allocations always exist, and even the three-player construction in §2 provides a very weak upper bound. Our lower bound, given by Theorem 3.1, is $2/3$. Lemma 3.2 shows that our technical approach cannot give a better lower bound. Narrowing this gap is, in our view, an important challenge.

As noted above, Budish [2011] introduced a different notion of MMS approximation. In its ideal form, we would ask for an allocation such that $v_i(A_i) \geq \text{MMS}_i^{(|N|+1)}(\mathcal{M})$. We have designed an algorithm that achieves this guarantee for the case of three players (it is already nontrivial). Proving or disproving the existence of such allocations for a general number of players remains an open problem; a positive result would provide a compelling alternative to Theorem 3.1.

Finally, our negative result, Theorem 2.1, requires a number of items that is exponential in the number of players. In contrast, if the number of items is only slightly larger than the number of players, an MMS allocation is guaranteed to exist [Bouveret

and Lemaître 2014]. What is the largest number of items for which an MMS allocation is guaranteed to exist?

2. NONEXISTENCE OF EXACT MMS ALLOCATIONS

In this section we will show that, in general, MMS allocations are not guaranteed to exist (even under our assumption of additive valuation functions). But, to give some context for this result, let us briefly discuss a case where they *do* exist. As pointed out by Bouveret and Lemaître [2014], when there are two players we can achieve an MMS allocation — essentially via an indivisible analog of the Cut and Choose protocol. Indeed, let player 1 divide the items according to his 2-maximin partition S_1, S_2 , i.e., the partition that maximizes $\min_{j \in [2]} v_1(S_j)$. Allocate to player 2 his preferred subset, and give the other subset to player 1. Player 1 clearly achieves his MMS guarantee, but what about player 2? By the additivity of v_2 , there exists $j \in [2]$ such that $v_2(S_j) \geq v_2(\mathcal{M})/2$. In addition, in any partition S'_1, S'_2 there exists $k \in [2]$ such that $v_2(S'_k) \leq v_2(\mathcal{M})/2$, hence $\text{MMS}_2^{(2)}(\mathcal{M}) \leq v_2(\mathcal{M})/2$. It follows that there exists $j \in [2]$ such that $v_2(S_j) \geq \text{MMS}_2^{(2)}(\mathcal{M})$.

In contrast, MMS allocations may not exist when the number of players is at least three.

THEOREM 2.1. *For any set of players \mathcal{N} such that $|\mathcal{N}| \geq 3$ there exist \mathcal{M} and (additive) valuation functions that do not admit an MMS allocation.*

In a nutshell, to prove the theorem, the players' valuation functions are defined using three matrices. One of these matrices is called the *oscillation* matrix, and its elements are chosen to satisfy a system of linear equalities and inequalities that guarantees that only very specific subsets of the elements sum up to 1. We then show that valid MMS allocations must include bundles of items whose values in the oscillation matrix sum up to 1, but (by perturbing players' values using an *epsilon matrix*) this precludes allocations that satisfy all players' MMS guarantees.

We prove Theorem 2.1 in §2.1, but, as the proof is rather intricate, we first illustrate the main ideas of our general construction by presenting an explicit counterexample construction for the case of three players. Let the set of items be $\mathcal{M} = \{(j, k) \mid j \in [3], k \in [4]\}$. The valuation functions of the three players are defined using the following five matrices: the *base matrix*

$$B = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix},$$

the *oscillation Matrix*

$$O = \begin{bmatrix} 17 & 25 & 12 & 1 \\ 2 & 22 & 3 & 28 \\ 11 & 0 & 21 & 23 \end{bmatrix},$$

and three *epsilon matrices*:

$$E^{(1)} = \begin{bmatrix} 3 & -1 & -1 & -1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad E^{(2)} = \begin{bmatrix} 3 & -1 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{bmatrix} \quad E^{(3)} = \begin{bmatrix} 3 & 0 & -1 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$$

For each item $(j, k) \in \mathcal{M}$, we let

$$v_i(\{(j, k)\}) = 10^6 \cdot B_{jk} + 10^3 \cdot O_{jk} + E_{jk}^{(i)}.$$

Our first goal is to compute the 3-MMS guarantee of each player. To this end, we will find it convenient to label each element of the oscillation matrix O with three of nine

possible labels (1, 2, 3, α , β , γ , +, -, *):

$$\begin{bmatrix} \alpha 17_+^1 & \alpha 25_-^1 & \beta 12_+^1 & \gamma 1_*^1 \\ \alpha 2_-^2 & \beta 22_*^2 & \gamma 3_+^2 & \gamma 28_-^2 \\ \alpha 11_*^3 & \beta 0_-^3 & \beta 21_*^3 & \gamma 23_+^3 \end{bmatrix}$$

The oscillation matrix has the following Sudoku-like property: For each label there are exactly four elements with that label, and the sum of these 4 elements is exactly 55. Moreover, any four elements whose sum is 55 must have the same label.

This observation facilitates a straightforward computation of MMS guarantees. Player 1 can divide the 12 items into three subsets: a subset consisting of the four elements labeled with 1 (the first row), a subset consisting of the four elements labeled by 2 (the second row), and a subset consisting of the four elements labeled by 3 (the third row). For each subset, the sum of its four elements in B , O and E_1 is 4, 55 and 0 respectively. Hence, $\text{MMS}_1^{(3)}(\mathcal{M}) = 4 \cdot 10^6 + 55 \cdot 10^3 + 0 = 4055000$. Player 2's maximin partition is obtained by dividing the items into three subsets according to the labels α , β and γ , and player 3's maximin partition corresponds to the labels +, - and *; all MMS guarantees are 4055000.

We next characterize MMS allocations of \mathcal{M} , with the goal of showing that no such allocations exist. First note that a valid MMS allocation of \mathcal{M} must allocate at least four items to each player. Indeed, for any bundle $S \subseteq \mathcal{M}$ such that $|S| = 3$ and each player $i = 1, 2, 3$, $v_i(S) \leq 3 \cdot 10^6 + 76 \cdot 10^3 + 3 < 4055000$. Because there are twelve items, each player must receive exactly four items.

We now claim that in an MMS allocation each player must receive four items with the same label. Indeed, as noted above, the only bundles whose values in O add up to 55 consist of four items with identical labels. Suppose that a player is allocated four items with different labels. Since the sum of all the elements in O is $165 = 55 \times 3$, there must be a player with four items whose sum in O is less than 55. This player's value is at most $4 \cdot 10^6 + 54 \cdot 10^3 + 3 < 4055000$.

It is easy to verify that there are only three ways to divide \mathcal{M} into three subsets such that the items in each subset have identical labels: (i) dividing according to the labels 1, 2, 3; (ii) according to the labels α, β, γ ; and (iii) according to the labels +, - and *. But all three ways will fail to give some player his MMS guarantee of 4055000. Indeed, in case (i), there is a player $i_1 \in \{2, 3\}$ who is allocated items labeled by 2 or 3. The sum of the corresponding elements in $E^{(i_1)}$ is -1 , hence the value i_1 obtains is $4 \cdot 10^6 + 55 \cdot 10^3 - 1 = 4054999 < 4055000$. In case (ii), a player $i_2 \in \{1, 3\}$ must be allocated a subset of items labeled with β or γ ; and in case (iii), a player $i_3 \in \{1, 2\}$ must be allocated a subset of items labeled with - or *. By the same reasoning as in case (i), in cases (ii) and (iii) player i_l , $l = 2, 3$, ends up with value at most 4054999. We conclude that it is impossible to satisfy the MMS guarantees of all three players.

2.1. Proof of Theorem 2.1

Let $n = |\mathcal{N}|$. We construct a counterexample by defining the set of items $\mathcal{M} = [n]^n$. Each item is associated with a vector

$$\mathbf{t} = (t_1, t_2, \dots, t_n) \in \mathcal{M},$$

that is, $t_i \in [n]$ for all $i \in [n]$. Let

$$S_{i,j} = \{\mathbf{t} \in \mathcal{M} \mid t_i = j\}$$

We also let \mathcal{X} be the collection of all such subsets, i.e.,

$$\mathcal{X} = \{S_{i,j} \mid i, j \in [n]\},$$

and use this to define

$$\mathcal{Y} = \{T \subseteq \mathcal{M} \mid |T| = n^{n-1}, T \notin \mathcal{X}\}.$$

Now consider the following system of linear equalities and linear inequalities,

$$\begin{aligned} \forall T \in \mathcal{X}, \sum_{\mathbf{t} \in T} x_{\mathbf{t}} &= 1 \\ \forall T \in \mathcal{Y}, \sum_{\mathbf{t} \in T} x_{\mathbf{t}} &\neq 1 \end{aligned} \tag{1}$$

We claim that the system has at least one solution \mathbf{x} ; this claim is proved later on. Currently we use such a solution \mathbf{x} to construct the counterexample.

For two sets $V, U \subseteq \mathcal{M}$, where $|V| = d$, let $\langle U | V \rangle = \sum_{u \in U \cap V} \mathbf{e}_u^V$ be a d -dimensional vector, where $\mathbf{e}_u^V = (0, \dots, 1, \dots, 0) \in \mathbb{R}^d$ is the i th unit vector, in which i is the position of element u in V according to a fixed order. Specifically, we sort V 's elements in increasing lexicographic order, i.e., \mathbf{t} appears before \mathbf{t}' if in the leftmost position i where the two vectors differ, $t_i < t'_i$. This ensures that $\langle U | V \rangle$ is well defined for any U .

Our construction is based on the following vectors:

- *Base vector*: $\mathbf{b} = (1, 1, \dots, 1)^\top \in \mathbb{R}^{n^n}$.
- *Oscillation vector*: \mathbf{x} .
- *Epsilon vector for player i* :

$$\mathbf{p}_i = (1, 0, 0, \dots, 0)^\top - n^{1-n} \langle S_{i,1} | \mathcal{M} \rangle,$$

where the left vector on the right hand side has 1 in the position corresponding to item $(1, \dots, 1)$.

We also let

$$\lambda = \min_{T \in \mathcal{Y}} |\langle T | \mathcal{M} \rangle^\top \mathbf{x} - 1| > 0.$$

Using these notations, we set the utility function of player $i \in [n]$ by defining it as a vector, with an entry for each of the n^n items:

$$\mathbf{u}_i = n\mathbf{b} + \mathbf{x} + \frac{1}{2}n^{1-n}\lambda\mathbf{p}_i.$$

Correctness of the Construction. We first claim that when the utility functions are $\mathbf{u}_1, \dots, \mathbf{u}_n$, the n^n items cannot be allocated so that each player achieves his MMS guarantee. We start by giving a lower bound on each player's n -MMS guarantee. Suppose that player i divides the items into n subsets $S_{i,1}, \dots, S_{i,n}$. Then for all $j \in [n]$,

$$v_i(S_{i,j}) = \sum_{\mathbf{t} \in S_{i,j}} v_i(\mathbf{t}) = \sum_{\mathbf{t} \in S_{i,j}} (n + x_{\mathbf{t}}) + [j = 1] \cdot \frac{1}{2}n^{1-n}\lambda \left(1 - \sum_{\mathbf{t} \in S_{i,j}} n^{1-n} \right) = n^n + 1,$$

where the second equality follows from the fact that \mathbf{p}_i is only nonzero in positions corresponding to $S_{i,1}$. We conclude that the n -MMS guarantee of each player is at least $n^n + 1$.

LEMMA 2.2. *Any n -MMS allocation must allocate exactly n^{n-1} items to each player.*

PROOF. Suppose there is a player that is allocated at most $n^{n-1} - 1$ items, then his utility is less than $(n^{n-1} - 1)n + \sum_{\mathbf{t} \in \mathcal{M}} x_{\mathbf{t}} + 1 = n^n + 1$, which violates his n -MMS guarantee. The lemma now follows from the fact there is a total of n^n items. \square

LEMMA 2.3. *Any n -MMS allocation must allocate a set in \mathcal{X} to each player.*

PROOF. By Lemma 2.2, each player is allocated exactly n^{n-1} items. Suppose for contradiction that one player acquires n^{n-1} elements T such that $T \notin \mathcal{X}$. The constraints of the linear system (1) for which \mathbf{x} is a solution imply that $\sum_{t \in T} x_t \neq 1$. Since $\sum_{t \in \mathcal{M}} x_t = n$, there must be a player i who receives a set T' , $|T'| = n^{n-1}$, such that $\sum_{t \in T'} x_t < 1$. By the definition of λ , the value player i obtains from the \mathbf{x} component of his utility function is $\sum_{t \in T'} x_t \leq 1 - \lambda$. Therefore, the allocation of player i is worth at most

$$nn^{n-1} + 1 - \lambda + \frac{1}{2}n^{1-n}\lambda \cdot 1 \cdot n^{n-1} = n^n + 1 - \frac{1}{2}\lambda < n^n + 1,$$

contradicting the lower bound on the n -MMS guarantee. \square

LEMMA 2.4. *For any n -MMS allocation there exists $i \in \mathcal{N}$ such that the allocation forms the partition $S_{i,1}, S_{i,2}, \dots, S_{i,n}$.*

PROOF. By Lemma 2.3, players can only obtain a set of items that forms an element of \mathcal{X}^n . Now consider the case that player 1 obtains a set of items represented by $S_{i,j}$; no other player receives the bundle $S_{i',j'}$ where $i \neq i'$, since $S_{i,j} \cap S_{i',j'} \neq \emptyset$ when $i \neq i'$. At the same time, $S_{i,1}, S_{i,2}, \dots, S_{i,n}$ forms a legal n -partition. \square

LEMMA 2.5. *There is no n -MMS allocation of \mathcal{M} .*

PROOF. By Lemma 2.4, a valid n -MMS allocation must be based on the n -partition $S_{i,1}, S_{i,2}, \dots, S_{i,n}$. Suppose player j obtains the item bundle $S_{i,1}$. Since there are at least three players, there is a player k that $k \neq i$ and $k \neq j$. For any bundle $S_{i,l}$ where $l \neq 1$, the value player k obtains when allocated the bundle $S_{i,l}$ is smaller than $nn^{n-1} + 1$, because there exist elements corresponding to $S_{i,l}$ in \mathbf{p}_i that are negative, and the other elements are zero. It follows that player k does not achieve his n -MMS guarantee. \square

Proof of Solution Existence. In the rest of the proof we show that the linear system (1) indeed has a solution. We first introduce some notation. Consider a system of linear equations $A\mathbf{x} = \mathbf{b}$ where $A = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m]^\top$, $\mathbf{b} \in \mathbb{R}^m$, and $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^n$. Let

$$\mathbb{U} = \{\mathbf{u} \in \mathbb{R}^n \mid \forall \mathbf{y} \in \mathbb{R}^m, A^\top \mathbf{y} \neq \mathbf{u}\},$$

and let $A^{-1}(\mathbf{b}) = \{\mathbf{x} \in \mathbb{R}^n \mid A\mathbf{x} = \mathbf{b}\}$.

LEMMA 2.6. *If $A^{-1}(\mathbf{b}) \neq \emptyset$ then for all $\mathbf{u} \in \mathbb{U}^k$ and $\mathbf{c} \in \mathbb{R}^k$ there exists $\mathbf{x} \in A^{-1}(\mathbf{b})$ such that for all $i \in [k]$, $\mathbf{u}_i^\top \mathbf{x} \neq c_i$.*

PROOF. Suppose for contradiction that $A^{-1}(\mathbf{b}) \neq \emptyset$ but there exist $\mathbf{u} \in \mathbb{U}^k$ and $\mathbf{c} \in \mathbb{R}^k$ such that for all $\mathbf{x} \in A^{-1}(\mathbf{b})$ there exists $i \in [k]$ such that $\mathbf{u}_i^\top \mathbf{x} = c_i$.

Let r be the dimension of the kernel of A . We can find an orthogonal basis $\mathbf{x}_1, \dots, \mathbf{x}_r, \mathbf{x}_{r+1}, \dots, \mathbf{x}_n$ for \mathbb{R}^n such that for each $i \in [r]$, $A\mathbf{x}_i = 0$. Then,

$$\text{span}(\mathbf{x}_{r+1}, \dots, \mathbf{x}_n) = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_m),$$

where $\text{span}(S)$ is the *span* of a set of vectors S , i.e.,

$$\text{span}(S) = \left\{ \sum_i \lambda_i v_i \mid \lambda_i \in \mathbb{R}, v_i \in S \right\}.$$

Using our assumption that $A^{-1}(\mathbf{b}) \neq \emptyset$, let \mathbf{x}_0 such that $A\mathbf{x}_0 = \mathbf{b}$. We can express the solution of this system using \mathbf{x}_0 and a linear combination of $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_r$, i.e., each vector $\mathbf{x} = \mathbf{x}_0 + [\mathbf{x}_1, \dots, \mathbf{x}_r] \mathbf{y}$, where $\mathbf{y} \in \mathbb{R}^r$, satisfies $A\mathbf{x} = \mathbf{b}$.

Next, we define a collection of r -dimensional vectors by $\mathbf{y}_i = (i, i, \dots, i)^\top$ for $i = 1, \dots, k+1$. By the assumption,

$$\forall \mathbf{y} \in \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{k+1}\}, \exists i \in [k], \mathbf{u}_i^\top (\mathbf{x}_0 + [\mathbf{x}_1, \dots, \mathbf{x}_r] \mathbf{y}) = c_i.$$

By the pigeonhole principle, there exist $i \in [k]$ and $p > q \in [k+1]$ such that

$$\mathbf{u}_i^\top (\mathbf{x}_0 + [\mathbf{x}_1, \dots, \mathbf{x}_r] \mathbf{y}_p) = c_i$$

and

$$\mathbf{u}_i^\top (\mathbf{x}_0 + [\mathbf{x}_1, \dots, \mathbf{x}_r] \mathbf{y}_q) = c_i.$$

It follows that

$$\mathbf{u}_i^\top [\mathbf{x}_1, \dots, \mathbf{x}_r] (\mathbf{y}_p - \mathbf{y}_q) = 0.$$

Since $\mathbf{y}_p - \mathbf{y}_q$ is strictly positive in every coordinate, it holds that for all $j \in [r]$, $\mathbf{u}_i^\top \mathbf{x}_j = 0$, that is, $\mathbf{u}_i \perp \text{span}(\mathbf{x}_1, \dots, \mathbf{x}_r)$. We conclude that

$$\mathbf{u}_i \in \text{span}(\mathbf{x}_{r+1}, \dots, \mathbf{x}_n) = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_m),$$

and hence there exists a vector $\mathbf{y} \in \mathbb{R}^m$ such that $A^\top \mathbf{y} = \mathbf{u}_i$, so $\mathbf{u}_i \notin \mathbb{U}$ — contradicting the assumption that $\mathbf{u} \in \mathbb{U}^k$. \square

LEMMA 2.7. For all $T \in \mathcal{Y}$,

$$\langle T | \mathcal{M} \rangle \notin \text{span}(\{\langle T' | \mathcal{M} \rangle | T' \in \mathcal{X}\}).$$

PROOF. The lemma is equivalent to claiming that for $m = n$,

$$\forall T \in \mathcal{Y}, \langle T | \mathcal{M} \rangle \notin \text{span}(\{\langle S_{i,j} | \mathcal{M} \rangle | i \in [m], j \in [n]\}).$$

We prove this claim by induction on m . The case of $m = 1$ is trivial.

Assume the claim holds for $m-1$; we prove it for m . Assume for contradiction that there exists $T \in \mathcal{Y}$ such that

$$\langle T | \mathcal{M} \rangle \in \text{span}(\{\langle S_{i,j} | \mathcal{M} \rangle | i \in [m], j \in [n]\}).$$

Then we can find numbers $\lambda_{i,j}$ such that

$$\langle T | \mathcal{M} \rangle = \sum_{i \in [m-1], j \in [n]} \lambda_{i,j} \langle S_{i,j} | \mathcal{M} \rangle + \sum_{j \in [n]} \lambda_{m,j} \langle S_{m,j} | \mathcal{M} \rangle.$$

Next, define the vectors β_1, \dots, β_n as

$$\beta_l = \sum_{i \in [m-1], j \in [n]} \lambda_{i,j} \langle S_{i,j} | S_{m,l} \rangle$$

for all $l \in [n]$. Note that the summation is only over $i \in [m-1]$, and for $i \in [m-1]$ and $j, l_1, l_2 \in [n]$, $\langle S_{i,j} | S_{m,l_1} \rangle = \langle S_{i,j} | S_{m,l_2} \rangle$. It follows that for all $l \in [n]$, $\beta_l = \beta$. Furthermore,

$$\langle S_{m,j} | S_{m,l} \rangle = \begin{cases} (1, \dots, 1) & j = l \\ (0, \dots, 0) & j \neq l. \end{cases}$$

Therefore, letting $\gamma = (1, \dots, 1)$ and $\mu_l = \lambda_{m,l}$, we can write

$$\omega_l = \langle T | S_{m,l} \rangle = \beta + \mu_l \gamma.$$

We consider two cases:

Case 1: There exist $l_1 \neq l_2 \in [n]$ such that $\mu_{l_1} < \mu_{l_2}$. Since ω_{l_1} and ω_{l_2} are 0-1 vectors, and $\gamma = (1, \dots, 1)$, this can only happen when $\omega_{l_1} = (0, \dots, 0)$ and $\omega_{l_2} = (1, \dots, 1) = \gamma$. In addition, $T \in \mathcal{Y}$, therefore $|T| = n^{n-1}$. It follows that $T = S_{m, l_2} \in \mathcal{X}$ — a contradiction to the assumption that $T \in \mathcal{Y}$.

Case 2: For all $l_1, l_2 \in [n]$, $\mu_{l_1} = \mu_{l_2} = \mu$. In this case, it holds that

$$\begin{aligned} \langle T | \mathcal{M} \rangle &= \sum_{i \in [m-1], j \in [n]} \lambda_{i,j} \langle S_{i,j} | \mathcal{M} \rangle + \mu \sum_{j \in [n]} \langle S_{m,j} | \mathcal{M} \rangle \\ &= \sum_{i \in [m-1], j \in [n]} \lambda_{i,j} \langle S_{i,j} | \mathcal{M} \rangle + \mu \sum_{j \in [n]} \langle S_{1,j} | \mathcal{M} \rangle \\ &= \sum_{i \in [m-1], j \in [n]} \lambda'_{i,j} \langle S_{i,j} | \mathcal{M} \rangle, \end{aligned}$$

where the second equality follows from the fact that

$$\sum_{j \in [n]} \langle S_{m,j} | \mathcal{M} \rangle = \langle \mathcal{M} | \mathcal{M} \rangle = \sum_{j \in [n]} \langle S_{1,j} | \mathcal{M} \rangle,$$

and for the last equality we define $\lambda'_{i,j} = \lambda_{i,j} + \mu [i = 1]$. Since the summation on the right hand side does not include $i = m$, it follows that

$$\langle T | \mathcal{M} \rangle \in \text{span}(\{\langle S_{i,j} | \mathcal{M} \rangle \mid i \in [m-1], j \in [n]\}).$$

By the induction assumption, $T \notin \mathcal{Y}$ — a contradiction to the assumption that $T \in \mathcal{Y}$. \square

To complete the theorem's proof, recall that there are n^2 different sets in \mathcal{X} , and denote

$$\mathcal{X} = \{T^{(1)}, T^{(2)}, \dots, T^{(n^2)}\}.$$

We introduce the matrix

$$A = \left[\langle T^{(1)} | \mathcal{M} \rangle, \langle T^{(2)} | \mathcal{M} \rangle, \dots, \langle T^{(n^2)} | \mathcal{M} \rangle \right]^\top.$$

As in the beginning of this subsection, we denote

$$\begin{aligned} \mathbb{U} &= \{\mathbf{u} \in \mathbb{R}^n \mid \forall \mathbf{y} \in \mathbb{R}^m, A^\top \mathbf{y} \neq \mathbf{u}\} \\ &= \left\{ \mathbf{u} \in \mathbb{R}^n \mid \mathbf{u} \notin \text{span} \left(\langle T^{(1)} | \mathcal{M} \rangle, \langle T^{(2)} | \mathcal{M} \rangle, \dots, \langle T^{(n^2)} | \mathcal{M} \rangle \right) \right\} \\ &= \{\mathbf{u} \in \mathbb{R}^n \mid \mathbf{u} \notin \text{span}(\langle T' | \mathcal{M} \rangle \mid T' \in \mathcal{X})\} \supseteq \{\langle T' | \mathcal{M} \rangle \mid T' \in \mathcal{Y}\}, \end{aligned}$$

where the containment follows from Lemma 2.7. Let $\mathbf{b} = (1, 1, \dots, 1)$; it is easy to see that $A^{-1}(\mathbf{b}) \neq \emptyset$, because we can set $x_{\mathbf{t}} = n^{1-n}$ for all $\mathbf{t} \in \mathcal{M}$. By Lemma 2.6, it follows that there is a solution \mathbf{x} such that $A\mathbf{x} = \mathbf{b}$ and for all $T \in \mathcal{Y}$, $\langle T | \mathcal{M} \rangle^\top \mathbf{x} \neq 1$. This completes the proof of Theorem 2.1. \square

3. EXISTENCE AND COMPUTATION OF APPROXIMATE MMS ALLOCATIONS

To circumvent Theorem 2.1 we introduce a new notion of *approximate* maximin share guarantee: rather than asking for an allocation $A_1, \dots, A_{|\mathcal{N}|}$ such that $v_i(A_i) \geq \text{MMS}_i^{(|\mathcal{N}|)}(\mathcal{M})$ for all $i \in \mathcal{N}$, we look for allocations such that $v_i(A_i) \geq c \cdot \text{MMS}_i^{(|\mathcal{N}|)}(\mathcal{M})$ for a constant $c > 0$. Our main result is that such allocations always exist when $c = 2/3$.

THEOREM 3.1. *Let there be $|\mathcal{N}| \geq 1$ players and a set of items \mathcal{M} . Then there exists an allocation $A_1, \dots, A_{|\mathcal{N}|}$ such that for all $i \in \mathcal{N}$, $v_i(A_i) \geq \frac{2}{3} \text{MMS}_i^{(|\mathcal{N}|)}(\mathcal{M})$. Moreover, such an allocation can be found in polynomial time if $|\mathcal{N}|$ is constant.*

At the center of our technical approach lies the *density balance parameter*. Intuitively, the smaller a subset's value, the higher the value of its complement — in this sense we are interested in the *balance* between the two. Formally, the N -density balance parameter ρ_N is given by

$$\rho_N = \max \left\{ \lambda \left| \begin{array}{l} \forall \mathcal{M}, \forall \text{additive } v_i \in (\mathbb{R}^+)^{2^{\mathcal{M}}}, \forall S \subseteq \mathcal{M}, \forall n, m \text{ s.t. } n + m = N, \\ v_i(\mathcal{M} \setminus S) \leq m \lambda \text{MMS}_i^{(N)}(\mathcal{M}) \Rightarrow \text{MMS}_i^{(n)}(S) \geq \lambda \text{MMS}_i^{(N)}(\mathcal{M}) \end{array} \right. \right\}$$

The crucial point is that, if there are $|\mathcal{N}| = N$ players and $v(\mathcal{M} \setminus S) \leq m \lambda \text{MMS}_i^{(N)}(\mathcal{M})$, giving player i any subset in his n -maximin partition of S (that is, his maximin partition if there were only n players and the available items were S) would be sufficient to guarantee a ρ_N fraction of his N -player MMS guarantee.

Our most significant technical tool is the following lemma, which uses $\lfloor N \rfloor_{\text{odd}} = N - 1 + (N \bmod 2)$ to denote the largest odd N' such that $N' \leq N$.

LEMMA 3.2 (DENSITY BALANCE LEMMA). *For all $N \geq 2$,*

$$\rho_N = \frac{2 \lfloor N \rfloor_{\text{odd}}}{3 \lfloor N \rfloor_{\text{odd}} - 1} \geq \frac{2}{3}.$$

The lemma's intricate proof is given in §3.3. But if we believe the lemma for now, to prove Theorem 3.1 it is sufficient to guarantee each player $\rho_{|\mathcal{N}|} \geq 2/3$ of his $|\mathcal{N}|$ -MMS guarantee. Importantly, this gives a stronger guarantee for a small number of players, and in particular for the cases of three and four players we get a $(3/4)$ -approximation.

We prove the theorem in two steps. First, we disregard time complexity issues and establish that a $\rho_{|\mathcal{N}|}$ -approximation is always feasible by constructing an exponential-time algorithm. Then, we explain how to convert the algorithm into a polynomial-time algorithm (assuming the number of players is constant).

3.1. An Exponential-Time Algorithm

Our goal is to construct an algorithm APX-MMS (n, N, S, \mathcal{M}) , where $N = |\mathcal{N}|$, by induction on n . The algorithm allocates the items $S \subseteq \mathcal{M}$ to n players denoted (for ease of exposition) by $[n]$, and guarantees each player a value of at least $\rho_N \text{MMS}_i^{(N)}(\mathcal{M})$, if the following assumption holds:

$$\forall i \in [n], v_i(\mathcal{M} \setminus S) \leq (N - n) \rho_N \text{MMS}_i^{(N)}(\mathcal{M}). \quad (2)$$

A subtle point is that while we build the algorithm from the bottom up, we also have to make sure that we satisfy Equation (2) when we recurse on a *smaller* subset of players — and we must make sure that this inequality holds when we first call the algorithm (we will do so later).

For $n = 1$ — the base of the induction — APX-MMS $(1, N, S, \mathcal{M})$ allocates S to the single player, denoted by 1. Note that for all $i \in \mathcal{N}$, $v_i(\mathcal{M}) \geq N \cdot \text{MMS}_i^{(N)}(\mathcal{M})$, and since we are assuming that

$$v_1(\mathcal{M} \setminus S) \leq (N - 1) \rho_N \text{MMS}_1^{(N)}(\mathcal{M}) \leq (N - 1) \text{MMS}_1^{(N)}(\mathcal{M}),$$

it follows that $v_1(S) \geq \text{MMS}_1^{(N)}(\mathcal{M}) \geq \rho_N \text{MMS}_1^{(N)}(\mathcal{M})$.

Now, suppose that for all $n \leq n'$ we have an algorithm APX-MMS (n, N, S, \mathcal{M}) that guarantees each of the n players a value of at least $\rho_N \text{MMS}_i^{(N)}(\mathcal{M})$ if

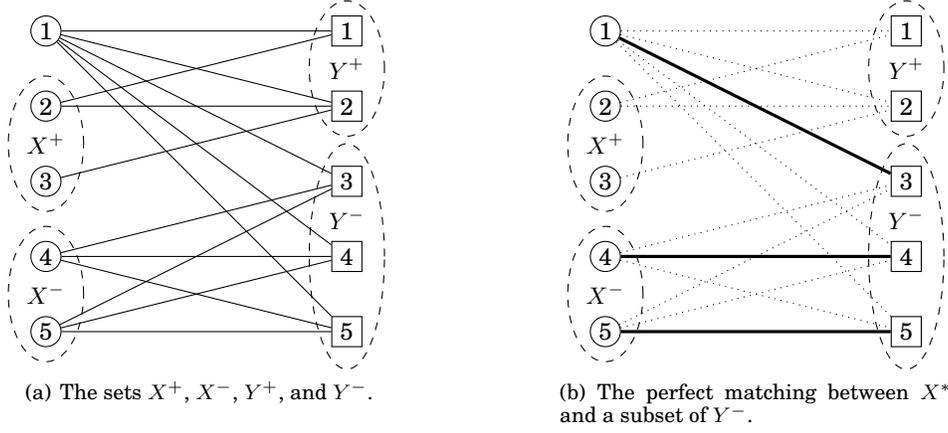


Fig. 1. An illustration of the proof of Theorem 3.1, with $n = 5$. Vertices in X (corresponding to players) are circles and vertices in Y (corresponding to subsets) are squares. Note that $1 \in X$ has an edge to every $j \in Y$.

Equation (2) holds. Our task in the following paragraphs is the construction of $\text{APX-MMS}(n, N, S, \mathcal{M})$ where $n = n' + 1$. $\text{APX-MMS}(n, N, S, \mathcal{M})$ starts out by computing an n -maximin partition of S with respect to player 1, denoted S_1, S_2, \dots, S_n . Using (2) we have that $v_1(\mathcal{M} \setminus S) \leq (N - n)\rho_N \text{MMS}_1^{(N)}(\mathcal{M})$. By the definition of ρ_N , $\text{MMS}_1^{(n)}(S) \geq \rho_N \text{MMS}_1^{(N)}(\mathcal{M})$. This means that for all $j \in [n]$, $v_1(S_j) \geq \rho_N \text{MMS}_1^{(N)}(\mathcal{M})$.

The next step is to draw an undirected bipartite graph $G = (V, E)$ where V is composed of two subsets of vertices $X = Y = [n]$, and

$$E = \left\{ (i, j) \mid i \in X, j \in Y, v_i(S_j) \geq \rho_N \text{MMS}_i^{(N)}(\mathcal{M}) \right\}.$$

Furthermore, for any subset $\Gamma \subseteq Y$, define the function f_Γ that maps subsets $Z \subseteq X$ to their neighbors in Y :

$$f_\Gamma(Z) = \bigcup_{i \in Z} \{j \in \Gamma \mid (i, j) \in E\}.$$

Finally, we denote

$$X^+ = \arg \max_{Z \subseteq \{2, 3, \dots, n\}} \{|Z| \mid |Z| \geq |f_Y(Z)|\},$$

$X^* = X \setminus X^+$, $X^- = X^* \setminus \{1\}$, $Y^+ = f_Y(X^+)$, and $Y^- = Y \setminus Y^+$ (see Figure 1(a) for an illustration).

LEMMA 3.3. *There exists a perfect matching between X^* and a subset of Y^- .*

PROOF. We prove the lemma using Hall's Theorem. We need to show that

$$\forall Z \subseteq X^*, |Z| \leq |f_{Y^-}(Z)|. \quad (3)$$

We first prove a restricted version of Equation (3) for subsets $Z \subseteq X^-$, that is,

$$\forall Z \subseteq X^-, |Z| \leq |f_{Y^-}(Z)|. \quad (4)$$

To see this, assume for contradiction that there exists $Z_0 \subseteq X^-$ such that $|Z_0| > |f_{Y^-}(Z_0)|$; then

$$\begin{aligned} |X^+ \cup Z_0| &= |X^+| + |Z_0| > |f_Y(X^+)| + |f_{Y^-}(Z_0)| \geq |f_Y(X^+) \cup f_{Y^-}(Z_0)| \\ &= |f_Y(X^+ \cup Z_0)|. \end{aligned}$$

This contradicts the definition of X^+ as the largest subset of $\{2, \dots, n\}$ whose cardinality is larger than that of its neighbors in Y , and establishes Equation (4).

We next expand Equation (4) to all $Z \subseteq X^*$. Since $X^* = X^- \cup \{1\}$, all subsets of X^* can be represented as Z or $Z \cup \{1\}$ for $Z \subseteq X^-$. For any subset $Z \subseteq X^-$, we have that $|Z| \leq |f_{Y^-}(Z)|$ by Equation (4). We have already argued that $v_1(S_j) \geq \rho_N \mathbf{MMS}_1^{(N)}(\mathcal{M})$ for all $j \in [n]$, meaning that $(1, j) \in E$ for all $j \in [n]$. Therefore

$$|Z \cup \{1\}| \leq |X^*| = n - |X^+| \leq n - |Y^+| = |Y^-| = |f_{Y^-}(\{1\})| \leq |f_{Y^-}(Z \cup \{1\})|,$$

establishing Equation (3) and completing the lemma's proof. \square

By Lemma 3.3 there exists a perfect matching between X^* and a subset of Y^- (see Figure 1(b) for an illustration). Mark the matched subset in Y^- as Y^* . If $i \in X^*$ is matched to $j \in Y^*$, APX-MMS (n, N, S, \mathcal{M}) will allocate subset S_j to player i .

To complete the allocation, let S^* be the subset of items that have not been allocated above, i.e., $S^* = \bigcup_{j \in Y \setminus Y^*} S_j$. Since $Y^* \subseteq Y^- = Y \setminus f_Y(X^+)$, for any $i \in X^+$ and $j \in Y^*$, $(i, j) \notin E$, that is, $v_i(S_j) < \rho_N \mathbf{MMS}_i^{(N)}(\mathcal{M})$. Therefore, for any player $i \in X^+$,

$$\begin{aligned} v_i(\mathcal{M} \setminus S^*) &= v_i\left(\left(\bigcup_{j \in Y^*} S_j\right) \cup (\mathcal{M} \setminus S)\right) = \sum_{j \in Y^*} v_i(S_j) + v_i(\mathcal{M} \setminus S) \\ &\leq |Y^*| \rho_N \mathbf{MMS}_i^{(N)}(\mathcal{M}) + (N - n) \rho_N \mathbf{MMS}_i^{(N)}(\mathcal{M}) \\ &= (|X^*| + N - n) \rho_N \mathbf{MMS}_i^{(N)}(\mathcal{M}) = (N - |X^+|) \rho_N \mathbf{MMS}_i^{(N)}(\mathcal{M}). \end{aligned}$$

We conclude that the conditions required to execute APX-MMS $(|X^+|, N, S^*, \mathcal{M})$ for players in X^+ (that is, the equivalent of Equation (2)) are satisfied, allowing us to divide the items in S^* between the players in X^+ in a way that each player receives value $\rho_N \mathbf{MMS}_i^{(N)}(\mathcal{M})$. Note that $X^+ \subseteq \{2, 3, \dots, n\}$, hence $|X^+| \leq n - 1$, i.e., APX-MMS $(|X^+|, N, S^*, \mathcal{M})$ provides the required guarantee by the induction assumption.

Initially we call APX-MMS $(|\mathcal{N}|, |\mathcal{N}|, \mathcal{M}, \mathcal{M})$. We have shown that (2) holds in subsequent calls, so it only remains to make sure that it also holds in this first call to the algorithm. This is clearly the case, because both sides of the inequality are zero.

3.2. A Polynomial-Time Algorithm

While the algorithm described above seems rather innocent at first glance, it does make one computational leap³ by letting one of the players compute an n -maximin partition of the current set of items S . It is easy to see that this is **NP-hard**; in fact, even when there are two players with identical valuations, it is **NP-hard** to determine whether the the MMS guarantee is $v_i(\mathcal{M})/2$ — this can be shown via an immediate reduction from PARTITION.

Woeginger [1997] studied the problem of computing a maximin partition, albeit under a different name: scheduling jobs on identical machines to maximize the minimum completion time. He gave a polynomial-time approximation scheme (PTAS), and

³The computation of the set X^+ may also be hard, but we are interested in the case of a constant number of players, for which it is obviously tractable.

showed that no fully polynomial-time approximation scheme (FPTAS) exists unless $P = NP$. Using our terminology, this means that given a constant $\epsilon > 0$ we can compute a partition S_1, \dots, S_n of the set of items S so that $\min_{i \in [n]} v_i(S) \geq (1 - \epsilon) \text{MMS}_i^{(n)}(S)$ in polynomial time.

The modified algorithm is almost identical, but the initial maximin partition, as well as edges in the bipartite graph, are computed based on $(1 - \epsilon)$ -approximations of MMS guarantees. The analysis essentially goes through unchanged, giving each player a bundle worth $(1 - \epsilon) \rho_N \text{MMS}_i^{(N)}(\mathcal{M})$. But, crucially, the exact value given by the Density Balance Lemma (Lemma 3.2) for ρ_N is slightly larger than $2/3$:

$$\rho_N = \frac{2 \lfloor N \rfloor_{\text{odd}}}{3 \lfloor N \rfloor_{\text{odd}} - 1} \geq \frac{2N}{3N - 1} = \frac{2}{3} \left(1 + \frac{1}{3N - 1} \right).$$

To get a $2/3$ -approximation, we set $(1 - \epsilon) = \left(1 + \frac{1}{3N - 1} \right)^{-1}$, that is, $\epsilon = \Theta(1/N)$ suffices. If N is constant, Woeginger's PTAS [Woeginger 1997] will run in polynomial time.

3.3. Proof of Lemma 3.2

For the purposes of this proof we can drop the subscript i , e.g., we can write $\text{MMS}^{(N)}(\mathcal{M})$ instead of $\text{MMS}_i^{(N)}(\mathcal{M})$ and v instead of v_i , because the properties of the density balance parameter ρ_N hold for any player and any possible valuation function.

Our first lemma proves a weaker result, implying an approximation ratio of $1/2$ instead of $2/3$. This lemma is technically required for the proof of the stronger bound, but it may also be of independent interest because the entire proof is quite long and intricate, whereas this weaker result is much easier to understand.

LEMMA 3.4 (WEAK DENSITY BALANCE LEMMA). *For all $N \geq 2$, $\rho_N \geq \frac{N}{2N-2} \geq \frac{1}{2}$.*

PROOF. Denote $\lambda_N = \frac{N}{2N-2}$. We will prove the lemma by establishing that $\rho_N \geq \lambda_N$, i.e., for any valuation function v , $S \subseteq \mathcal{M}$, and $n + m = N$,

$$v(\mathcal{M} \setminus S) \leq m \lambda_N \text{MMS}^{(N)}(\mathcal{M}) \Rightarrow \text{MMS}^{(n)}(S) \geq \lambda_N \text{MMS}^{(N)}(\mathcal{M}).$$

For any subset $S \subseteq \mathcal{M}$ and n, m such that $n + m = N$, suppose the condition $v(\mathcal{M} \setminus S) \leq m \lambda_N \text{MMS}^{(N)}(\mathcal{M})$ is satisfied. We only need to consider the case where $m \geq 1$ since the implication is obviously correct when $m = 0$.

We start by collecting all the items with value with at least $\lambda_N \text{MMS}^{(N)}(\mathcal{M})$ into a set T :

$$T = \{t_1, t_2, \dots, t_k\} = \left\{ t \in S \mid v(t) \geq \lambda_N \text{MMS}^{(N)}(\mathcal{M}) \right\}.$$

If $k \geq n$, we can create an n -partition such that each bundle has value at least $\lambda_N \text{MMS}^{(N)}(\mathcal{M})$ by placing the n items in T into n different subsets and then adding the other items to any of the subsets. Formally:

$$\text{MMS}^{(n)}(S) \geq \min \{v(t_1), v(t_2), \dots, v(t_k)\} \geq \lambda_N \text{MMS}^{(N)}(\mathcal{M}).$$

Otherwise, $|T| = k \leq n$. It holds that

$$\text{MMS}^{(N-k)}(\mathcal{M} \setminus T) \geq \text{MMS}^{(N)}(\mathcal{M}),$$

because we can always find k subsets of items to discard from an N -maximin partition such that all the items in T are discarded, and then the $(N - k)$ -MMS guarantee of the

remaining items (which would be a subset of $\mathcal{M} \setminus T$) is at least as large as the original N -MMS guarantee. Therefore,

$$\begin{aligned}
v(S \setminus T) &= v(\mathcal{M} \setminus T) - v(\mathcal{M} \setminus S) \geq (N - k) \text{MMS}^{(N-k)}(\mathcal{M} \setminus T) - m\lambda_N \text{MMS}^{(N)}(\mathcal{M}) \\
&\geq (N - k) \text{MMS}^{(N)}(\mathcal{M}) - m\lambda_N \text{MMS}^{(N)}(\mathcal{M}) = (N - k - m\lambda_N) \text{MMS}^{(N)}(\mathcal{M}) \\
&\geq (2(n - k) - 1) \lambda_N \text{MMS}^{(N)}(\mathcal{M}),
\end{aligned} \tag{5}$$

where the last inequality holds because

$$\begin{aligned}
(N - k - m\lambda_N) - (2(n - k) - 1) \lambda_N &= (2\lambda_N - 1)k + N + (1 - m - 2n) \lambda_N \\
&\geq N + (1 - m - 2n) \lambda_N = \frac{(m - 1)N}{2N - 2} \geq 0.
\end{aligned}$$

Since any item in $S \setminus T$ has value at most $\lambda_N \text{MMS}^{(N)}(\mathcal{M})$ and the total value of $S \setminus T$ is at least $(2(n - k) - 1) \lambda_N \text{MMS}^{(N)}(\mathcal{M})$, we can divide items in S into n subsets using the following algorithm. First, each item in T is a singleton subset. For each of the $n - k$ other subsets, we iteratively assign items in $S \setminus T$ to that subset until its value is at least $\lambda_N \text{MMS}^{(N)}(\mathcal{M})$; note that the value in each of these subsets will not exceed $2\lambda_N \text{MMS}^{(N)}(\mathcal{M})$, which by Equation (5) means that there is enough value in $S \setminus T$ to make sure each subset has value at least $\lambda_N \text{MMS}^{(N)}(\mathcal{M})$. We have proved that $\text{MMS}^{(n)}(S) \geq \lambda_N \text{MMS}^{(N)}(\mathcal{M})$, which completes the lemma's proof. \square

With the goal of obtaining a stronger bound in mind, we introduce the concept of *water*. An item is considered to be (made of) water if it is divisible, that is, it can be divided arbitrarily between multiple players without losing value. Equivalently, it may be intuitive to think of water as cash — although the water interpretation is more natural in our figures.

We show that an equivalent definition for the density balance parameter relies on the condition that for all $S \subseteq \mathcal{M}$, if $v(\mathcal{M} \setminus S) = (N - n) \text{MMS}^{(n)}(S)$ then $\text{MMS}^{(n)}(S) \geq \rho_N \text{MMS}^{(N)}(\mathcal{M})$. Moreover, this is true when $\mathcal{M} \setminus S$ is water. Therefore, we need to be able to reason about the MMS guarantee over indivisible items and water.

Suppose that the items in S have been partitioned into subsets. How should we allocate water W in order to maximize the minimum value of any subset? On an intuitive level we can regard each indivisible items in S as a brick of height equal to its value and width 1. Bricks that belong to the same subset are placed on top of each other in a container. Now (this is where the water analogy is useful) we pour the water W into the container; its volume is $v(W)$. The height of the water line is exactly the value we are interested in, and each subset now contains bricks and water in a single column of width 1 (see Figure 2 for an illustration). By enumerating all possible ways to place the items in S into the container we can obtain an N -maximin partition over $S \cup W$ — it would be the one with the highest water line.

We formalize this intuition and develop it into a series of lemmas, which culminates in the proof of the following lemma — a lower bound on ρ_N . The proof is relegated to the full version of the paper.⁴

LEMMA 3.5. *Let $\rho_N^* = \frac{2 \lfloor N \rfloor_{\text{odd}}}{3 \lfloor N \rfloor_{\text{odd}} - 1}$. Then $\rho_N \geq \rho_N^*$.*

⁴Available at: <http://www.cs.cmu.edu/~arielpro/papers.html>.

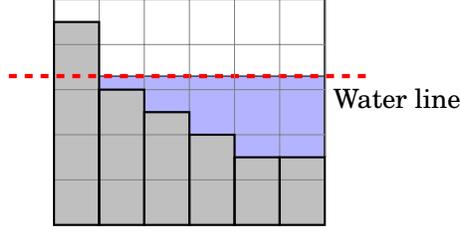


Fig. 2. Filling water to maximize the minimum value of any subset.

By Lemma 3.5, to show that $\rho_N = \rho_N^*$ and complete the proof of Lemma 3.2, it is sufficient to show that $\rho_N \leq \rho_N^*$. We establish this inequality by constructing an example for any N .

Let $\mathcal{M}^* = \{i_1, \dots, i_{\lceil N \rceil_{\text{odd}}}, W\}$, where $\lceil N \rceil_{\text{odd}} = N + 1 - (N \bmod 2)$ indicates the smallest odd number N' such that $N' \geq N$. For all $j \in [n]$, $v(i_j) = 1$, and W is water such that $v(W) = \lfloor \frac{N-1}{2} \rfloor$. We also let $S^* = \{i_1, \dots, i_{\lceil N \rceil_{\text{odd}}}\}$, so $\mathcal{M}^* \setminus S^* = \{W\}$. Finally, we let $n^* = \lceil \frac{N+1}{2} \rceil$ and $m^* = \lfloor \frac{N-1}{2} \rfloor$.

We claim that

$$\text{MMS}^{(N)}(\mathcal{M}^*) = 1 + \frac{\lfloor \frac{N-1}{2} \rfloor}{\lceil N \rceil_{\text{odd}}}.$$

This is obviously true for an odd N : each subset contains one item i_j and $1/N$ of the water W . For an even N , we can place the $N + 1$ indivisible items into N subsets such that one subset has two items (and value 2). We then divide the water between the remaining $N - 1 = \lceil N \rceil_{\text{odd}}$ subsets, so each has value equal to the N -MMS guarantee.

Since

$$\rho_N^* \text{MMS}^{(N)}(\mathcal{M}^*) = \frac{2\lceil N \rceil_{\text{odd}}}{3\lceil N \rceil_{\text{odd}} - 1} \cdot \left(1 + \frac{\lfloor \frac{N-1}{2} \rfloor}{\lceil N \rceil_{\text{odd}}}\right) = 1$$

and $v(W) = \lfloor \frac{N-1}{2} \rfloor = m^*$, it holds that $v(\mathcal{M}^* \setminus S^*) = v(W) \cdot 1 = m^* \cdot \rho_N^* \text{MMS}^{(N)}(\mathcal{M}^*)$. Moreover, there are $\lceil N \rceil_{\text{odd}}$ identical indivisible items in subset S^* , and

$$2n^* - 1 = 2 \left\lceil \frac{N+1}{2} \right\rceil - 1 = \lceil N \rceil_{\text{odd}}.$$

Thus in any n^* -maximin partition of S^* there is a subset with only one item. It follows that $\text{MMS}^{(n^*)}(S^*) = 1 = \rho_N^* \text{MMS}^{(N)}(\mathcal{M}^*)$. This shows that $\rho_N \leq \rho_N^*$, and completes the proof of Lemma 3.2. \square

4. ACKNOWLEDGMENTS

Procaccia was supported in part by the NSF under grant CCF-1215883; Wang was supported in part by the National Basic Research Program of China grant 2011CBA00300, 2011CBA00301, and the National Natural Science Foundation of China grant 61033001, 61361136003. We thank Eric Budish, Sylvain Bouveret, Steven Brams, Ioannis Caragiannis, Jonathan Goldman, Ian Kash, Jeremy Karp, Alex Kazachkov, Michel Lemaître, Omer Lev, Hervé Moulin, and Erel Segal Halevi for helpful discussions and comments.

REFERENCES

ABDULKADIROĞLU, A., SÖNMEZ, T., AND ÜNVER, M. U. 2004. Room assignment-rent division: A market approach. *Social Choice and Welfare* 22, 3, 515–538.

- ALON, N. 1987. Splitting necklaces. *Advances in Mathematics* 63, 241–253.
- ASADPOUR, A. AND SABERI, A. 2007. An approximation algorithm for max-min fair allocation of indivisible goods. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*. 114–121.
- BANSAL, N. AND SVIRIDENKO, M. 2006. The Santa Claus problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*. 31–40.
- BEZÁKOVÁ, I. AND DANI, V. 2005. Allocating indivisible goods. *SIGecom Exchanges* 5, 3, 11–18.
- BOUVERET, S. AND LEMAÎTRE, M. 2014. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. Forthcoming.
- BRAMS, S. J., EDELMAN, P. H., AND FISHBURN, P. C. 2003. Fair division of indivisible items. *Theory and Decision* 55, 2, 147–180.
- BRAMS, S. J., KILGOUR, M., AND KLAMLER, C. 2014. Two-person fair division of indivisible items: An efficient, envy-free algorithm. *Notices of the AMS* 61, 2, 130–141.
- BRAMS, S. J. AND TAYLOR, A. D. 1995. An envy-free cake division protocol. *The American Mathematical Monthly* 102, 1, 9–18.
- BRAMS, S. J. AND 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., KAKLAMANIS, C., KANELLOPOULOS, P., AND KYROPOULOU, M. 2009. The efficiency of fair division. In *Proceedings of the 5th International Workshop on Internet and Network Economics (WINE)*. 475–482.
- DE CLIPPEL, G., MOULIN, H., AND TIDEMAN, N. 2008. Impartial division of a dollar. *Journal of Economic Theory* 139, 176–191.
- DEMERS, A., KESHAV, S., AND SHENKER, S. 1989. Analysis and simulation of a fair queueing algorithm. In *Proceedings of the ACM Symposium on Communications Architectures & Protocols (SIGCOMM)*. 1–12.
- EDMONDS, J. AND PRUHS, K. 2006a. Balanced allocations of cake. In *Proceedings of the 47th Symposium on Foundations of Computer Science (FOCS)*. 623–634.
- EDMONDS, J. AND PRUHS, K. 2006b. Cake cutting really is not a piece of cake. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 271–278.
- GHOSSI, A., ZAHARIA, M., HINDMAN, B., KONWINSKI, A., SHENKER, S., AND STOICA, I. 2011. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)*. 24–37.
- HILL, T. 1987. Partitioning general probability measures. *Annals of Probability* 15, 2, 804–813.
- LIPTON, R. J., MARKAKIS, E., MOSSEL, E., AND SABERI, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Electronic Commerce (EC)*. 125–131.
- MAGDON-ISMAIL, M., BUSCH, C., AND KRISHNAMOORTHY, M. S. 2003. Cake cutting is not a piece of cake. In *Proceedings of the 20th International Symposium on Theoretical Aspects of Computer Science (STACS)*. 596–607.
- MARKAKIS, E. AND PSOMAS, C.-A. 2011. On worst-case allocations in the presence of indivisible goods. In *Proceedings of the 5th International Workshop on Internet and Network Economics (WINE)*. 278–289.
- MOULIN, H. 1990. Uniform externalities: Two axioms for fair allocation. *Journal of Public Economics* 43, 3, 305–326.
- MOULIN, H. 2003. *Fair Division and Collective Welfare*. MIT Press.
- PARKES, D. C., PROCACCIA, A. D., AND SHAH, N. 2012. Beyond Dominant Resource Fairness: Extensions, limitations, and indivisibilities. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. 808–825.
- PAZNER, E. AND SCHMEIDLER, D. 1978. Egalitarian equivalent allocations: A new concept of economic equity. *Quarterly Journal of Economics* 92, 4, 671–687.
- PROCACCIA, A. D. 2009. Thou shalt covet thy neighbor’s cake. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*. 239–244.
- PROCACCIA, A. D. 2013. Cake cutting: Not just child’s play. *Communications of the ACM* 56, 7, 78–87.
- WÖEGINGER, G. 1997. A polynomial-time approximation scheme for maximizing the minimum machine completion time. *Operations Research Letters* 20, 4, 149–154.
- WÖEGINGER, G. J. AND SGALL, J. 2007. On the complexity of cake cutting. *Discrete Optimization* 4, 213–220.