When Can the Maximin Share Guarantee Be Guaranteed?

David Kurokawa
Computer Science Department
Carnegie Mellon University

Ariel D. Procaccia
Computer Science Department
Carnegie Mellon University

Junxing Wang
Computer Science Department
Carnegie Mellon University

Abstract

The fairness notion of *maximin share (MMS) guarantee* underlies a deployed algorithm for allocating indivisible goods under additive valuations. Our goal is to understand when we can expect to be able to give each player his MMS guarantee. Previous work has shown that such an *MMS allocation* may not exist, but the counterexample requires a number of goods that is exponential in the number of players; we give a new construction that uses only a linear number of goods. On the positive side, we formalize the intuition that these counterexamples are very delicate by designing an algorithm that provably finds an MMS allocation with high probability when valuations are drawn at random.

1 Introduction

We study the classic problem of *fairly* allocating *indivisible* goods among players with *additive* valuation functions. Specifically, let the set of players be $N = \{1, \ldots, n\}$, and let the set of goods be G, with $|G| = m$. We denote the value of player $i \in N$ for good $g \in G$ by $V_i(g) \geq 0$. For a bundle of items $S \subseteq G$, we assume that $V_i(S) = \sum_{g \in S} V_i(g)$ (i.e. additive valuations). We are interested in finding an *allocation* $A_1, \ldots, A_n$ — this is a partition of G where A_i is the bundle of goods allocated to player $i \in N$.

Let us now revisit the first sentence above — what do we mean by "fairly"? Before presenting the fairness notion we are interested in, let us briefly discuss two others. An allocation is *envy free* if for all $i, j \in N$, $V_i(A_i) \geq V_i(A_j)$; and it is *proportional* if for all $i \in N$, $V_i(A_i) \geq V_i(G)/n$. Note that, in our setting, any envy-free allocation is also proportional. While these notions are compelling — and provably feasible in some fair division settings, such as cake cutting (Brams and Taylor 1996; Procaccia 2013) — they cannot always be achieved in our setting (say for example when there are two players and one good).

We therefore focus on a third fairness notion: *maximin share (MMS)* guarantee, introduced by Budish (2011). The MMS guarantee of player $i \in N$ is

$$\text{MMS}(i) = \max_{S_1, \ldots, S_n} \min_{j \in N} V_i(S_j),$$

where $S_1, \ldots, S_n$ is a partition of the set of goods G; a partition that maximizes this value is known as an *MMS partition*. In words, this is the value player i can achieve by dividing the goods into n bundles, and receiving his least desirable bundle. Alternatively, this is the value i can *guarantee* by partitioning the items, and then letting all other players choose a bundle before he does. An *MMS allocation* is an allocation $A_1, \ldots, A_n$ such that for all $i \in N$, $V_i(A_i) \geq \text{MMS}(i)$. In contrast to work on maximizing the minimum value of any player (Bansal and Sviridenko 2006; Asadpour and Saberi 2007; Roos and Rothe 2010), MMS is a "Boolean" fairness notion. Also note that a proportional allocation is always an MMS allocation, that is, proportionality is a stronger fairness property than MMS.

It is tempting to think that in our setting (additive valuations), an MMS allocation always exists. In fact, extensive experiments by Bouveret and Lemaître (2014) did not yield a single counterexample. Alas, it turns out that (intricate) counterexamples do exist (Procaccia and Wang 2014). On the positive side, *approximate* MMS allocations are known to exist. Specifically, it is always possible to give each player a bundle worth at least $2/3$ of his MMS guarantee, that is, there exists an allocation $A_1, \ldots, A_n$ such that for all $i \in N$, $V_i(A_i) \geq \frac{2}{3}\text{MMS}(i)$ (Procaccia and Wang 2014). Furthermore, very recent work by Amanatidis et al. (2015) achieves the same approximation ratio in polynomial time.

These theoretical results have already made a significant real-world impact through *Spliddit* (`www.spliddit.org`), a not-for-profit fair division website (Goldman and Procaccia 2014). Since its launch in November 2014, Spliddit has attracted more than 52,000 users. The website currently offers five applications, for dividing goods, rent, credit, chores, and fare. Spliddit's algorithm for dividing goods, in particular, elicits additive valuations (which is easy to do), and maximizes social welfare (the total value players receive) subject to the highest feasible level of fairness among envy-freeness, proportionality, and MMS. If envy-freeness and proportionality are infeasible, the algorithm computes the maximum α such that all players can receive an α fraction of their MMS guarantee; since $\alpha \geq 2/3$ (Procaccia and Wang 2014), the solution is, in a sense, provably fair. The website summarizes the method's fairness guarantees as follows:

> *"We guarantee each participant at least two thirds of her maximin share. In practice, it is extremely likely that each participant will receive at least her full maximin share."*

Our goal in this paper is to better understand the second sentence of this quote: When is it possible to find an (exact) MMS allocation? And how "likely" is it?

Our results. Our first set of results has to do with the following question: what is the maximum $f(n)$ such that every instance with n players and $m \leq f(n)$ goods admits an MMS allocation? The previously known counterexample to the existence of MMS allocations uses a huge number of goods — n^n, to be exact (Procaccia and Wang 2014). Hence, $f(n) \leq n^n - 1$. Our first major result drastically improves this upper bound: an MMS allocation may not exist even when the number of goods is *linear* in the number of players.

Theorem 2.1. *For all $n \geq 3$, there is an instance with n players and $m \leq 3n+4$ goods such that an MMS allocation does not exist.*

That is, $f(n) \leq 3n + 3$. On the other hand, Bouveret and Lemaître (2014) show that $f(n) \geq n+3$. As a bonus result, we show in Appendix C that $f(n) \geq n + 4$.

The counterexamples to the existence of MMS allocations are extremely delicate, in the sense that an MMS allocation does exist if the valuations are even slightly perturbed. In addition, as mentioned above, randomly generated instances did not contain any counterexamples (Bouveret and Lemaître 2014). We formalize these observations by considering the regime where for each $i \in N$ there is a distribution $\mathcal{D}_i$ such that the values $V_i(g)$ are drawn independently from $\mathcal{D}_i$.

Theorem 3.1 *Assume that for all $i \in N$, $\mathbb{V}[\mathcal{D}_i] \geq c$ for a constant $c > 0$. Then for all $\varepsilon > 0$ there exists $K = K(c, \varepsilon)$ such that if $\max(n, m) \geq K$, then the probability that an MMS allocation exists is at least $1 - \varepsilon$.*

In words, an MMS allocation exists with high probability as the number of players *or* the number of goods goes to infinity. It was previously known that an envy-free allocation (and, hence, an MMS allocation) exists with high probability when $m \in \Omega(n \ln n)$ (Dickerson et al. 2014). Our analysis therefore focuses on the case of $m \in O(n \ln n)$. In this case, an envy-free allocation is unlikely to exist (such an allocation certainly does not exist when $m < n$), but (as we show) the existence of an MMS allocation is still likely. Specifically, we develop an allocation algorithm and show that it finds an MMS allocation with high probability. The algorithm's design and analysis leverage techniques for matching in random bipartite graphs.

2 Dependence on the Number of Goods

The main result of this section is the following theorem:

Theorem 2.1. *For all $n \geq 3$, there is an instance with n players and $m \leq 3n+4$ goods such that an MMS allocation does not exist.*

Note that when $n = 2$, an MMS allocation is guaranteed to exist: simply let player 1 divide the goods into two bundles according to his MMS partition, and let player 2 choose. Player 1 then obviously receives his MMS guarantee, whereas player 2 receives a bundle worth at least $V_2(G)/2 \geq \text{MMS}(2)$. The result of Procaccia and Wang (2014) shows that an MMS allocation may not exist even when $n = 3$ and $m = 12$ which proves the theorem for $n = 3$, but, as noted in Section 1, their construction requires n^n goods in general.

Because the new construction that proves Theorem 2.1 is somewhat intricate, we relegate the detailed proof to Appendix A. Here we explicitly provide the special case of $n = 4$. To this end, let us define the following two matrices, where ε is a very small positive constant ($\varepsilon = 1/16$ will suffice).

$$S = \begin{bmatrix} \frac{7}{8} & 0 & 0 & \frac{1}{8} \\ 0 & \frac{3}{4} & 0 & \frac{1}{4} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{8} & \frac{1}{4} & \frac{1}{2} & \frac{1}{8} \end{bmatrix},$$

$$T = \begin{bmatrix} 0 & \varepsilon^4 & 0 & -\varepsilon^4 \\ \varepsilon^3 & 0 & -\varepsilon^3 + \varepsilon^2 & -\varepsilon^2 \\ 0 & -\varepsilon^4 + \varepsilon & 0 & \varepsilon^4 - \varepsilon \\ -\varepsilon^3 & -\varepsilon & \varepsilon^3 - \varepsilon^2 & \varepsilon^2 + \varepsilon \end{bmatrix}$$

Let $M = S + T$. Crucially, the rows and columns of M sum to 1. Let G contain goods that correspond to the nonzero elements of M, that is, for every entry $M_{i,j} > 0$ we have a good $g_{i,j}$; note that $|G| = 14 \leq 3n + 4$.

Next, partition the 4 players into $P = \{1, 2\}$ and $Q = \{3, 4\}$. Define the valuations of the players in P as follows where $0 < \tilde{\varepsilon} \ll \varepsilon$ ($\tilde{\varepsilon} = 1/64$ will suffice).

$$M + \begin{bmatrix} 0 & 0 & 0 & -\tilde{\varepsilon} \\ 0 & 0 & 0 & -\tilde{\varepsilon} \\ 0 & 0 & 0 & -\tilde{\varepsilon} \\ 0 & 0 & 0 & 3\tilde{\varepsilon} \end{bmatrix}$$

That is, the values of the rightmost column are perturbed. For example, for $i \in P$, $V_i(g_{1,4}) = 1/8 - \varepsilon^4 - \tilde{\varepsilon}$. Similarly, for players in Q, the values of the bottom row are perturbed:

$$M + \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -\tilde{\varepsilon} & -\tilde{\varepsilon} & -\tilde{\varepsilon} & 3\tilde{\varepsilon} \end{bmatrix}$$

It is easy to verify that the MMS guarantee of all players is 1. Moreover, the unique MMS partition of the players in P (where every subset has value 1) corresponds to the columns of M, and the unique MMS partition of the players in Q corresponds to the rows of M. If we divide the goods by columns, one of the two players in Q will end up with a bundle of goods worth at most $1 - \tilde{\varepsilon}$ — less than his MMS value of 1. Similarly, if we divide the goods by rows, one of the players in P will receive a bundle worth only $1 - \tilde{\varepsilon}$. Any other partition of the goods will ensure that some party does not achieve their MMS value due to the relative size of $\tilde{\varepsilon}$.

3 Random Valuations

The counterexamples to the existence of MMS allocations — Theorem 2.1 and the construction of Procaccia and Wang (2014) — are very sensitive: tiny random perturbations are extremely likely to invalidate them. Our goal in this section is to prove MMS allocations do, in fact, exist with high probability, if a small amount of randomness is present.

To this end, let us consider a probabilistic model with the following features:

1. For all $i \in N$, $\mathcal{D}_i$ denotes a probability distribution over $[0, 1]$.
2. For all $i \in N, g \in G$, $V_i(g)$ is randomly sampled from $\mathcal{D}_i$.
3. The set of random variables $\{V_i(g)\}_{i\in N, g\in G}$ is mutually independent.

We will establish the following theorem:

Theorem 3.1. *Assume that for all $i \in N$, $\mathbb{V}[\mathcal{D}_i] \geq c$ for a constant $c > 0$. Then for all $\varepsilon > 0$ there exists $K = K(c, \varepsilon)$ such that if $\max(n, m) \geq K$, then the probability that an MMS allocation exists is at least $1 - \varepsilon$.*

In words, as long as each $\mathcal{D}_i$ has constant variance, if either the number of players or the number of goods goes to infinity, there exists an MMS allocation with high probability. In parallel, independent work, Amanatidis et al. (2015) establish (as one of several results) a special case of Theorem 3.1 where each $\mathcal{D}_i$ is the uniform distribution over $[0, 1]$. Dealing with arbitrary distributions presents significant technical challenges, and is also important in terms of explaining the abovementioned experiments, which cover a wide range of distributions. Yet the result of Amanatidis et al. is not completely subsumed by Theorem 3.1, as they carefully analyze the rate of convergence to 1.

Our starting point is a result by Dickerson et al. (2014), who study the existence of envy-free allocations. They show that an envy-free allocation exists with high probability as $m \to \infty$, as long as $n \in O(m/\ln m)$, and the distributions $\mathcal{D}_i$ satisfy the following conditions for all $i, j \in N$:

1. $\mathbb{P}\left[\arg\max_{k\in N} V_k(g) = \{i\}\right] = 1/n$.
2. There exist constants μ, μ^* such that

$$0 < \mathbb{E}\left[V_i(g) \;\middle|\; \arg\max_{k\in N} V_k(g) = \{j\}\right] \leq \mu < \mu^*$$
$$\leq \mathbb{E}\left[V_i(g) \;\middle|\; \arg\max_{k\in N} V_k(g) = \{i\}\right].$$

The proof uses a naïve allocation algorithm: simply give each good to the player who values it most highly. The first condition then implies that each player receives roughly $1/n$ of the goods, and the second condition ensures that each player has higher expected value for each of his own goods compared to goods allocated to other players.

It turns out that, via only slight modifications, their theorem can largely work in our setting. That is, alter their allocation algorithm to give a good g to a player i who believes g is in the top $1/n$ of their probability distribution $\mathcal{D}_i$. If there are multiple such players, choose one uniformly at random and if no such player exists, give it to any player uniformly at random.

This procedure is fairly straightforward for continuous probability distributions. For example, if player i's distribution $\mathcal{D}_i$ is uniform over the interval $[0, 1]$ then he believes g is in the top $1/n$ of $\mathcal{D}_i$ if $V_i(g) \geq (n-1)/n$. However, distributions with atoms require more care. For example, suppose D_i is $1/3$ with probability $7/8$ and uniform over $[1/2, 1]$ with probability $1/8$. Then if $n = 3$, i believes g is in the top $1/n$ of $\mathcal{D}_i$ if $V_i(g) > 1/3$ or if $V_i(g) = 1/3$ he should believe it is in his top $1/n$ only $1/n - 1/8 = 5/24$ of the time. To implement such a procedure, when sampling from $\mathcal{D}_i$, we should first sample from the uniform distribution over $[0, 1]$. If our sampled value is at least $(n - 1)/n$ we will say i has drawn from his top $1/n$. We then convert our sampled value to a sampled value from $\mathcal{D}_i$ by applying the inverse CDF.

Utilizing the observation that any envy-free allocation is also an MMS allocation we can then restate the result of Dickerson et al. (2014) as the following lemma, whose proof is relegated to Appendix B.

Lemma 3.2 ((Dickerson et al. 2014))**.** *Assume that for all $i \in N$, $\mathbb{V}[\mathcal{D}_i] \geq c$ for a constant $c > 0$. Then for all $\varepsilon > 0$ there exists $K = K(\varepsilon)$ such that if $m \geq K$ and $m \geq \alpha n \ln n$, for some $\alpha = \alpha(c)$, then the probability that an MMS allocation exists is at least $1 - \varepsilon$.*

Note that the statement of Lemma 3.2 is identical to that of Theorem 3.1, except for two small changes: only m is assumed to go to infinity, and the additional condition $m \geq \alpha n \ln n$. So it only remains to deal with the case of $m < \alpha n \ln n$. We can handle this scenario via consideration of the case $m < n^{8/7}$ — formalized in the following lemma.

Lemma 3.3. *For all $\varepsilon > 0$ there exists $K = K(\varepsilon)$ such that if $n \geq K$ and $m < n^{8/7}$, then the probability that an MMS allocation exists is at least $1 - \varepsilon$.*

Note that this lemma actually does not even require the minimum variance assumption, that is, we are proving a stronger statement than is needed for Theorem 3.1.

It is immediately apparent that when the number of goods is relatively small, we will not be able to prove the existence of MMS allocations via the existence of envy-free allocations. For example, envy-free allocations certainly do not exist if $m < n$, and are provably highly unlikely to exist if $m = n + o(n)$ (Dickerson et al. 2014). Our approach, to which we devote the remainder of this section, is significantly more intricate.

3.1 Proof of Lemma 3.3

We assume that $m > n$, because an MMS allocation always exists when $m \leq n$ (in fact, when $m \leq n + 4$, as Theorem C.1 shows). We will require the following notions and lemma.

Definition 3.4. A *ranking* of the goods G for some player $i \in N$ is the order of the goods by value from most valued to least. Ties are broken uniformly at random. Furthermore, a good g's *rank* for a player i is the position of g in i's ranking.

An important observation of the rankings that we will use often throughout this section is that the players' rankings are independent of each other.

Definition 3.5. Suppose $X \subseteq N$ and $Y \subseteq G$ where $|X| \leq |Y|$. Let

$$s = |X|\lceil |Y|/|X|\rceil - |Y|$$

and G be the bipartite graph where:

1. L represents the vertices on the left, and R on the right.
2. L is comprised of $\lfloor |Y|/|X| \rfloor$ copies of the first s players of X and $\lceil |Y|/|X| \rceil$ copies of the other players.
3. $R = Y$.
4. The i^{th} copy of a player has an edge to a good g iff g's rank is in $((i-1)\Delta, i\Delta]$ in the player's ranking where $\Delta = \ln^3 n$.

Note that $|L| = |R|$ since if we let $x = |X|$ and $y = |Y|$ (and therefore $s = x\lceil y/x \rceil - y$). Then

$$\begin{aligned} |L| &= s\lfloor y/x \rfloor + (x-s)\lceil y/x \rceil \\ &= x\lceil y/x \rceil - s\left(\lceil y/x \rceil - \lfloor y/x \rfloor\right). \end{aligned}$$

If x divides y, then we have that $\lceil y/x \rceil = \lfloor y/x \rfloor = \frac{y}{x}$ and so $|L| = y$. If, on the other hand, x does not divide y, then we have that $\lceil y/x \rceil - \lfloor y/x \rfloor = 1$ and so we have

$$\begin{aligned} |L| &= x\lceil y/x \rceil - s \\ &= x\lceil y/x \rceil - (x\lceil y/x \rceil - y) \\ &= y. \end{aligned}$$

Therefore, in either case, $|L| = y = |Y| = |R|$.

The *matched draft* on X and Y is the process of constructing G and producing an allocation corresponding to a perfect matching of G. That is, if a perfect matching exists then a player in X is given all goods the copies of it are matched to. In the event that no perfect matching exists, the matched draft is said to fail.

Lemma 3.6. *Suppose of the $m < n^{8/7}$ goods $x = \gamma\lfloor m/n \rfloor$ are randomly chosen and removed, where $\gamma \leq n^{1/3}$, and the remaining $\tilde{m} := m - x$ goods are allocated via a matched draft to $\tilde{n} := n - \gamma$ players. Then this matched draft succeeds with probability $\to 1$ as $n \to \infty$ (note that as $n \to \infty$, so too do $\tilde{n}, \tilde{m}$).*

Proof. Define d as the minimum degree of a vertex of L in G and $D = 2\lg n \ln n$. Then we have

$$\begin{aligned} &\mathbb{P}\left[\text{matched draft fails}\right] \\ &= \mathbb{P}\left[\text{matched draft fails} \mid d < D\right]\mathbb{P}\left[d < D\right] \\ &\quad + \mathbb{P}\left[\text{matched draft fails} \mid d \geq D\right]\mathbb{P}\left[d \geq D\right] \\ &\leq \mathbb{P}\left[d < D\right] + \mathbb{P}\left[\text{matched draft fails} \mid d \geq D\right]. \end{aligned}$$

Let us consider these two terms separately and show they $\to 0$ as $n \to \infty$.

If $x = 0$ we have that $\mathbb{P}\left[d < D\right] = 0$ for sufficiently large n, so let us assume $x > 0$. Denoting by p_D^{ij} the probability that player i has less than D of the goods ranked in positions $((j-1)\Delta, j\Delta]$ remaining, we have

$$\mathbb{P}\left[d < D\right] \leq \sum_{i=1}^{\tilde{n}} \sum_{j=1}^{\substack{\text{number of goods} \\ i \text{ receives}}} p_D^{ij}.$$

The right hand side is equal to $\tilde{m}$ times the probability that player 1 has less than D of the goods ranked in the top Δ positions remaining, which is equal to $\tilde{m}$ times the probability that of the x randomly chosen goods, more than $\Delta - D$ are ranked in the top Δ positions for player 1.

Now let the random variable X denote the number of the x random goods ranked in the top Δ for player 1. Clearly $\mathbb{E}\left[X\right] = \frac{\Delta x}{\tilde{m}}$. Thus by Markov's inequality we have that

$$\begin{aligned} &\mathbb{P}\left[X > \Delta - D\right] \\ &= \mathbb{P}\left[X > \mathbb{E}\left[X\right]\frac{\tilde{m}(\Delta - D)}{\Delta x}\right] \\ &\leq \frac{\mathbb{E}\left[X\right]}{\tilde{m}(\Delta - D)/(\Delta x)} \\ &= \left(\frac{\Delta x}{\tilde{m}}\right)^2 \frac{1}{\Delta - D} \\ &= \left(\frac{(\ln^3 n)(\gamma\lfloor m/n \rfloor)}{m - \gamma\lfloor m/n \rfloor}\right)^2 \frac{1}{\ln^3 n - 2\lg n \ln n} \\ &\leq \left(\frac{n^{10/21}\ln^3 n}{n - n^{10/21}}\right)^2 \frac{1}{\ln^3 n - 2\lg n \ln n} \\ &\to 0. \end{aligned}$$

Next let us consider $\mathbb{P}\left[\text{matched draft fails} \mid d \geq D\right]$. We would like to appeal to the plethora of results on perfect matchings in bipartite Erdös-Rényi graphs (Bollobás 2001) or random bipartite k-out graphs (McDiarmid 1980), but due to the lack of independence on the edge existences we do not satisfy a crucial assumption of much of this literature, and more importantly its proofs. We will therefore prove this in full here via an approach that allows us to ignore the dependence. We will utilize Hall's theorem and denote by $N(X)$ the set of neighbors of X in the bipartite graph G.

$$\begin{aligned} &\mathbb{P}\left[\text{matched draft fails} \mid d \geq D\right] \\ &= \mathbb{P}\left[\exists\, X \subseteq L \text{ s.t. } |X| < |N(X)| \mid d \geq D\right] \\ &\leq \sum_{X \subseteq L} \mathbb{P}\left[|X| < |N(X)| \mid d \geq D\right] \\ &\leq \sum_{i=D}^{\tilde{m}} \sum_{\{X \subseteq L \mid |X| = i\}} \sum_{\{Y \subseteq R \mid |Y| = i-1\}} \mathbb{P}\left[N(X) \subseteq Y \mid d \geq D\right]. \end{aligned}$$

If the edges of G were independent then we would find that for $|X| = i$ and $|Y| = i - 1$,

$$\mathbb{P}\left[N(X) \subseteq Y\right] = \left(\frac{i-1}{\tilde{m}}\right)^{\sum_{x \in X} |N(x)|},$$

and more importantly

$$\mathbb{P}\left[N(X) \subseteq Y \mid d \geq D\right] \leq \left(\frac{i-1}{\tilde{m}}\right)^{iD}. \qquad (1)$$

Via our independence assumptions in our randomized setting there is only one form of dependence in the edges of G. Specifically, if we take all copies of any player $i \in L$, then their neighbors in R never intersect. Though this does indeed introduce dependence into our system, note that we still have that Equation (1) as the dependence only lowers the probability of $N(X)$ "fitting" into Y. We therefore find

$$\begin{aligned}
&\mathbb{P}\left[\text{matched draft fails} \mid d \geq D\right] \\
&\leq \sum_{i=D}^{\tilde{m}} \sum_{\{X \subseteq L \mid |X|=i\}} \sum_{\{Y \subseteq R \mid |Y|=i-1\}} \left(\frac{i-1}{\tilde{m}}\right)^{iD} \\
&= \sum_{i=D}^{\tilde{m}} \binom{\tilde{m}}{i}\binom{\tilde{m}}{i-1}\left(\frac{i-1}{\tilde{m}}\right)^{iD} \\
&\leq \sum_{i=D}^{\lfloor \tilde{m}/2 \rfloor} \binom{\tilde{m}}{i}\binom{\tilde{m}}{i-1}\left(\frac{i-1}{\tilde{m}}\right)^{iD} \\
&\quad + \sum_{i=\lceil \tilde{m}/2 \rceil}^{\tilde{m}} \binom{\tilde{m}}{\tilde{m}-i}\binom{\tilde{m}}{\tilde{m}-i+1}\left(\frac{i-1}{\tilde{m}}\right)^{iD} \\
&= \sum_{i=D}^{\lfloor \tilde{m}/2 \rfloor} \binom{\tilde{m}}{i}\binom{\tilde{m}}{i-1}\left(\frac{i-1}{\tilde{m}}\right)^{iD} \\
&\quad + \sum_{j=0}^{\lfloor \tilde{m}/2 \rfloor} \binom{\tilde{m}}{j}\binom{\tilde{m}}{j+1}\left(\frac{\tilde{m}-j-1}{n}\right)^{(\tilde{m}-j)D}.
\end{aligned}$$

We now show both of these terms separately $\to 0$ as $n \to \infty$.

First,

$$\begin{aligned}
&\sum_{i=D}^{\lfloor \tilde{m}/2 \rfloor} \binom{\tilde{m}}{i}\binom{\tilde{m}}{i-1}\left(\frac{i-1}{\tilde{m}}\right)^{iD} \\
&\leq \sum_{i=D}^{\lfloor \tilde{m}/2 \rfloor} \left(\frac{\tilde{m}e}{i}\right)^{i}\left(\frac{\tilde{m}e}{i-1}\right)^{i-1}\left(\frac{i-1}{\tilde{m}}\right)^{iD} \\
&\leq \sum_{i=D}^{\lfloor \tilde{m}/2 \rfloor} \left(\frac{\tilde{m}e}{i-1}\right)^{2i-1}\left(\frac{i-1}{\tilde{m}}\right)^{iD} \\
&= \sum_{i=D}^{\lfloor \tilde{m}/2 \rfloor} \left(\frac{i-1}{\tilde{m}}\right)^{i(D-2)+1} e^{2i-1} \\
&\leq \sum_{i=D}^{\lfloor \tilde{m}/2 \rfloor} \frac{e^{2i-1}}{2^{i(D-2)+1}} \\
&\leq \sum_{i=D}^{\lfloor \tilde{m}/2 \rfloor} \frac{4e^2}{2^D} \\
&\leq \frac{2e^2 n^{8/7}}{n^{2\ln n}} \\
&\to 0,
\end{aligned}$$

where the first inequality follows from the fact that $\binom{a}{b} \leq \left(\frac{ae}{b}\right)^b$ for $b > 0$, and the third inequality follows from the fact that $i \leq \lfloor \tilde{m}/2 \rfloor$.

Second,

$$\begin{aligned}
&\sum_{j=0}^{\lfloor \tilde{m}/2 \rfloor} \binom{\tilde{m}}{j}\binom{\tilde{m}}{j+1}\left(\frac{\tilde{m}-j-1}{\tilde{m}}\right)^{(\tilde{m}-j)D} \\
&\leq \tilde{m}\left(\frac{\tilde{m}-1}{\tilde{m}}\right)^{\tilde{m}D} \\
&\quad + \sum_{j=1}^{\lfloor \tilde{m}/2 \rfloor} \left(\frac{\tilde{m}e}{j}\right)^{j}\left(\frac{\tilde{m}e}{j+1}\right)^{j+1}\left(\frac{\tilde{m}-j-1}{\tilde{m}}\right)^{(\tilde{m}-j)D} \\
&\leq \tilde{m}\left(1-\frac{1}{\tilde{m}}\right)^{\tilde{m}D} \\
&\quad + \sum_{j=1}^{\lfloor \tilde{m}/2 \rfloor} \left(\frac{\tilde{m}e}{j}\right)^{2j+1}\left(1-\frac{j+1}{\tilde{m}}\right)^{(\tilde{m}-j)D} \\
&\leq \tilde{m}e^{-D} + \sum_{j=1}^{\lfloor \tilde{m}/2 \rfloor} \left(\frac{\tilde{m}e}{j}\right)^{2j+1} e^{-D(j+1)(\tilde{m}-j)/\tilde{m}} \\
&\leq \tilde{m}e^{-D} + \sum_{j=1}^{\lfloor \tilde{m}/2 \rfloor} \left(\frac{\tilde{m}e}{j}\right)^{2j+1} e^{-D(j+1)/2} \\
&\leq \tilde{m}e^{-D} + \sum_{j=1}^{\lfloor \tilde{m}/2 \rfloor} \left(\frac{\tilde{m}^2e^2}{j^2e^{D/2}}\right)^{j+1} \\
&\leq \frac{n^{8/7}}{n^{2\lg n}} + \sum_{j=1}^{\lfloor n^{8/7}/2 \rfloor} \left(\frac{(n^{8/7})^2e^2}{j^2n^{\lg n}}\right)^{j+1} \\
&\leq \frac{n^{8/7}}{n^{2\lg n}} + \lfloor n^{8/7}/2 \rfloor \left(\frac{(n^{8/7})^2e^2}{n^{\lg n}}\right) \\
&\to 0,
\end{aligned}$$

where the first inequality follows from $\binom{a}{b} \leq \left(\frac{ae}{b}\right)^b$ for $b > 0$ and the third inequality follows from $1 + x \leq e^x$ for all x.

Thus, we find that as $n \to \infty$ the matched draft succeeds with probability $\to 1$. ■

We are now ready to prove the lemma.

Proof of Lemma 3.3. Recall that we may assume that $m > n$. We will ensure every player has at most one less good than any other player. Let s then represent the number of players that receive one less good than any other player, that is,

$$s = n\lceil m/n \rceil - m.$$

We consider two separate cases here.

Case 1: $s \leq n^{1/3}$. In this scenario we do the following.

1. If possible, give each of the first s players their top $\lfloor m/n \rfloor$ goods. Otherwise, fail to produce any allocation.
2. Hold a matched draft for the remaining $(n - s)\lceil m/n \rceil$ goods and $n - s$ players.

We first show that as $n \to \infty$ this procedure successfully produces an allocation with probability $\to 1$.

Consider the probability that the first step of the procedure successfully completes. That is, the first s players each get their top $\lfloor m/n \rfloor$ goods. Similarly to a birthday paradox like argument we get that this occurs with probability at least

$$\prod_{i=1}^{s\lfloor m/n \rfloor} \left(1 - \frac{i-1}{m}\right) > \left(1 - \frac{sm/n}{m}\right)^{sm/n}$$
$$\geq \left(1 - \frac{1}{n^{2/3}}\right)^{n^{10/21}}.$$

But as

$$\lim_{x\to\infty} \left(1 - \frac{1}{\omega(x)}\right)^x = 1$$

we find that this too goes to 1 as $n \to \infty$.

Now consider the second step of the procedure. By Lemma 3.6 with $\gamma = s$, we know that this succeeds with probability 1 as $n \to \infty$. Therefore the entire procedure will successfully complete with probability $\to 1$ as $n \to \infty$.

Therefore, to prove the theorem, it suffices to show that if the procedure successfully completes, then we have an MMS allocation. Since for every player any MMS partition must include a subset with at most $\lfloor m/n \rfloor$ goods and the first s players are given their top $\lfloor m/n \rfloor$ goods, they must receive their MMS value.

Let us turn our attention then to the remaining $n-s$ players. Upon successful completion of the matched draft, we know that all of these players will receive goods ranked in their top $\Delta\lceil m/n \rceil$. We claim that for sufficiently large n any player's MMS partition must include a subset of at most $\lceil m/n \rceil$ goods where each good is ranked lower than $\Delta\lceil m/n \rceil$. Suppose this were not true for purposes of contradiction. Then each of the n subsets in an offending player's MMS partition must include either one of the top $\Delta\lceil m/n \rceil$ goods or $\lceil m/n \rceil + 1$ goods. We then see that for sufficiently large n, the number of such subsets is bounded by

$$\begin{aligned}
&\Delta\lceil m/n \rceil + \frac{m - \Delta\lceil m/n \rceil}{\lceil m/n \rceil + 1} \\
&= \Delta\lceil m/n \rceil + \frac{s(\lceil m/n \rceil - 1) + (n-s)\lceil m/n \rceil - \Delta\lceil m/n \rceil}{\lceil m/n \rceil + 1} \\
&= \frac{\Delta\lceil m/n \rceil^2 + n\lceil m/n \rceil - s}{\lceil m/n \rceil + 1} \\
&\leq \frac{\lceil m/n \rceil}{\lceil m/n \rceil + 1} n + \Delta\lceil m/n \rceil \\
&\leq \frac{n^{1/7}}{n^{1/7}+1} n + n^{1/7}\ln^3 n \\
&< n.
\end{aligned}$$

Thus the offending player cannot produce such an MMS partition which proves the claim.

Now note that the $n-s$ players of interest have MMS partitions that include the same number of goods they received, but all of which are worth strictly less than every good in their bundle. They therefore must have achieved their MMS value.

Case 2: $s > n^{1/3}$. In this scenario we simply run a matched draft. Similarly to the previous case we know from Lemma 3.6 with $\gamma = 0$ that all the players will receive goods ranked in their top $\Delta\lceil m/n \rceil$ with probability $\to 1$ as $n \to \infty$.

In this case for sufficiently large n any player's MMS partition must include a subset of at most $\lfloor m/n \rfloor$ goods where each good is ranked lower than $\Delta\lceil m/n \rceil$. Again, suppose this were not true for purposes of contradiction. Then each of the n subsets in a player's MMS partition must include either one of the top $\Delta\lceil m/n \rceil$ goods or $\lfloor m/n \rfloor + 1 = \lceil m/n \rceil$ goods (in this case $m \not\equiv 0 \pmod n$). We then see that for sufficiently large n, the number of subsets is at most

$$\begin{aligned}
&\Delta\lceil m/n \rceil + \frac{m - \Delta\lceil m/n \rceil}{\lceil m/n \rceil} \\
&= \Delta\lfloor m/n \rfloor + \frac{s(\lceil m/n \rceil - 1) + (n-s)\lceil m/n \rceil}{\lceil m/n \rceil} \\
&= n + \Delta\lfloor m/n \rfloor - \frac{s}{\lceil m/n \rceil} \\
&\leq n + n^{1/7}\ln^3 n - \frac{n^{1/3}}{n^{1/7}} \\
&< n.
\end{aligned}$$

Via logic similar to the previous case, we conclude that all players must have achieved their MMS value. ■

4 Discussion

Theorem 3.1, together with the extensive experiments of Bouveret and Lemaître (2014), tells us that an MMS allocation is very likely to exist *ex post*, that is, after the players report their preferences. But imagine a parallel universe where, unless the number of goods is huge — e.g., n^n like in the construction of Procaccia and Wang (2014) — an MMS allocation must exist. In such a universe one would typically be able to *guarantee* to players an MMS allocation *a priori* — before they even enter their preferences (simply based on n and m). In contrast, in our universe we can typically (unless m is quite small) only guarantee an approximate MMS allocation. As far as people's perception is concerned, we believe that the difference between the two universes is significant — and Theorem 2.1 places us, for the first time, in the latter universe.

While Theorem 2.1 essentially settles one of the main open problems of Procaccia and Wang (2014), it sheds no light on the other: For each number of players n, what is the maximum $g(n) \in (0,1)$ such that it is always possible to achieve a $g(n)$-approximate MMS allocation, that is, an allocation satisfying $V_i(A_i) \geq g(n) \cdot \mathrm{MMS}(i)$ for all players i. Procaccia and Wang prove that

$$g(n) \geq \frac{2\lfloor n \rfloor_{\mathrm{odd}}}{3\lfloor n \rfloor_{\mathrm{odd}} - 1},$$

where $\lfloor n \rfloor_{\mathrm{odd}}$ is the largest odd n' such that $n' \leq n$. In particular, for all n we have that $g(n) > 2/3$, and $g(3) \geq 3/4$. Amanatidis et al. (2015) establish (among their other results) an improved bound of $g(3) \geq 7/8$, but do not improve the general lower bound. On the other hand, counterexamples to the existence of MMS allocations — the construction of Procaccia and Wang (2014), and the proof of Theorem 2.1 — only imply that $g(n) \leq 1 - o(1)$, that is, they give an upper bound that is extremely close to 1. The challenge of closing this gap is, in our view, both technically fascinating and practically significant.

References

Amanatidis, G.; Markakis, E.; Nikzad, A.; and Saberi, A. 2015. Approximation algorithms for computing maximin share allocations. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP)*, 39–51.

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.

Bollobás, B. 2001. *Random Graphs*. Cambridge University Press, 2nd edition.

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)*, 1321–1328.

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.

Dickerson, J. P.; Goldman, J.; Karp, J.; Procaccia, A. D.; and Sandholm, T. 2014. The computational rise and fall of fairness. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 1405–1411.

Goldman, J., and Procaccia, A. D. 2014. Spliddit: Unleashing fair division algorithms. *SIGecom Exchanges* 13(2):41–46.

McDiarmid, C. J. H. 1980. Clutter percolation and random graphs. *Mathematical Programming Study* 13:17–25.

Procaccia, A. D., and Wang, J. 2014. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, 675–692.

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

Roos, M., and Rothe, J. 2010. Complexity of social welfare optimization in multiagent resource allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 641–648.

A Proof of Theorem 2.1

With the illustrative $n = 4$ example under our belt, we now prove the general case where $n \geq 4$ (recall that Procaccia and Wang (2014) establish the result for $n = 3$). The crux of the argument is proving the existence of a matrix $M \in \mathbb{R}^{n\times n}$ with the following properties:

1. $\forall i, j : M_{i,j} \geq 0$.
2. $\forall i : M_{i,n}, M_{n,i} > 0$.
3. The sum of a row or column of M is 1 (i.e. $M\mathbf{1} = M^T\mathbf{1} = \mathbf{1}$).
4. Define M^+ as the set of all positive entries in M. Then if we wish to partition M^+ into n subsets that sum to exactly 1 then our partition must correspond to the rows of M or the columns of M.

To begin, let $S \in \mathbb{R}^{n\times n}$ be the following matrix.

$$\begin{bmatrix} \frac{2^{n-1}-1}{2^{n-1}} & 0 & 0 & \cdots & 0 & 0 & \frac{1}{2^{n-1}} \\ 0 & \frac{2^{n-2}-1}{2^{n-2}} & 0 & \cdots & 0 & 0 & \frac{1}{2^{n-2}} \\ 0 & 0 & \frac{2^{n-3}-1}{2^{n-3}} & \cdots & 0 & 0 & \frac{1}{2^{n-3}} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \\ 0 & 0 & 0 & \cdots & \frac{3}{4} & 0 & \frac{1}{4} \\ 0 & 0 & 0 & \cdots & 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2^{n-1}} & \frac{1}{2^{n-2}} & \frac{1}{2^{n-3}} & \cdots & \frac{1}{4} & \frac{1}{2} & \frac{1}{2^{n-1}} \end{bmatrix},$$

i.e.,

$$S_{i,j} = \begin{cases} \frac{2^{n-i}-1}{2^{n-i}} & \text{if } i = j \neq n \\ \frac{1}{2^{n-j}} & \text{if } i = n \text{ and } j \neq n \\ \frac{1}{2^{n-i}} & \text{if } j = n \text{ and } i \neq n \\ \frac{1}{2^{n-1}} & \text{if } i = j = n \\ 0 & \text{otherwise.} \end{cases}$$

Now for $\varepsilon \approx 0$ where $\varepsilon > 0$, and for all $i \in \{1, ..., n-2\}$, let $r_i = \varepsilon^{2n-2i-2}$, and $c_i = \varepsilon^{2n-2i-3}$. Specifically, this implies:

$$0 < r_1 \ll c_1 \ll r_2 \ll c_2 \ll ... \ll r_{n-2} \ll c_{n-2} = \varepsilon \approx 0.$$

Furthermore, let $T \in \mathbb{R}^{n\times n}$ be the matrix given by:

$$\begin{bmatrix} 0 & v_1 & 0 & \cdots & 0 & 0 & -r_1 \\ u_1 & 0 & v_2 & \cdots & 0 & 0 & -r_2 \\ 0 & u_2 & 0 & \cdots & 0 & 0 & -r_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & v_{n-2} & -r_{n-2} \\ 0 & 0 & 0 & \cdots & u_{n-2} & 0 & -y \\ -c_1 & -c_2 & -c_3 & \cdots & -c_{n-2} & -x & z \end{bmatrix}$$

where the only nonzero entries are on the first diagonals above and below the main diagonal, and the last row and column.

Assign positive values to the u_i, v_i, x, y, and z such that all rows and columns sum to zero. A bit of arithmetic then gives:

$$u_i = \left(\sum_{j \leq i, j \equiv i \pmod 2} c_j \right) - \left(\sum_{j \leq i, j \not\equiv i \pmod 2} r_j \right) \approx c_i$$

$$v_i = \left(\sum_{j \leq i, j \equiv i \pmod 2} r_j \right) - \left(\sum_{j \leq i, j \not\equiv i \pmod 2} c_j \right) \approx r_i$$

$$x = v_{n-2} \approx r_{n-2}$$

$$y = u_{n-2} \approx c_{n-2}$$

$$z = \left(\sum_{j \leq n-2, j \equiv n \pmod 2} r_j + c_j \right) \approx c_{n-2}.$$

Now define $M = S + T$ and M^+ as the set of nonzero elements of M. Moreover, for a (finite) set $X \subset \mathbb{R}$, let $\sum X = \sum_{x \in X} x$. Then we see for sufficiently small ε that the following properties hold.

[P1] $0 \leq M_{i,j}$ (elements of M^+ are positive), and if $S_{i,j} \neq 0$ or $T_{i,j} \neq 0$, then $M_{i,j} > 0$.

[P2] $M_{i,j} \approx S_{i,j}$.

[P3] The sum of a row or column of M is 1 (i.e. $M\mathbf{1} = M^T\mathbf{1} = \mathbf{1}$).

[P4] $\forall i \in [n-1]$ if we have $X \subseteq M^+$ s.t. $M_{i,i} \in X$ and $\sum X = 1$ then one of the following is true:

(a) $M_{i,n} \in X$.
(b) $M_{n,i} \in X$.
(c) $M_{1,n}, M_{2,n}, ..., M_{i-1,n}, M_{n,n} \in X$.
(d) $M_{n,1}, M_{n,2}, ..., M_{n,i-1}, M_{n,n} \in X$.
(e) $\exists j, k < i$ s.t. $M_{j,n}, M_{n,k} \in X$.

This is easy to see when we take note that $S \approx M$ by [P2].

[P5] If $X \subseteq M^+$ s.t. $\sum X = r_i$, then $X = \{M_{i,i-1}, M_{i,i+1}\}$.

[P6] If $X \subseteq M^+$ s.t. $\sum X = c_i$, then $X = \{M_{i-1,i}, M_{i+1,i}\}$.

[P7] If $X \subseteq M^+$ s.t. $\sum X = x$, then $X = \{M_{n-2,n-1}\}$.

[P8] If $X \subseteq M^+$ s.t. $\sum X = y$, then $X = \{M_{n-1,n-2}\}$.

We now make a key observation with respect to M.

Lemma A.1. *Suppose we partition M^+ into n subsets such that the elements in each subset of the partition sum exactly to 1. Then for sufficiently small ε, the partition must correspond to the rows of M or the columns of M.*

Proof. Let us first consider the subset in the partition which includes $M_{1,1}$, call X_1. We wish to prove that X_1 is either:

1. the first row: $\{M_{1,1}, M_{1,2}, M_{1,n}\}$
2. the first column: $\{M_{1,1}, M_{2,1}, M_{n,1}\}$.

By [P4] we see that exactly one of $M_{n,1}$, $M_{1,n}$, and $M_{n,n}$ must be part of our subset X_1.

1. Suppose $M_{n,n} \in X_1$. Then $\sum X_1 \geq M_{1,1} + M_{n,n} = 1 + z > 1$. This is therefore impossible.

2. Suppose $M_{1,n} \in X_1$. As $M_{1,1} + M_{1,n} = 1 - r_1$ we see that by [P5] we must have $M_{1,2} \in X_1$. Then X_1 corresponds to the first row.
3. Suppose $M_{n,1} \in X_1$. As $M_{1,1} + M_{n,1} = 1 - c_1$ we see that by [P6] we must have $M_{2,1} \in X_1$. Then X_1 corresponds to the first column.

Now suppose we wish to find a partition as in the lemma's statement such that the first $i-1$ rows are in the partition where $i \in \{2, ..., n\}$. Then we claim row i must be in the partition as well. Importantly, this implies that if the first row is to be in the partition, then the partition must be the rows.

We first consider the case where $i \leq n-1$. Let X_i denote the subset in the partition that includes $M_{i,i}$. By [P4] we see that we must have one of the following.

1. $M_{i,n} \in X_i$.
If $n \leq n-2$ we find that $M_{i,i} + M_{i,n} = -r_i$ and so by [P5] we have $M_{i,i-1}, M_{i,i+1} \in X_i$. We therefore find that $X_1 = \{M_{i,i-1}, M_{i,i}, M_{i,i+1}, M_{i,n}\}$. On the other hand, if $i = n-1$ we find that $M_{i,i} + M_{i,n} = -y$ and so by [P8] we have $M_{n-1,n-2} \in X_i$. Thus $X_1 = \{M_{n-1,n-2}, M_{n-1,n-1}, M_{n-1,n}\}$. In either case X_i is the i^{th} row.
2. $M_{n,i} \in X_i$.
If $n \leq n-2$ we find that $M_{i,i} + M_{n,i} = -c_i$ and so by [P6] we have $M_{i-1,i} \in X_i$. But $M_{i-1,i}$ is in a previous row, which by our assumption is already assigned to a subset in the partition. On the other hand, if $i = n-1$ we have $M_{i,i} + M_{n,i} = -x$ and so by [P7] we have $M_{n-2,n-1} \in X_i$. Similarly to before, this element is in a previous row and thus is already assigned to a subset in the partition.
3. $M_{1,n}, M_{2,n}, ..., M_{i-1,n}, M_{n,n} \in X_i$.
As $M_{1,n}$ is in a previous row it is already assigned to a subset in the partition.
4. $M_{n,1}, M_{n,2}, ..., M_{n,i-1}, M_{n,n} \in X_i$.
This is impossible because
$$\begin{aligned}\sum X_i &\geq M_{i,i} + M_{n,1} + M_{n,2} + ... + M_{n,i-1} + M_{n,n} \\ &= 1 - r_1 - r_2 - ... - r_{i-1} + z \\ &= 1 + r_i + r_{i+1} + ... + r_{n+2} + y \\ &> 1.\end{aligned}$$
5. $\exists j, k < i$ s.t. $M_{j,n}, M_{n,k} \in X_i$.
As $M_{j,n}$ is in a previous row it is already assigned to a subset in the partition.

Next, suppose $i = n$. In this case, since we are only allowed n subsets in this partition, all remaining entries (i.e. the last row) must be in the last set. By [P3] we know this last row sums to 1. We therefore have shown that if the first row is in the partition, then the partition simply corresponds to the rows. A similar argument gives an analogous result for columns. As the first row or first column must be a subset in the partition (namely as X_1) we are done. ■

To prove the theorem, we now consider our construction through the lens of MMS allocations.

Proof of Theorem 2.1. We first show that there exists a set of $5n-6$ such goods for $n \geq 4$.

Partition the n players into two groups P and Q such that $|P|, |Q| \geq 2$. For each element such that $M_{i,j} > 0$, we will define a good $g_{i,j}$ (note that there are $5n-6$ such goods). For $k \in P$, we define

$$V_k(g_{i,j}) = \begin{cases} M_{i,j} & \text{if } j < n \\ M_{i,j} - \tilde{\varepsilon} & \text{if } j = n \text{ and } i < n \\ M_{i,j} + (n-1)\tilde{\varepsilon} & \text{if } j = n \text{ and } i = n \end{cases}$$

and similarly, for $k \in Q$, let

$$V_k(g_{i,j}) = \begin{cases} M_{i,j} & \text{if } i < n \\ M_{i,j} - \tilde{\varepsilon} & \text{if } i = n \text{ and } j < n \\ M_{i,j} + (n-1)\tilde{\varepsilon} & \text{if } i = n \text{ and } j = n \end{cases}$$

where $\tilde{\varepsilon}$ is small enough to ensure all $V_k(g_{i,j}) \geq 0$.

As all players in P (respectively Q) can partition the goods into columns (respectively rows) such that the value of each subset in the partition is exactly 1, the MMS guarantee of all players in P (respectively Q) must remain 1.

Next, let us consider an allocation of the goods. Lemma A.1 tells us that if the $V_k(g_{i,j})$ were exactly equal to the $M_{i,j}$ there are only two ways to allocate the goods such that every subset in the partition has value 1 (i.e. we get an MMS allocation): via the rows or via the columns. But note that the alteration to the value of a good $g_{i,j}$ from $M_{i,j}$ is at most $(n-1)\tilde{\varepsilon}$ and indeed no subset of goods can have its total value altered by more than $(n-1)\tilde{\varepsilon}$ for any player. Therefore, we claim that if we wish to have any hope of achieving an MMS allocation we must still partition according to the rows or columns (assuming $\tilde{\varepsilon}$ is sufficiently small). To see this, define

$$\gamma = \max_{X_1,...,X_n \in \mathcal{X}} \min_{i \in [n]} \sum X_i$$

where $\mathcal{X}$ is the set of partitions of M^+ excluding the rows and the columns. Importantly, via Lemma A.1 and the finite nature of $\mathcal{X}$ we know that $\gamma < 1$. Now suppose $\tilde{\varepsilon} < \frac{1-\gamma}{n-1}$. Then for any allocation that did not correspond to the rows or columns some player must have value at most $\gamma + (n-1)\tilde{\varepsilon} < 1$. This proves the claim.

Now note that if we split via rows the players of P will believe only the last row is worth at least 1 and all other rows are worth strictly less than 1. As there are at least two players in P, not all players can receive their MMS guarantee. A similar issue occurs when we split via the columns for the players in Q. Therefore, there exists no MMS allocation in this setting.

We have just shown the result for $5n-6$ goods (for $n \geq 4$) and now set our sights on $3n+4$ goods. Let $\tilde{n} = \lceil (n+4)/2 \rceil \geq 4$. We know that we can find $5\tilde{n}-6$ goods that do not admit an MMS allocation for $\tilde{n}$ players. Take this set of goods, and let there be n players such that $\lfloor n/2 \rfloor$ players are in group P and the remaining $\lceil n/2 \rceil$ are in group Q. Finally, add $n - \tilde{n}$ goods each of value 1 to all players. Note that the number of goods is:

$$m = (5\tilde{n}-6) + (n-\tilde{n}) = 4\lceil (n+4)/2 \rceil + n - 6 \leq 3n+4.$$

We then see that every player's MMS value remains 1, but even after the $n - \tilde{n}$ goods of value 1 are allocated we still must have at least 2 players in both P and Q when we allocate the original $5\tilde{n} - 6$ goods. Therefore, there cannot exist an MMS allocation. ■

B Proof of Lemma 3.2

The crux of the proof of Dickerson et al. (2014) relies on the allocation algorithm only satisfying the following two properties.

1. For any good g, if we do not condition on the $V_i(g)$, then every player has a $1/n$ probability of receiving g.
2. For some constant Δ, we have that

   $$\mathbb{E}\left[V_i(g) \mid i \text{ receives } g\right] - \mathbb{E}\left[V_i(g) \mid i \text{ does not receive } g\right] \geq \Delta.$$

We must show that our allocation algorithm implies these two properties in our setting. The first is clear via symmetry and so we turn our attention to the second. We claim that $\Delta = c/16$ suffices (recall that $\mathbb{V}\left[\mathcal{D}_i\right] \geq c$).

Let $X \sim \mathcal{D}_i$, $\bar{X} = \mathbb{E}\left[X\right]$, $p = \mathbb{P}\left[X < \bar{X}\right]$, and γ represent the value such that $\mathbb{P}\left[X \geq \gamma\right] = 1/n$.[1] We first show that $\mathbb{E}\left[X \mid X \geq \gamma\right] - \mathbb{E}\left[X \mid X < \gamma\right] \geq c/2$.

$$\begin{aligned} c &\leq \mathbb{V}\left[X\right] \\ &= \mathbb{E}\left[(X - \bar{X})^2\right] \\ &\leq \mathbb{E}\left[|X - \bar{X}|\right] \\ &= p\mathbb{E}\left[\bar{X} - X \mid X < \bar{X}\right] + (1-p)\mathbb{E}\left[X - \bar{X} \mid X \geq \bar{X}\right] \\ &= -p\mathbb{E}\left[X \mid X < \bar{X}\right] + (1-p)\mathbb{E}\left[X \mid X \geq \bar{X}\right] + (2p-1)\bar{X} \end{aligned}$$

Knowing that $\bar{X} = p\mathbb{E}\left[X \mid X < \bar{X}\right] + (1-p)\mathbb{E}\left[X \mid X \geq \bar{X}\right]$ there are two cases.

1. $\gamma < \bar{X}$. It follows that

   $$c \leq 2p\left(\bar{X} - \mathbb{E}\left[X \mid X < \bar{X}\right]\right),$$

   and therefore

   $$\frac{c}{2p} \leq \bar{X} - \mathbb{E}\left[X \mid X < \bar{X}\right].$$

   We conclude that

   $$\begin{aligned} c/2 &\leq \bar{X} - \mathbb{E}\left[X \mid X < \bar{X}\right] \\ &\leq \mathbb{E}\left[X \mid X \geq \gamma\right] - \mathbb{E}\left[X \mid X < \gamma\right]. \end{aligned}$$

2. $\gamma \geq \bar{X}$. It follows that

   $$c \leq 2(1-p)\left(\mathbb{E}\left[X \mid X \geq \bar{X}\right] - \bar{X}\right),$$

   and therefore

   $$\frac{c}{2(1-p)} \leq \mathbb{E}\left[X \mid X \geq \bar{X}\right] - \bar{X}.$$

   We conclude that

   $$\begin{aligned} c/2 &\leq \mathbb{E}\left[X \mid X \geq \bar{X}\right] - \bar{X} \\ &\leq \mathbb{E}\left[X \mid X \geq \gamma\right] - \mathbb{E}\left[X \mid X < \gamma\right]. \end{aligned}$$

[1] As discussed previously, such a γ may not exist in distributions with atoms, but we ignore this possibility purely for ease of exposition.

We now can show the second property. Since

$$\begin{aligned} \mathbb{E}\left[V_i(g) \mid i \text{ receives } g\right] = \left(1 - \frac{1}{n}\right)^n \mathbb{E}\left[X \mid X < \gamma\right] \\ + \left(1 - \left(1 - \frac{1}{n}\right)^n\right)\mathbb{E}\left[X \mid X \geq \gamma\right], \end{aligned}$$

and

$$\begin{aligned} &\mathbb{E}\left[V_i(g) \mid i \text{ does not receive } g\right] \leq \mathbb{E}\left[V_i(g)\right] \\ &= \left(1 - \frac{1}{n}\right)\mathbb{E}\left[X \mid X < \gamma\right] + \frac{1}{n}\mathbb{E}\left[X \mid X \geq \gamma\right], \end{aligned}$$

we have that

$$\begin{aligned} &\mathbb{E}\left[V_i(g) \mid i \text{ receives } g\right] - \mathbb{E}\left[V_i(g) \mid i \text{ does not receive } g\right] \\ &\geq \left(\left(1 - \frac{1}{n}\right) - \left(1 - \frac{1}{n}\right)^n\right)\left(\mathbb{E}\left[X \mid X \geq \gamma\right] - \mathbb{E}\left[X \mid X < \gamma\right]\right) \\ &\geq \left(\frac{1}{2} - \frac{1}{e}\right)\frac{c}{2} \\ &\geq \frac{c}{16}. \end{aligned}$$

■

C A Small Number of Goods Guarantees the Existence of an MMS Allocation

In Section 2 we have seen that even with $m = 3n+4$ goods, we cannot guarantee an MMS allocation. On the other hand, note that when $m < n$, the MMS guarantee of each player is 0, so any allocation is an MMS allocation. If $m = n$, we have that $\text{MMS}(i) = \min_{g \in G} V_i(g)$, so any allocation that gives a single good to each player fits the bill. The case of $m = n + 1$ is still trivial: the MMS partition of a player puts his two least desirable goods in one bundle, and every other good in a singleton bundle. Therefore, it is sufficient to let the players choose a single good each in the order $1, \ldots, n-1$, and allocate to player n the two remaining goods. Bouveret and Lemaître (2014) extend this argument to show that an MMS allocation exists whenever $m \leq n+3$.

In this appendix, which we view as an aside, we slightly improve the bound of Bouveret and Lemaître (2014) to $m \leq n + 4$. While the improvement is not of major excitement, we include it as we believe the approach is quite interesting and its ideas may be used to further hone the bounds.

Theorem C.1. *If $m \leq n + 4$ then there exists an MMS allocation.*

Proof. We give a detailed algorithm (with some commentary) to handle the case of $m \leq n + 4$ as Algorithm 1, but we highlight some of the nuances here.

Observe that whenever a player believes a single good is worth at least his MMS value we can give that good to him and in the reduced problem (where there is one less player and one less good) every player's MMS value has not decreased. Thus, so long as the reduced problem has an MMS allocation we will have an MMS allocation overall.

Algorithm 1 MMS Allocation for $m \leq n+4$.

while $\exists i \in N$ s.t. $\exists g \in G$ where $V_i(g) \geq \text{MMS}(i)$ **do**
  Give g to player i.
  $N \leftarrow N \setminus \{i\}$
  $G \leftarrow G \setminus \{g\}$
// Note that at this point $|N| \leq 4$.
For convenience, relabel the players so the first is 1, the second is 2, and so on.
if $|N| = 1$ **then**
  Give all of G to player 1.
else if $|N| = 2$ **then**
  Let player 1 produce an MMS partition (for 2 players).
  Let player 2 choose a single subset of the partition.
  Give player 1 the remaining subset.
else if $|N| = 3$ **then**
  // $|G| \in \{6, 7\}$, so every MMS partition of any player has two subsets of size 2.
  Let X_1, X_2, X_3 be an MMS partition of player 1 where $|X_1| = |X_2| = 2$.
  if $\exists i \in \{2, 3\}$ s.t. $V_i(X_1) < \text{MMS}(i)$ or $V_i(X_2) < \text{MMS}(i)$ **then**
    WLOG assume $V_2(X_1) < \text{MMS}(2)$.
    Let Y_1, Y_2, Y_3 be an MMS partition of player 2.
    WLOG assume $X_1 \subseteq Y_1 \cup Y_2$.
    $Z := (Y_1 \cup Y_2) \setminus X_1$.
    Let player 3 choose one of X_1, Y_3, Z.
    if Player 3 chooses X_1 **then**
      // As $V_2(X_1) < \text{MMS}(2)$ at least one of $V_2(X_2), V_2(X_3) \geq \text{MMS}(2)$.
      Give player 2 one of X_2 and X_3 such that he achieves his MMS.
      Give player 1 the other subset.
    else if Player 3 chooses Y_3 **then**
      Give X_1 to player 1.
      // Note that $V_2(Y_1), V_2(Y_2) \geq \text{MMS}(2)$ and $V_2(X_1) < \text{MMS}(2)$.
      // Thus, $V_2(Z) = V_2(Y_1) + V_2(Y_2) - V_2(X_1) > \text{MMS}(2)$.
      Give Z to player 2.
    else if Player 3 chooses Z **then**
      Give X_1 to player 1.
      Give Y_3 to player 2.
  else
    Give X_3 to player 1.
    Give X_1 to player 2.
    Give X_2 to player 3.
else if $|N| = 4$ **then**
  // $|G| = 8$, so every MMS partition of any player has only subsets of size 2.
  Let the players choose a single good one at a time in the order: $1, 2, 3, 4, 4, 3, 2, 1$.

Now note that as long as $m < 2n$ every MMS partition of any player must include a subset with a single good. We can therefore utilize the observation repeatedly until there are at most 4 players left. In the event that there are at most 2 players left, we know this is easily handled and so only the cases where there are 3 or 4 players remaining are of interest.

In the more complex outcome where there are 3 players we essentially have an intricate case analysis that is best understood via the fully explicit treatment given in Algorithm 1. We therefore only consider the case where there are 4 players left here. In such an outcome exactly 8 goods remain and in every MMS partition of any player all subsets of the partition are of size 2 (as otherwise some player achieves his MMS value with a single good).

We claim then that for any player i if the goods $g_1, ..., g_8$ were sorted by value in that order (i.e. $V_i(g_j) \geq V_i(g_k)$ for all $j \leq k$) then $\{g_1, g_8\}, \{g_2, g_7\}, \{g_3, g_6\}, \{g_4, g_5\}$ is an MMS partition for i (i.e. the partition where g_j paired with g_{9-j}). Suppose this were not true, then let $j \in \arg\min_{k \leq 4}(V_i(g_k) + V_i(g_{9-k}))$. Now let S_1, S_2, S_3, S_4 be any MMS partition for i and consider the good g_j is paired with, say g_k. We know that that $k < 9 - j$ as otherwise we would have

$$\begin{aligned}\text{MMS}(i) = \min_k V_i(S_k) &\leq V_i(g_j) + V_i(g_k) \\ &\leq V_i(g_j) + V_i(g_{9-j}) = \min_{k \leq 4}(V_i(g_k) + V_i(g_{9-k})).\end{aligned}$$

Now consider the j goods $g_{9-j}, g_{10-j}, ..., g_8$. As we require $\text{MMS}(i) = \min_k V_i(S_k) > V_i(g_j) + V_i(g_{9-j})$ we must have that the goods they are paired with in the MMS partition chosen have value greater than $V_i(g_j)$. Unfortunately, there are at most $j - 1$ such goods — a clear contradiction.

Thus, if we allow players to choose one good at a time, we find that so long as a player gets to make the j^{th} and $(9-j)^{th}$ choice, he will have his MMS value.

As mentioned above, see Algorithm 1 for the complete approach. ■