



# Optimal Bounds on the Price of Fairness for Indivisible Goods

Siddharth Barman<sup>1(✉)</sup>, Umang Bhaskar<sup>2</sup>, and Nisarg Shah<sup>3</sup>

<sup>1</sup> Indian Institute of Science, Bangalore, India  
barman@iisc.ac.in

<sup>2</sup> Tata Institute of Fundamental Research, Mumbai, India  
umang@tifr.res.in

<sup>3</sup> University of Toronto, Toronto, Canada  
nisarg@cs.toronto.edu

**Abstract.** In the allocation of resources to a set of agents, how do fairness guarantees impact social welfare? A quantitative measure of this impact is the price of fairness, which measures the worst-case loss of social welfare due to fairness constraints. While initially studied for divisible goods, recent work on the price of fairness also studies the setting of indivisible goods.

In this paper, we resolve the price of two well-studied fairness notions in the context of indivisible goods: envy-freeness up to one good (EF1) and approximate maximin share (MMS). For both EF1 and  $1/2$ -MMS we show, via different techniques, that the price of fairness is  $O(\sqrt{n})$ , where  $n$  is the number of agents. From previous work, it follows that these guarantees are tight. We, in fact, obtain the price-of-fairness results via efficient algorithms. For  $1/2$ -MMS our bound holds for additive valuations, whereas for EF1, it holds for the more general class of subadditive valuations. This resolves an open problem posed by Bei et al. (2019).

## 1 Introduction

What does it mean for an allocation of resources among a set of agents to be fair? The most compelling notion of fairness advocated in prior work is *envy-freeness* (EF) [15], which demands that no agent envy any other (i.e., value the resources allocated to any other agent more than those allocated to herself). When the resources to be allocated contain *indivisible* goods, guaranteeing envy-freeness is impossible.<sup>1</sup> Hence, researchers have sought relaxations such as *envy-freeness up to one good* (EF1) [8, 10], which states that it should be possible to eliminate any envy one agent has toward another by removal of at most one good from the latter’s allocation.

A different notion advocated in the context of indivisible goods is the *maximin share guarantee* (MMS) [8]. When there are  $n$  agents, the maximin share

---

<sup>1</sup> The canonical example is that of a single indivisible good and two agents—the agent that does not receive the good will inevitably envy the other.

of an agent  $a$  is defined as the maximum value  $a$  could obtain if she were to divide the goods into  $n$  bundles and then receive the least-valued (according to her) bundle. An MMS allocation is an allocation where each agent receives at least her maximin share. This generalizes the classical cut-and-choose protocol for dividing a cake between two agents, where one agent cuts the cake into two pieces, the other agent chooses a piece, and the first agent then receives the remaining piece. Even for agents with additive valuations over the goods,<sup>2</sup> MMS allocations may not exist [18, 23]. However, under additive valuations, an allocation where each agent receives at least  $3/4$ -th of her maximin share—i.e., a  $3/4$ -MMS allocation—always exists [16, 17].

While fairness is important, it is clearly not the only objective of interest. Another central criterion in resource allocation is the aggregate value of the agents from the resources they receive, or the *social welfare* of the allocation. The tradeoff between fairness and social welfare is quantitatively measured by the *price of fairness*, the supremum over all instances of the ratio between the maximum welfare of any allocation and the maximum welfare of any allocation satisfying the desired fairness notion. Intuitively, this quantity measures the factor by which welfare may be lost to achieve desired fairness.

Caragiannis et al. [9] initiated the study of the price of fairness in the canonical setting of cake-cutting, in which a heterogeneous *divisible* good is to be allocated (hence, envy-freeness can be guaranteed) and agents have additive valuations. They proved that the price of envy-freeness is between  $\Omega(\sqrt{n})$  and  $O(n)$ , where  $n$  is the number of agents. Later, Bertsimas et al. [6] closed the gap by proving that the correct bound is  $\Theta(\sqrt{n})$ , and that the matching upper bound can be achieved by maximizing the Nash welfare [6].

For *indivisible* goods and additive valuations, Bei et al. [5] studied the price of fairness for various notions of fairness. They showed that the price of envy-freeness up to one good (EF1) is between  $\Omega(\sqrt{n})$  and  $O(n)$ . One might immediately wonder if maximizing the Nash welfare (which is known to satisfy EF1 in the case of indivisible goods and additive valuations [10]) can be used to derive a matching  $O(\sqrt{n})$  upper bound, as in cake-cutting. Unfortunately, Bei et al. also showed that maximizing the Nash welfare can result in an  $\Omega(n)$ -factor loss with respect to social welfare, and posed settling the price of EF1 as a significant open question. For approximate MMS allocations (since exact MMS allocations may not exist), the price of fairness has not been studied earlier, though a lower bound of  $\Omega(\sqrt{n})$  can be obtained from previous constructions [9].

**Our Contributions.** We consider the allocation of  $m$  indivisible goods among  $n$  agents. An allocation  $\mathcal{A} = (A_1, \dots, A_n)$  partitions the goods into  $n$  bundles, where  $A_i$  is the bundle assigned to agent  $i$ . The preferences of each agent  $i$  are specified via a valuation function  $v_i$ , which assigns a non-negative value to every subset of goods. Then,  $v_i(A_i)$  is the value to agent  $i$  under  $\mathcal{A}$  and  $\sum_{i=1}^n v_i(A_i)$  is the social welfare of  $\mathcal{A}$ . The price of a fairness notion, as defined above, is the

---

<sup>2</sup> A valuation is additive iff the value of a bundle of goods is equal to the sum of the values of the individual goods in the bundle.

supremum of the ratio between the maximum social welfare obtainable, and the maximum social welfare subject to the fairness notion.

Our main contribution is to comprehensively settle the price of EF1 and  $1/2$ -MMS fairness notions. First, we show that the price of EF1 is  $O(\sqrt{n})$ . This matches the  $\Omega(\sqrt{n})$  lower bound due to Bei et al. [5]. The lower bound is for additive valuations, whereas our upper bound holds for the more general class of subadditive valuations. Hence, our work settles this open question for all valuation classes between additive and subadditive. Computationally, we obtain this upper bound via a polynomial-time algorithm. For subadditive valuations, given the absence of succinct representation, we assume access to a *demand-query oracle*.<sup>3</sup> As a consequence of this result, we also settle the price of a weaker fairness notion—proportionality up to one good (Prop1)—as  $\Theta(\sqrt{n})$  for additive valuations.

For the  $1/2$ -MMS fairness notion and additive valuations, we similarly establish, via a different algorithm, that the price of fairness is  $\Theta(\sqrt{n})$ . We show that for a fixed  $\varepsilon > 0$ , a  $(1/2 - \varepsilon)$ -MMS allocation with welfare within  $O(\sqrt{n})$  factor of the optimal can be computed in polynomial time.

**Related Work.** For resource allocation, the price of fairness was first studied by Caragiannis et al. [9] and Bertsimas et al. [6]. Caragiannis et al. [9] left open the question of the price of envy-freeness in cake-cutting, later settled by Bertsimas et al. [6] as  $\Theta(\sqrt{n})$ . Bertsimas et al. also showed that the price of proportionality is  $\Theta(\sqrt{n})$  and the price of equitability is  $\Theta(n)$ . In addition, they extended their analysis to the case where agents *dislike* the cake (i.e., a divisible *chore* is being allocated), and the case with indivisible goods or chores. However, in case of indivisible goods, notions such as envy-freeness, proportionality, and equitability cannot be guaranteed. The analysis of Caragiannis et al. [9] simply excluded instances which do not admit allocations satisfying these criteria.

Bei et al. [5] instead focused on notions that can be guaranteed with indivisible goods and chores, such as envy-freeness up to one good (EF1). While they did not settle the price of EF1 (which is the focus of our work), they showed that the price of popular allocation rules such as the maximum Nash welfare rule [10], the egalitarian rule [24], and the leximin rule [7, 19] is  $\Theta(n)$ . They also considered allocations that are balanced, i.e., give all agents an approximately equal number of goods, and settled the price of this guarantee as  $\Theta(\sqrt{n})$ . To the best of our knowledge, our work is the first to consider the price of approximate maximin share guarantee.

## 2 Preliminaries

For  $k \in \mathbb{N}$ , define  $[k] := \{1, \dots, k\}$ . We study discrete fair division problems, wherein a set  $[m]$  of indivisible goods need to be partitioned in a fair manner among a set  $[n]$  of agents.

<sup>3</sup> Section 2 provides a formal description of the valuation classes and the query models.

**Agent Valuations.** The cardinal preference of each agent  $i \in [n]$  (over the goods) is specified via the valuation function  $v_i : 2^{[m]} \mapsto \mathbb{R}_{\geq 0}$ , where  $v_i(S) \in \mathbb{R}_{\geq 0}$  is the value that agent  $i$  has for the subset of goods  $S \subseteq [m]$ . We assume valuations are normalized ( $v_i(\emptyset) = 0$ ), nonnegative ( $v_i(S) \geq 0$  for all  $S \subseteq [m]$ ), and monotone ( $v_i(A) \leq v_i(B)$  for  $A \subseteq B \subseteq [m]$ ). A fair-division instance is given by a tuple  $\langle [n], [m], \{v_i\}_{i \in [n]} \rangle$ . We primarily consider two valuation classes:

- Additive:  $v_i(S) = \sum_{g \in S} v_i(g)$  for each agent  $i \in [n]$  and subset of goods  $S \subseteq [m]$ . Here,  $v_i(g)$  denotes the value that agent  $i$  has for good  $g \in [m]$ .
- Subadditive:  $v_i(S \cup T) \leq v_i(S) + v_i(T)$  for each agent  $i \in [n]$  and subsets  $S, T \subseteq [m]$ . Note that the family of subadditive valuations encompasses additive valuations.

**Oracle Access.** Since describing subadditive valuations may require size exponential in the number of goods, to design efficient algorithms, we assume oracle access to the valuation functions. The literature focuses on two query models.

- Value queries: Given an agent  $i \in [n]$  and a subset of goods  $S \subseteq [m]$ , the oracle returns  $v_i(S)$ .
- Demand queries: Given an agent  $i \in [n]$  and a price  $p_g \in \mathbb{R}_{\geq 0}$  for each good  $g \in [m]$ , the oracle returns a “profit-maximizing” set  $S^* \in \arg \max_{S \subseteq [m]} v_i(S) - \sum_{g \in S} p_g$ .

Demand queries are strictly more powerful than value queries [22, Section 11.5].

**Allocations.** Write  $\Pi_n([m])$  to denote the set of all  $n$ -partitions of the set of goods  $[m]$ . An allocation  $\mathcal{A} = (A_1, A_2, \dots, A_n) \in \Pi_n([m])$  corresponds to an  $n$ -partition wherein the subset  $A_i \subseteq [m]$  is assigned to agent  $i \in [n]$ ; such a subset is called a bundle. The term partial allocation denotes an  $n$ -partition  $(P_1, P_2, \dots, P_n) \in \Pi_n(S)$  of a subset of goods  $S \subseteq [m]$ . Here, as before, subset  $P_i$  is assigned to agent  $i \in [n]$ .

**Fairness.** The notions of fairness considered in this work are defined next.

**Definition 1 (Prop1).** *An allocation  $\mathcal{A}$  is called proportional up to one good (Prop1) if for each agent  $i \in [n]$ , there exists a good  $g \in [m]$ , such that  $v_i(A_i \cup \{g\}) \geq v_i([m])/n$ .*

**Definition 2 (EF1).** *An allocation  $\mathcal{A}$  is called envy-free up to one good (EF1) iff for all agents  $i, j \in [n]$  with  $A_j \neq \emptyset$ , there exists a good  $g \in A_j$  such that  $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ .*

It is easy to check that EF1 implies Prop1 for additive valuations. In a fair-division instance  $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$ , the maximin share of agent  $i \in [n]$  is defined as

$$\text{MMS}_i := \max_{(P_1, \dots, P_n) \in \Pi_n([m])} \min_{j \in [n]} v_i(P_j)$$

**Definition 3** ( $\alpha$ -MMS). For  $\alpha \in [0, 1]$ , an allocation  $\mathcal{A}$  is called  $\alpha$ -approximate maximin share fair ( $\alpha$ -MMS) if  $v_i(A_i) \geq \alpha \cdot \text{MMS}_i$  for each agent  $i \in [n]$ .

**Social Welfare.** The social welfare of an allocation  $\mathcal{A}$ , denoted  $\text{SW}(\mathcal{A})$ , is defined as the sum of the values that  $\mathcal{A}$  generates among the agents:  $\text{SW}(\mathcal{A}) = \sum_{i=1}^n v_i(A_i)$ . We will use  $\mathcal{W}^* = (W_1^*, W_2^*, \dots, W_n^*)$  to denote a social welfare maximizing allocation, i.e.,  $\mathcal{W}^* \in \arg \max_{\mathcal{A} \in \Pi_n([m])} \text{SW}(\mathcal{A})$ , and  $\text{OPT}$  to denote the optimal social welfare,  $\text{OPT} = \text{SW}(\mathcal{W}^*)$ .

**Price of Fairness.** Given a fairness property  $X$ , the price of  $X$  is the supremum, over all fair division instances, of the ratio between the maximum social welfare of any allocation and the maximum social welfare of any allocation satisfying property  $X$ .

**Scaling.** To ensure that valuations are on the same scale, much of the literature on fair division assumes that agents' valuations are *scaled*, i.e.,  $v_i([m]) = 1$  for all  $i \in [n]$ . Noting that unscaled valuations are common in other areas of social choice (e.g., [13]), we consider both scaled and unscaled valuations.

All missing proofs appear in the full version of the paper [4].

### 3 Price of Envy-Freeness Up to One Good (EF1)

We begin by studying the price of fairness for EF1 allocations for agents with subadditive valuations. For scaled additive valuations, Bei et al. [5] show that the price of EF1 is between  $\Omega(\sqrt{n})$  and  $O(n)$ . We tighten their upper bound to  $O(\sqrt{n})$  (thus matching their lower bound) even when the valuations are subadditive. For unscaled valuations, we show that the bound is  $\Theta(n)$ . Our main result in this section is as follows.

**Theorem 1.** *theorem The price of EF1 is  $O(\sqrt{n})$  for scaled subadditive valuations and  $O(n)$  for unscaled subadditive valuations. Both bounds are tight even when the valuations are additive.*

We begin by proving the upper bounds for EF1, using Algorithms 1 and 2.

#### 3.1 An Absolute Welfare Guarantee

First, we show (via Algorithm 1) that when agents have subadditive valuations  $\{v_i\}_{i \in [n]}$  (not necessarily scaled), there always exists an EF1 allocation  $\mathcal{A}$  with social welfare  $\text{SW}(\mathcal{A}) \geq \frac{1}{2n} \sum_{i=1}^n v_i([m])$ .

This absolute welfare guarantee has two implications. First, since  $\sum_{i=1}^n v_i([m])$  is a trivial upper bound on the optimal social welfare,  $\text{OPT}$ , the result establishes an  $O(n)$  upper bound on the price of EF1. For unscaled valuations, this is exactly the bound we need. For scaled valuations, we need to improve this to  $O(\sqrt{n})$ . Since  $\sum_{i=1}^n v_i([m]) = n$  under scaled valuations, the result gives  $\text{SW}(\mathcal{A}) \geq 1/2$ . Hence, if  $\text{OPT} = O(\sqrt{n})$ , then we have the desired  $O(\sqrt{n})$  upper bound. We analyse the case when  $\text{OPT} = \Omega(\sqrt{n})$  in Sect. 3.2.

---

**Algorithm 1.** ALG-EF1-ABS
 

---

**Input:** Fair-division instance  $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$  with value-query oracle access to the subadditive valuations  $v_i$ s.

**Output:** An EF1 allocation  $\mathcal{B}$  with social welfare  $\text{SW}(\mathcal{B}) \geq \frac{1}{2n} \sum_{i \in [n]} v_i([m])$ .

- 1: Consider the weighted bipartite graph  $G = ([n] \cup [m], [n] \times [m])$  with weight of each edge  $(i, g) \in [n] \times [m]$  set as  $v_i(g)$ . Let  $\pi$  be a maximum-weight matching in  $G$  that matches all nodes in  $[n]$ .
  - 2: Construct the partial allocation  $\mathcal{B}'$  such that  $B'_i = \{\pi(i)\}$  for each  $i \in [n]$ . {Note that  $\mathcal{B}'$  is trivially EF1 because each agent is assigned a single good.}
  - 3: Use the algorithm of Lipton et al. [20] to extend the partial EF1 allocation  $\mathcal{B}'$  into a complete EF1 allocation  $\mathcal{B}$  such that  $v_i(B_i) \geq v_i(B'_i)$  for each agent  $i \in [n]$ .
  - 4: **return** Allocation  $\mathcal{B}$
- 

**Lemma 1.** *Let  $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$  be a fair-division instance in which agent valuations are subadditive. Then, given value-query oracle access to the valuations, ALG-EF1-ABS (Algorithm 1) efficiently computes an EF1 allocation  $\mathcal{B}$  with social welfare  $\text{SW}(\mathcal{B}) \geq \frac{1}{2n} \sum_{i \in [n]} v_i([m])$ .*

### 3.2 The Case of High Optimal Welfare

As noted in Sect. 3.1, Lemma 1 allows us to focus on fair division instances with scaled subadditive valuations in which the optimal social welfare is  $\text{OPT} = \Omega(\sqrt{n})$ . This allows us to sacrifice  $O(\sqrt{n})$  of the welfare in  $\text{OPT}$ , obtain  $O(\sqrt{n})$  approximation to the *remaining* welfare through an EF1 allocation, and yet achieve  $O(\sqrt{n})$  approximation to  $\text{OPT}$ . In particular, we present an algorithm that efficiently finds an EF1 allocation  $\mathcal{A}$  with social welfare  $\text{SW}(\mathcal{A}) \geq \frac{\text{OPT} - 2\sqrt{n}}{12\sqrt{n}}$ .

Ideally, we would like to use as reference an allocation  $\mathcal{W}^*$  with the optimal social welfare  $\text{OPT}$ . However, for subadditive valuations, computing such an allocation is NP-hard under both value queries and demand queries [12]. For values queries,  $\Theta(\sqrt{m})$  is the best possible approximation to  $\text{OPT}$  with a polynomial number of queries [12, 21]. Instead, we turn to demand queries, for which a 2-approximation algorithm is given by Feige [14]. In particular, this algorithm efficiently computes an allocation  $\mathcal{W}$  with social welfare at least  $\frac{1}{2} \text{OPT}$ .<sup>4</sup> Outside of the black-box use of this algorithm, the rest of our algorithm uses value queries (which are a special case of demand queries). We emphasize that our use of value or demand queries is only for computation. Our main result—the price of EF1—is existential and independent of any query model.

Starting with the high-welfare allocation  $\mathcal{W}$  returned by Feige’s algorithm, our algorithm works as follows. It first indexes the  $m$  goods as  $g_1, g_2, \dots, g_m$  such that the goods in each  $W_i$  receive consecutive indices.<sup>5</sup> Alternatively,

<sup>4</sup> This allocation serves as a reference in our algorithm, and may not be EF1 itself.

<sup>5</sup> For example, we can index the goods such that  $W_i = \left\{ g_k : 1 + \sum_{j=1}^{i-1} |W_j| \leq k \leq \sum_{j=1}^i |W_j| \right\}$  for each agent  $i \in [n]$ .

consider a line graph  $L = ([m], E)$  over the set of goods with edges  $E = \{(g_k, g_{k+1}) : k \in [m - 1]\}$ . Then, each  $W_i$  induces a connected subgraph of  $L$ .

**Definition 4.** Let  $L = ([m], E)$  be a line graph over the goods. We say that  $S \subseteq [m]$  is a connected bundle in  $L$  if  $S$  induces a connected subgraph of  $L$ . Given a partial allocation  $\mathcal{P}$ , define  $\mathcal{U}(\mathcal{P})$  as the set of connected components of  $L$  that remain after removing the allocated goods  $\cup_{i \in [n]} P_i$ . We refer to  $U \in \mathcal{U}(\mathcal{P})$  as an unassigned connected bundle.

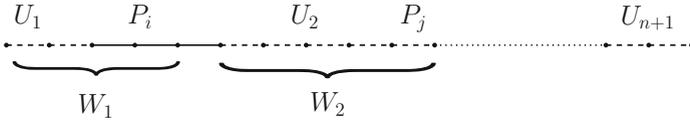


Fig. 1. Line graph over the goods and unassigned components.

In this terminology, each  $W_i$  is a connected bundle in  $L$  (see Fig. 1). ALG-EF1-HIGH builds a partial allocation  $\mathcal{P}$  by giving each agent  $i$  her most valuable good from  $W_i$ . This allocation trivially satisfies two properties: it is EF1, and each  $P_i$  is a connected bundle in  $L$ . The algorithm then iteratively updates  $\mathcal{P}$  to improve its social welfare while maintaining both these properties. This iterative process is inspired by a similar algorithm for (divisible) cake-cutting (Algorithm 1, Arunachaleswaran et al. [3]).

In particular, at every iteration, our algorithm computes the set of unassigned connected bundles  $\mathcal{U}(\mathcal{P})$ ; note that removing  $n$  connected bundles from  $\mathcal{P}$  can create at most  $n + 1$  unassigned connected bundles, as shown in Fig. 1. If there is an unassigned connected bundle  $U \in \mathcal{U}(\mathcal{P})$  that some agent envies, then the algorithm finds an *inclusion-wise minimal* subset of  $U$  that is envied and allocates it to an envious agent. While this preserves exact envy-freeness in the cake-cutting setting of Arunachaleswaran et al. [3], in our setting we argue that inclusion-wise minimality preserves the EF1 property of  $\mathcal{P}$ . We also argue that this iterative process terminates at a partial allocation  $\mathcal{P}$  that satisfies the desired social welfare guarantee. Finally, we use the algorithm of Lipton et al. [20] (see also [2]) to extend this partial EF1 allocation into a complete EF1 allocation without losing social welfare. The detailed algorithm is presented as Algorithm 2.

We start by proving some relevant properties of ALG-EF1-HIGH.

**Lemma 2.** When ALG-EF1-HIGH is run on an instance  $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$  with subadditive valuations, the following hold regarding the partial allocation  $\mathcal{P}^t$  constructed after  $t$  iterations of the while loop.

1. If  $t \geq 1$ , then for each agent  $i \in [n]$ , either  $P_i^t = P_i^{t-1}$  or  $v_i(P_i^t) > v_i(P_i^{t-1})$ .
2. For each agent  $i \in [n]$ ,  $P_i$  is a connected bundle under the line graph  $L$  constructed in Line 6.
3.  $\mathcal{P}^t$  is EF1.

---

**Algorithm 2.** ALG-EF1-HIGH
 

---

**Input:** A fair division instance  $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$  with demand-query oracle access to the subadditive valuations  $v_i$ s.

**Output:** An EF1 allocation  $\mathcal{A}$  with social welfare  $\text{SW}(\mathcal{A}) \geq \frac{\text{OPT}}{12\sqrt{n}} - \frac{1}{6}$ .

- 1: Use Feige's algorithm [14] to compute an allocation  $\mathcal{W}$  with social welfare  $\text{SW}(\mathcal{W}) \geq 1/2 \cdot \text{OPT}$
  - 2: Re-index the  $m$  goods as  $g_1, g_2, \dots, g_m$  such that in the line graph  $L$  over the goods containing edges  $(g_k, g_{k+1})$  for  $k \in [m-1]$ , each  $W_i$  forms a connected subgraph
  - 3: Initialize  $t = 0$
  - 4: For each agent  $i \in [n]$  with  $W_i \neq \emptyset$ , pick  $g_i^* \in \arg \max_{g \in W_i} v_i(g)$  and set  $P_i^t = \{g_i^*\}$
  - 5: For each agent  $i \in [n]$  with  $W_i = \emptyset$ , set  $P_i^t = \emptyset$   
 {For a partial allocation  $\mathcal{P}^t$ , let  $\mathcal{U}(\mathcal{P}^t)$  denote the collection of connected components in the line graph  $L$  that remain after the removal of  $\cup_{i \in [n]} P_i^t$ }
  - 6: **while**  $\exists$  agent  $i \in [n]$  and connected component  $U \in \mathcal{U}(\mathcal{P}^t)$  with  $v_i(P_i^t) < v_i(U)$   
**do**
  - 7:   Let  $U$  consist of goods  $\{g_a, g_{a+1}, \dots, g_b\}$
  - 8:   Let  $c \in [a, b]$  be the smallest index such that  $v_k(P_k^t) < v_k(\{g_a, g_{a+1}, \dots, g_c\})$  for some  $k \in [n]$   
 {The choice of  $c$  ensures that no agent  $k$  values  $\{g_a, \dots, g_{c-1}\}$  more than her current bundle, i.e.,  $v_k(P_k^t) \geq v_k(\{g_a, \dots, g_{c-1}\})$  for every agent  $k$ .}
  - 9:   Pick an arbitrary agent  $k$  for which  $v_k(P_k^t) < v_k(\{g_a, g_{a+1}, \dots, g_c\})$
  - 10:   Set  $P_k^{t+1} = \{g_a, \dots, g_c\}$  and  $P_j^{t+1} = P_j^t$  for all  $j \neq k$
  - 11:   Update  $t \leftarrow t + 1$
  - 12: **end while**  
 {At this point,  $v_i(P_i^t) \geq v_i(U)$  for all  $i \in [n]$  and  $U \in \mathcal{U}(\mathcal{P}^t)$ . We will show that  $\mathcal{P}^t$  is a partial EF1 allocation.}
  - 13: Use the algorithm of Lipton et al. [20] to extend the partial EF1 allocation  $\mathcal{P}^t$  into a complete EF1 allocation  $\mathcal{A}$  such that  $v_i(A_i) \geq v_i(P_i^t)$  for each  $i \in [n]$
  - 14: **return** Allocation  $\mathcal{A}$
- 

**Lemma 3.** *When ALG-EF1-HIGH is run on an instance  $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$  with subadditive valuations, the following hold.*

1. *The while loop terminates after  $T = O(nm^2)$  iterations.*
2. *The partial allocation  $\mathcal{P}^T$  constructed at the end of the while loop satisfies  $v_i(P_i^T) \geq v_i(U)$  for every agent  $i \in [n]$  and unassigned connected bundle  $U \in \mathcal{U}(\mathcal{P}^T)$ .*

Using Lemmas 2 and 3, we can derive the key technical result of this section.

**Lemma 4.** *Given a fair division instance  $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$  with scaled subadditive valuations, ALG-EF1-HIGH terminates in polynomial time and returns an EF1 allocation  $\mathcal{A}$  satisfying*

$$\text{SW}(\mathcal{A}) \geq \frac{\text{OPT}}{12\sqrt{n}} - \frac{1}{6},$$

where OPT is the optimal social welfare achievable in instance  $\mathcal{I}$ .

We now prove the upper bounds in Theorem 1, using Lemmas 1 and 4. As mentioned earlier, the lower bound for scaled valuations is already obtained by Bei et al. [5]. The lower bound for unscaled valuations is shown in the full version [4].

*Proof (of Theorem 1).* For unscaled valuations, Lemma 1 shows that there exists an EF1 allocation  $\mathcal{B}$  with  $\text{SW}(\mathcal{B}) \geq \frac{1}{2n} \sum_{i \in [n]} v_i([m]) \geq \frac{1}{2n} \text{OPT}$ , which yields the desired  $O(n)$  bound.

For an instance  $\mathcal{I}$  with scaled valuations, we consider two cases: either  $\text{OPT} \leq 8\sqrt{n}$  or  $\text{OPT} > 8\sqrt{n}$ .

**Case 1:** Suppose  $\text{OPT} \leq 8\sqrt{n}$ . Because valuations are scaled (i.e.,  $v_i([m]) = 1$  for each agent  $i \in [n]$ ), Lemma 1 implies that the EF1 allocation  $\mathcal{B}$  returned by ALG-EF1-ABS satisfies  $\text{SW}(\mathcal{B}) \geq \frac{1}{2n} \sum_{i \in [n]} v_i([m]) = \frac{1}{2}$ . Hence,  $\text{SW}(\mathcal{B}) \geq \frac{1}{16\sqrt{n}} \text{OPT}$ .

**Case 2:** Suppose  $\text{OPT} > 8\sqrt{n}$ . Then,  $\frac{1}{6} < \frac{\text{OPT}}{48\sqrt{n}}$ . Now, Lemma 4 implies that the EF1 allocation  $\mathcal{A}$  returned by ALG-EF1-HIGH satisfies

$$\text{SW}(\mathcal{A}) \geq \frac{\text{OPT}}{12\sqrt{n}} - \frac{1}{6} \geq \frac{\text{OPT}}{12\sqrt{n}} - \frac{\text{OPT}}{48\sqrt{n}} = \frac{\text{OPT}}{16\sqrt{n}}.$$

Thus, in either case, there exists an EF1 allocation with social welfare at least  $\frac{1}{16\sqrt{n}} \text{OPT}$ , which yields the desired  $O(\sqrt{n})$  bound. To obtain the required EF1 allocation, we simply run both algorithms and take the allocation with higher social welfare.

Before concluding this section, we recall that for additive valuations, EF1 implies proportionality up to one good (Prop1). Hence, our upper bounds of  $O(\sqrt{n})$  and  $O(n)$  on the price of EF1 under scaled and unscaled additive valuations, respectively, carry over to the price of Prop1 as well. These bounds are tight: Our lower bound construction from the proof of Theorem 1 for unscaled valuations and the construction due to Bei et al. [5] for scaled valuations can be modified slightly for Prop1 as well.

**Corollary 1.** *The price of Prop1 is  $\Theta(\sqrt{n})$  for scaled additive valuations and  $\Theta(n)$  for unscaled additive valuations.*

## 4 Price of $1/2$ -Approximate Maximin Share Guarantee

We now study the price of approximate MMS for additive valuations. Our main result settles the price of  $1/2$ -MMS for scaled and unscaled additive valuations.

**Theorem 2.** *The price of  $1/2$ -MMS is  $\Theta(\sqrt{n})$  for scaled additive valuations and  $\Theta(n)$  for unscaled additive valuations.*

We sketch the proof of the upper bound, based on Algorithms 3 and 4. Recall that  $\mathcal{W}^* = (W_1^*, \dots, W_n^*) \in \arg \max_{\mathcal{A} \in \Pi_n([m])} \text{SW}(\mathcal{A})$  is a social welfare maximizing allocation, and  $\text{OPT} = \text{SW}(\mathcal{W}^*)$  is the maximum social welfare.

### 4.1 An Absolute Welfare Guarantee

First, we show that when agents have additive valuations  $\{v_i\}_{i \in [n]}$  (not necessarily scaled), there always exists a  $1/2$ -MMS allocation  $\mathcal{A}$  with social welfare  $\text{SW}(\mathcal{A}) \geq \frac{1}{3n} \sum_{i=1}^n v_i([m])$ . Algorithm 3 computes such an allocation efficiently.

As for EF1 allocations, there are two implications of this. First, for both scaled and unscaled additive valuations, this establishes that the price of  $1/2$ -MMS is  $O(n)$ . For unscaled valuations, this is the bound we seek. For scaled valuations, since  $v_i([m]) = 1$  for each agent  $i \in [n]$ , so  $\text{SW}(\mathcal{A}) \geq 1/3 = \Omega(1)$ . Thus, if  $\text{OPT} = O(\sqrt{n})$ , this gives the desired upper bound of  $O(\sqrt{n})$ . In Sect. 4.2, we thus limit our attention to instances with  $\text{OPT} = \Omega(\sqrt{n})$ .

---

**Algorithm 3.** ALG-MMS-ABS

---

**Input:** A fair division instance  $\mathcal{I} = \langle [n], [m], \{v_i\}_i \rangle$  with additive valuations  $\{v_i\}_{i \in [n]}$ .  
**Output:** A  $1/2$ -MMS allocation  $\mathcal{B}$  with  $\text{SW}(\mathcal{B}) \geq \frac{1}{3n} \sum_{i=1}^n v_i([m])$ .

- 1: Initialize set of agents  $A = [n]$ , set of goods  $G = [m]$ , and bundles  $B_i = \emptyset$ , for all  $i \in [n]$
- 2: **while** there exists agent  $i \in A$  and good  $g \in G$  such that  $v_i(g) \geq \frac{1}{2|A|} v_i(G)$  **do**
- 3:   Set  $(i', g') \in \arg \max_{\{(i,g) \in A \times G : v_i(g) \geq \frac{1}{2|A|} v_i(G)\}} v_i(g)$
- 4:   Set  $B_{i'} = \{g'\}$  and update  $A \leftarrow A \setminus \{i'\}$  along with  $G \leftarrow G \setminus \{g'\}$
- 5: **end while**
- 6: Efficiently compute a Prop1 allocation  $(B_i)_{i \in A}$  of the fair division instance  $\langle A, G, \{v_i\}_{i \in A} \rangle$
- 7: **return** allocation  $\mathcal{B} = (B_1, B_2, \dots, B_n)$

---

Algorithm 3 is a refinement of the algorithm of Amanatidis et al. [1] for computing a  $1/2$ -MMS allocation: in Line 3, we use an arg max to break ties, whereas they use an arbitrary pair  $(i, g)$ . Since Amanatidis et al. [1] prove that their algorithm always returns a  $1/2$ -MMS allocation, regardless of any tie-breaking, it follows that our refinement, Algorithm 3, also always returns a  $1/2$ -MMS allocation. It remains to prove that it also provides the desired welfare guarantee, which is shown in the following lemma.

**Lemma 5.** *Given any fair division instance  $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$  with additive valuations, ALG-MMS-ABS (Algorithm 3) efficiently computes a  $1/2$ -MMS allocation  $\mathcal{B}$  with social welfare  $\text{SW}(\mathcal{B}) \geq \frac{1}{3n} \sum_{i=1}^n v_i([m])$ .*

### 4.2 The Case of High Optimal Welfare

As argued in Sect. 4.1, Lemma 5 allows us to restrict our attention to scaled additive valuations in which the optimal social welfare  $\text{OPT} = \Omega(\sqrt{n})$ . Similar to the case of EF1, we can now safely sacrifice  $O(\sqrt{n})$  welfare, and simply achieve

$O(\sqrt{n})$  approximation of the *remaining* welfare. This is achieved by ALG-MMS-HIGH (Algorithm 4).

However, ALG-MMS-HIGH requires knowledge of the maximin share  $MMS_i$  of each agent  $i$ . While computing this quantity is known to be strongly NP-hard, there exists a polynomial-time approximation scheme (PTAS) for it [25]. For a fixed  $\epsilon \in (0, 1)$ , this PTAS can compute an estimate  $Z_i \in [(1 - \epsilon)MMS_i, MMS_i]$  for each agent  $i$  in polynomial time. We pass these estimates as input to ALG-MMS-HIGH, which runs in polynomial time and yields a  $(\frac{1}{2} - \epsilon)$ -MMS allocation with the desired welfare guarantee.

We emphasize that the approximation here is solely for computational purposes. To derive our main existential result about the price of  $1/2$ -MMS, we can simply pass  $Z_i = MMS_i$  to ALG-MMS-HIGH, and it returns an exact  $1/2$ -MMS allocation with the desired welfare guarantee.

The intuition behind the algorithm is as follows. Our goal is to assign each agent  $i$  a bundle of value at least  $\frac{1}{2} \cdot Z_i \geq (\frac{1}{2} - \epsilon) \cdot MMS_i$ , and ensure that most (all but  $\sqrt{n}$ ) agents  $i$  achieve a value at least  $\Omega\left(\frac{1}{\sqrt{n}}v_i(W_i^*)\right)$ ; here,  $\mathcal{W} = (W_1^*, \dots, W_n^*)$  denotes a social welfare maximizing allocation. Therefore, throughout the algorithm, we keep track of two sets of agents,  $T$  (temporary) and  $P$  (permanent). Each agent  $i \in T \cup P$  must have received a bundle  $B_i$  worth  $v_i(B_i) \geq \frac{1}{2} \cdot Z_i$ , and agent  $i \in P$  further has  $v_i(B_i) \geq \frac{1}{3\sqrt{n}} \cdot v_i(W_i^*)$ . The algorithm ensures that an agent is never removed after being added to  $P \cup T$ . Specifically, once she is added to  $T$ , she can only be moved to  $P$ . Additionally, once an agent is included in the set  $P$ , her assignment is never updated and she remains in  $P$ .

Lines 5 and 9 handle easy cases, when an agent  $i$  either has  $MMS_i = 0$ , or can be added directly to  $P$  by giving her a single good from  $W_i^*$ . Line 13 addresses agents to whom giving a single good (not necessarily from  $W_i^*$ ) is sufficient to add them to  $P \cup T$ . These steps leverage the fact that the maximin share is maintained while assigning away singleton bundles.

Finally, Line 17 leverages an idea similar to what we utilized for EF1. This step slowly grows a bundle by iteratively adding goods ordered according to  $\mathcal{W}^*$ , and assigns the bundle as soon as its value for some agent is at least half of her maximin share. Line 24 plays a key role in bookkeeping, as we show in our proofs. The following lemma states the main result for Algorithm 4.

**Lemma 6.** *Given a fair division instance  $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$  with scaled additive valuations, and, for a fixed  $\epsilon \in [0, 1)$ , an estimate  $Z_i \in [(1 - \epsilon)MMS_i, MMS_i]$  for each agent  $i$ , ALG-MMS-HIGH (Algorithm 4) efficiently computes a  $(\frac{1}{2} - \epsilon)$ -MMS allocation  $\mathcal{B}$  with the property that  $3\sqrt{n} \cdot SW(\mathcal{B}) + 4\sqrt{n} \geq OPT$ ; here,  $OPT$  denotes the optimal social welfare in  $\mathcal{I}$ .*

Using the properties of ALG-MMS-ABS and ALG-MMS-HIGH from Sect. 4.1 and 4.2, we can directly prove the upper bounds on the price of  $1/2$ -MMS in Theorem 2. For the lower bounds, we use constructions that appeared in the work of Caragiannis et al. [11]. The complete proofs appear in the full version [4].

---

**Algorithm 4.** ALG-MMS-HIGH

---

**Input:** A fair division instance  $\mathcal{I} = \langle [m], [n], \{v_i\}_i \rangle$  with scaled additive valuations, and for a fixed  $\epsilon \in [0, 1)$ , an estimate  $Z_i \in [(1 - \epsilon)\text{MMS}_i, \text{MMS}_i]$  for each agent  $i$ .

**Output:** A  $(\frac{1}{2} - \epsilon)$ -MMS allocation  $\mathcal{B}$  with  $\text{SW}(\mathcal{B}) \geq \frac{1}{3\sqrt{n}} \text{OPT} - \frac{4}{3}$ .

- 1: Compute a social welfare maximizing allocation  $\mathcal{W}^* \in \arg \max_{\mathcal{A} \in \Pi_n([m])} \text{SW}(\mathcal{A})$
  - 2: Index the goods as  $g_1, \dots, g_m$  so that, for each  $i \in [n]$ , the goods in  $W_i^*$  receive consecutive indices
  - 3: Initialize  $\mathcal{B}$  with  $B_i = \emptyset$  for each  $i \in [n]$
  - 4: Initialize  $P = T = \emptyset$
  - 5: For each  $i \in [n]$  with  $\text{MMS}_i = 0$ , update  $P \leftarrow P \cup \{i\}$  if  $v_i(W_i^*) = 0$ , and  $T \leftarrow T \cup \{i\}$  otherwise. {MMS<sub>*i*</sub> = 0 iff agent *i* has positive value for less than *n* goods, which can be checked efficiently.}
  - 6: Let  $\Gamma_{\text{single}} = \left\{ i \in [n] : Z_i < \frac{2}{3\sqrt{n}} v_i(W_i^*) \text{ and } \exists g \in W_i^* \text{ s.t. } v_i(g) \geq \frac{1}{3\sqrt{n}} v_i(W_i^*) \right\}$
  - 7: **for** each  $i \in \Gamma_{\text{single}}$  **do**
  - 8: Pick  $g_i \in \arg \max_{g \in W_i^*} v_i(g)$ , and set  $B_i = \{g_i\}$
  - 9: Update  $P \leftarrow P \cup \{i\}$ , and if  $i \in T$ , update  $T \leftarrow T \setminus \{i\}$
  - 10: **end for**
  - 11: **while** there exists an agent  $a \in [n] \setminus (P \cup T)$  and a good  $h \in [m] \setminus \cup_{b \in [n]} B_b$  such that  $v_a(h) \geq \frac{1}{2} \cdot Z_a$  **do**
  - 12: Set  $B_a = \{h\}$ .
  - 13: If  $v_a(B_a) \geq \frac{1}{3\sqrt{n}} v_a(W_a^*)$ , update  $P \leftarrow P \cup \{a\}$ , else update  $T \leftarrow T \cup \{a\}$
  - 14: **end while**
  - 15: Let  $R \leftarrow [m] \setminus \cup_{a \in [n]} B_a$  be the set of remaining goods
  - 16: Initialize  $K = \emptyset$  and index  $t = 1$
  - 17: **for**  $t = 1, \dots, m$  **do**
  - 18: **if**  $g_t \notin R$  **then**
  - 19: **continue**
  - 20: **end if**
  - 21: Update  $K \leftarrow K \cup \{g_t\}$
  - 22: **if** there exists an agent  $i \in T$  such that  $g_t \in W_i^*$  and  $v_i(K) \geq \frac{1}{3\sqrt{n}} v_i(W_i^*)$  **then**
  - 23: Set  $(B_i, K) \leftarrow (K, B_i)$  {Swap  $B_i$  and  $K$ }
  - 24: Update  $P \leftarrow P \cup \{i\}$  and  $T \leftarrow T \setminus \{i\}$
  - 25: **end if**
  - 26: **if** there exists an agent  $a \in [n] \setminus (P \cup T)$  such that  $v_a(K) \geq \frac{1}{2} \cdot Z_a$  **then**
  - 27: Set  $B_a = K$  and update  $K = \emptyset$
  - 28: If  $v_a(B_a) \geq \frac{1}{3\sqrt{n}} v_a(W_a^*)$ , update  $P \leftarrow P \cup \{a\}$ , else update  $T \leftarrow T \cup \{a\}$
  - 29: **end if**
  - 30: **end for**
  - 31: Let  $X = [m] \setminus \cup_{a \in [n]} B_a$  be the set of unassigned goods. Assign each  $g \in X$  to agent  $i$  such that  $g \in W_i^*$ , i.e.,  $B_i \leftarrow B_i \cup (W_i^* \cap X)$ .
  - 32: **return** allocation  $(B_1, B_2, \dots, B_n)$
- 

**Acknowledgements.** SB gratefully acknowledges the support of a Ramanujan Fellowship (SERB - SB/S2/RJN-128/2015) and a Pratiksha Trust Young Investigator Award. UB's research is generously supported the Department of Atomic Energy, Government of India (project no. RTI4001), a Ramanujan Fellowship (SERB - SB/S2/RJN-

055/2015), and an Early Career Research Award (SERB - ECR/2018/002766). NS was partially supported by an NSERC Discovery Grant.

## References

1. Amanatidis, G., Markakis, E., Nikzad, A., Saberi, A.: Approximation algorithms for computing maximin share allocations. *ACM Trans. Algorithms (TALG)* **13**(4), 1–28 (2017)
2. Amanatidis, G., Markakis, E., Ntokos, A.: Multiple birds with one stone: beating  $1/2$  for EFX and GMMS via envy cycle elimination. In: *AAAI 2020*, pp. 1790–1797. AAAI Press (2020)
3. Arunachaleswaran, E.R., Barman, S., Kumar, R., Rathi, N.: Fair and efficient cake division with connected pieces. In: Caragiannis, I., Mirrokni, V., Nikolova, E. (eds.) *WINE 2019*. LNCS, vol. 11920, pp. 57–70. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-35389-6\\_5](https://doi.org/10.1007/978-3-030-35389-6_5)
4. Barman, S., Bhaskar, U., Shah, N.: Optimal bounds on the price of fairness for indivisible goods. *CoRR abs/2007.06242* (2020)
5. Bei, X., Lu, X., Manurangsi, P., Suksompong, W.: The price of fairness for indivisible goods. In: Kraus, S. (ed.) *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, 10–16 August 2019*, pp. 81–87 (2019)
6. Bertsimas, D., Farias, V.F., Trichakis, N.: The price of fairness. *Oper. Res.* **59**(1), 17–31 (2011)
7. Bogomolnaia, A., Moulin, H.: Random matching under dichotomous preferences. *Econometrica* **72**, 257–279 (2004)
8. Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. *J. Polit. Econ.* **119**(6), 1061–1103 (2011)
9. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M.: The efficiency of fair division. In: *Proceedings of the 5th Conference on Web and Internet Economics (WINE)*, pp. 475–482 (2009)
10. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum Nash welfare. *ACM Trans. Econ. Comput. (TEAC)* **7**(3), 1–32 (2019)
11. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M.: The efficiency of fair division. *Theory Comput. Syst.* **50**(4), 589–610 (2012)
12. Dobzinski, S., Nisan, N., Schapira, M.: Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.* **35**(1), 1–13 (2010)
13. Dutta, B., Peters, H., Sen, A.: Strategy-proof cardinal decision schemes. *Soc. Choice Welfare* **28**(1), 163–179 (2007)
14. Feige, U.: On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.* **39**(1), 122–142 (2009)
15. Foley, D.: Resource allocation and the public sector. *Yale Economics Essays*, vol. 7, pp. 45–98 (1967)
16. Garg, J., Taki, S.: An improved approximation algorithm for maximin shares. In: *EC 2020*, pp. 379–380. ACM (2020)
17. Ghodsi, M., Hajiaghayi, M., Seddighin, M., Seddighin, S., Yami, H.: Fair allocation of indivisible goods: improvements and generalizations. In: *Proceedings of the 2018 ACM Conference on Economics and Computation*, pp. 539–556 (2018)

18. Kurokawa, D., Procaccia, A.D., Wang, J.: When can the maximin share guarantee be guaranteed? In: Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI), pp. 523–529 (2016)
19. Kurokawa, D., Procaccia, A.D., Shah, N.: Leximin allocations in the real world. *ACM Trans. Econ. Comput. (TEAC)* **6**(3–4), 1–24 (2018)
20. Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: Proceedings of the 5th ACM Conference on Electronic Commerce, pp. 125–131 (2004)
21. Mirrokni, V., Schapira, M., Vondrák, J.: Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In: Proceedings of the 9th ACM Conference on Electronic Commerce, pp. 70–77 (2008)
22. Suzuki, M., Vetta, A.: How many freemasons are there? The consensus voting mechanism in metric spaces. In: Harks, T., Klimm, M. (eds.) SAGT 2020. LNCS, vol. 12283, pp. 322–336. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-57980-7\\_21](https://doi.org/10.1007/978-3-030-57980-7_21)
23. Procaccia, A.D., Wang, J.: Fair enough: guaranteeing approximate maximin shares. In: Proceedings of the 14th ACM Conference on Economics and Computation (EC), pp. 675–692 (2014)
24. Rawls, J.: *A Theory of Justice*. Harvard University Press, Cambridge (1971)
25. Woeginger, G.J.: A polynomial-time approximation scheme for maximizing the minimum machine completion time. *Oper. Res. Lett.* **20**(4), 149–154 (1997)