

Optimal Bounds on the Price of Fairness for Indivisible Goods

Siddharth Barman
Indian Institute of Science, India
barman@iisc.ac.in

Umang Bhaskar
Tata Institute of Fundamental Research, India
umang@tifr.res.in

Nisarg Shah
University of Toronto, Canada
nisarg@cs.toronto.edu

Abstract

In the allocation of resources to a set of agents, how do fairness guarantees impact the 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 for the allocation of indivisible goods: envy-freeness up to one good (EF1), and approximate maximin share (MMS). For both EF1 and $1/2$ -MMS guarantees, 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 our bounds are tight. Our bounds are obtained via efficient algorithms. For $1/2$ -MMS, our bound holds for additive valuations, whereas for EF1, our bound 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 to a set of agents to be fair? The most compelling notion of fairness advocated in prior work is *envy-freeness* (EF) [23], which demands that no agent envy another agent (i.e., value the resources allocated to the other agent more than those allocated to herself). When the resources to be allocated contain *indivisible* goods, guaranteeing envy-freeness is impossible.¹ Thus researchers have sought relaxations such as *envy-freeness up to one good* [14; 17], 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 relaxation of envy-freeness advocated in the context of indivisible goods is the *maximin share guarantee* (MMS) [14]. When there are n agents, the maximin share value of an agent is defined as the maximum value she could obtain if she were to divide the goods into n bundles, and then receive the least-valued bundle according to her. An MMS allocation is an allocation where each agent receives at least her maximin share value. 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,² MMS allocations may not exist [28; 35]. However, an allocation where each agent receives at least $3/4$ -th of her maximin share value — i.e., a $3/4$ -MMS allocation — always exists [24; 25].

¹The canonical example is that of a single indivisible good and two agents; the agent who does not receive the good will inevitably envy the other.

²A valuation function is additive if the value of a bundle of goods is the sum of the values of the individual goods in the bundle.

While fairness is important, it is clearly not the only objective of interest. A competing criterion 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 ratio of the maximum welfare of any allocation to the maximum welfare of any allocation satisfying the desired fairness notion, where the supremum is over all possible problem instances. Intuitively, this measures the factor by which welfare may be lost to achieve the desired fairness.

Caragiannis et al. [15] 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 (thus, 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. [11] 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 [11].

For *indivisible* goods and additive valuations, Bei et al. [10] 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 when allocating indivisible goods [17], can be used to derive a matching $O(\sqrt{n})$ upper bound, like in cake-cutting. Unfortunately, Bei et al. also showed that maximizing the Nash welfare results in a loss of welfare by $\Omega(n)$ factor in the worst case, 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 is not studied earlier, though a lower bound of $\Omega(\sqrt{n})$ can be obtained from previous constructions [15].

Our contributions. We first describe our model more formally. We consider the allocation of m indivisible goods to n agents. An allocation \mathcal{A} partitions the goods into n bundles, where A_i is the bundle assigned to agent i . The preferences of agent i are encoded by 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 $\sum_i v_i(A_i)$ is the social welfare of \mathcal{A} . The price of a fairness notion, as defined above, is the supremum ratio of the maximum social welfare to 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. [10]. 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. Our upper bound is obtained via a polynomial-time algorithm. For subadditive valuations, given the lack of succinct representation, we assume access to a *demand-query oracle*.³ 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 show, via a different algorithm, that the price of fairness is $\Theta(\sqrt{n})$. We show that for a fixed $\epsilon > 0$, a $(1/2 - \epsilon)$ -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. [15] and Bertsimas et al. [11]. While Caragiannis et al. [15] left open the question of the price of envy-freeness in cake-cutting, later settled by Bertsimas et al. [11] as $\Theta(\sqrt{n})$, they proved that the price of a weaker fairness notion — proportionality — is $\Theta(\sqrt{n})$, and the price of an incomparable fairness notion — equitability — is $\Theta(n)$. They also extended their analysis to the case where the 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. [15] simply excluded instances which do not admit allocations meeting these criteria.

Bei et al. [10] 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

³A formal description of valuation classes and query models is in Section 2.

work), they showed that the price of popular allocation rules such as the maximum Nash welfare rule [17], the egalitarian rule [36], and the leximin rule [12; 27] 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. Recent work has also analyzed the price of fairness in cake-cutting when the piece allocated to each agent is required to be contiguous [4]. Suksompong [37] extended this analysis to indivisible goods.

The price of fairness has also been analyzed in other paradigms at the intersection of computer science and economics, such as in kidney exchange [19], and is inspired from other similar worst-case measures such as the price of anarchy and of stability in game theory [26]. More distantly, the price of fairness has been extensively analyzed in machine learning (see, e.g., Mehrabi et al. [32]).

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.

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 all $A \subseteq B \subseteq [m]$). A fair-division instance is represented 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 all 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 agents and goods, to design efficient algorithms, we assume oracle access to the valuation functions. The literature focuses on two prominent query models.

- Value query: Given an agent $i \in [n]$ and a subset of goods $S \subseteq [m]$, the oracle returns $v_i(S)$.
- Demand query: 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 (e.g., [34, 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 $S \subseteq [m]$; 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 an 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 (α -MMS). For $\alpha \in [0, 1]$, an allocation \mathcal{A} is called α -approximate maximin share fair (α -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 the agents derive from the allocation: $\text{SW}(\mathcal{A}) = \sum_{i=1}^n v_i(A_i)$. 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 OPT to denote the optimal social welfare, i.e., $\text{OPT} = \text{SW}(\mathcal{W}^*)$.

Price of fairness. Given a fairness property X , the price of X is the supremum, over all fair division instances with n agents and m goods, of the ratio between the the maximum social welfare of any allocation to the maximum social welfare of any allocation satisfying property X .

Scaling. To ensure that agent valuations are on the same scale, much of the literature on fair division assumes that agent 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., [21]), we consider both scaled and unscaled valuations.

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. [10] 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. *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, which are established using Algorithms 1 and 2.

3.1 An Absolute Welfare Guarantee

First, we show 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 result has two implications. First, since $\sum_{i=1}^n v_i([m])$ is a trivial upper bound on the optimal social welfare 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 Section 3.2. Our absolute welfare guarantee is derived through Algorithm 1.

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 $\text{SW}(\mathcal{B}) \geq \frac{1}{2n} \sum_{i \in [n]} v_i([m])$.*

Proof. For each agent $i \in [n]$, sort the goods as g_{i_1}, \dots, g_{i_m} in a non-decreasing order of their value to agent i (with ties broken arbitrarily), i.e., so that $v_i(g_{i_1}) \geq \dots \geq v_i(g_{i_m})$, and let $G_i = \{g_{i_1}, \dots, g_{i_n}\}$ be the set of n most valuable goods to agent i .

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 π 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. [30] 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}
-

Consider a subgraph $G' = ([n] \cup [m], E')$ of the weighted bipartite graph G constructed in Line 1 of Algorithm 1, in which we only retain the edges $E' = \{(i, g) : i \in [n], g \in G_i\}$ (i.e., from each agent i to her n most valuable goods). Because G' is n -left-regular, it can be decomposed into n disjoint left-perfect matchings M_1, \dots, M_n ; this is an easy consequence of Hall's theorem [13]. Due to the pigeonhole principle, one of M_1, \dots, M_n must have weight at least $\frac{1}{n} \cdot \sum_{i \in [n], g \in G_i} v_i(g)$. Since this matching is also a left-perfect matching in G , and π constructed in Line 1 is a maximum-weight left-perfect matching in G , we have $\sum_{i \in [n]} v_i(\pi(i)) = \text{SW}(\mathcal{B}') \geq \frac{1}{n} \cdot \sum_{i \in [n], g \in G_i} v_i(g)$.

Note that \mathcal{B}' assigns a single good to each agent; hence, it is trivially EF1. It is known that for any monotonic valuations (thus, in particular, for subadditive valuations), the algorithm of Lipton et al. [30] can extend a partial EF1 allocation \mathcal{B}' into a complete EF1 allocation \mathcal{B} such that $v_i(B_i) \geq v_i(B'_i)$. Thus, for the EF1 allocation \mathcal{B} returned by the algorithm, we have

$$\text{SW}(\mathcal{B}) \geq \frac{1}{n} \sum_{i \in [n]} \sum_{g \in G_i} v_i(g). \quad (1)$$

Using the fact that \mathcal{B} is EF1, we get that, for each $i, j \in [n]$, there exists $S_j \subseteq B_j$ with $|S_j| \leq 1$ such that $v_i(B_i) \geq v_i(B_j \setminus S_j) \geq v_i(B_j) - v_i(S_j)$; here the last inequality follows from the subadditivity of v_i . Summing this inequality over $j \in [n]$, we get

$$n \cdot v_i(B_i) \geq \sum_{j \in [n]} v_i(B_j) - \sum_{j \in [n]} v_i(S_j) \geq v_i([m]) - \sum_{j \in [n]} v_i(S_j) \geq v_i([m]) - \sum_{g \in G_i} v_i(g),$$

where the penultimate inequality holds because B_j s form a partition of $[m]$ and v_i is subadditive, and the last inequality holds because the n sets S_j s are disjoint and have cardinality at most 1, and G_i is the set of the n most valuable goods to i . Summing over $i \in [n]$, we get

$$n \cdot \sum_{i \in [n]} v_i(B_i) = n \cdot \text{SW}(\mathcal{B}) \geq \sum_{i \in [n]} v_i([m]) - \sum_{i \in [n]} \sum_{g \in G_i} v_i(g).$$

Substituting Equation (1), we get the desired bound

$$n \cdot \text{SW}(\mathcal{B}) + n \cdot \text{SW}(\mathcal{B}) \geq \sum_{i \in [n]} v_i([m]) - \sum_{i \in [n]} \sum_{g \in G_i} v_i(g) + \sum_{i \in [n]} \sum_{g \in G_i} v_i(g) = \sum_{i \in [n]} v_i([m]),$$

which implies $\text{SW}(\mathcal{B}) \geq \frac{1}{2n} \sum_{i \in [n]} v_i([m])$, as needed. \square

3.2 The Case of High Optimal Welfare

As noted in Section 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 OPT , obtain $O(\sqrt{n})$ approximation to the *remaining* welfare through an EF1 allocation, and yet achieve $O(\sqrt{n})$ approximation to 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, our algorithm would like to use as reference an allocation \mathcal{W}^* with the optimal social welfare OPT . However, for subadditive valuations, computing such an allocation is known to be NP-hard under both value queries and demand queries [20]. Using a polynomial number of value queries, it is known that $\Theta(\sqrt{m})$ is the best possible approximation to the optimal social welfare [20; 33]. But this is too loose for our purpose. Hence, we instead turn to demand queries, for which an algorithm of Feige [22] yields the optimal 2-approximation. In particular, it efficiently computes an allocation \mathcal{W} with social welfare at least $1/2 \cdot \text{OPT}$.⁴ Outside of the black-box use of Feige’s algorithm, the rest of our algorithm uses value queries (which are a special case of demand queries). We note that for the special case of submodular valuations, it is possible to efficiently compute an allocation with social welfare at least $1/2 \cdot \text{OPT}$ using only value queries [29]; hence, for such special cases, our algorithm would not require access to demand queries. Even more importantly, we emphasize that our use of value or demand queries is only for the efficiency purpose; 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.⁵ Alternatively, 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*.




Figure 1: Line graph over the goods and unassigned components.

In this terminology, each W_i is a connected bundle in L (see Figure 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 by Arunachaleswaran et al. [3, Algorithm 1].

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 Figure 1. If there is an unassigned connected bundle $U \in \mathcal{U}(\mathcal{P})$ that some agent envies,

⁴We note that this allocation only serves as a reference in our algorithm, and may not be EF1 itself.

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

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. [30] to extend this partial EF1 allocation into a complete EF1 allocation without losing social welfare (while the algorithm of Lipton et al. may not preserve the partial allocation to each agent, it does ensure that the utility of each agent from her final allocation is at least that from her partial allocation; see, e.g., [2]). The detailed algorithm is presented as Algorithm 2.

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 [22] 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. [30] 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}
-

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

Lemma 2. *When ALG-EF1-HIGH is run on a fair division 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.*

Proof. The first property trivially holds because Line 14 updates the bundle of exactly one agent k in each iteration, and k is chosen such that she values her new bundle strictly more than her previously assigned

bundle. The second property also trivially holds because Line 14 always assigns a connected bundle to an agent.

For the third property, we use induction on t . For the base case of $t = 0$, note that \mathcal{P}^0 is trivially EF1 because $|P_i^0| \leq 1$ for each $i \in [n]$. Suppose \mathcal{P}^t is EF1 for some $t \in \mathbb{N}$. To see why \mathcal{P}^{t+1} is EF1, suppose agent k is assigned $\{g_a, \dots, g_c\}$ in iteration $t + 1$. Because $v_k(\{g_a, \dots, g_c\}) > v_k(P_k^t)$ by design, the only envy that may exist under \mathcal{P}^{t+1} is from an agent $i \neq k$ toward agent k . However, due to the minimality of c chosen in Line 12, no agent i can envy the bundle $\{g_a, \dots, g_{c-1}\}$. Hence, agent i may only envy agent k by up to one good (g_c in particular) under \mathcal{P}^{t+1} . Hence, \mathcal{P}^{t+1} is EF1. \square

Lemma 3. *When ALG-EF1-HIGH is run on a fair division instance $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$ with sub-additive 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 $U \in \mathcal{U}(\mathcal{P}^T)$.*

Proof. Suppose Line 14, in iteration t of the while loop, updates the allocation to agent k by assigning her the bundle $\{g_a, \dots, g_c\}$. Then, by Lemma 2, $\{g_a, \dots, g_c\}$ must be a connected bundle in the line graph L defined in Line 6 and agent k must strictly prefer this bundle to her previously assigned bundle.

Because there are $O(m^2)$ possible connected bundles in L , and because an agent's value for her assigned bundle strictly increases every time her assignment is updated, an agent's assignment can only be updated $O(m^2)$ times. Because each iteration of the while loop updates the assignment of one agent, there can be at most $O(nm^2)$ iterations of the while loop.

The second statement follows directly from the condition of the while loop; the while loop would not have terminated if $v_i(P_i^T) < v_i(U)$ for any agent $i \in [n]$ and any unassigned connected bundle $U \in \mathcal{U}(\mathcal{P}^T)$. \square

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} .

Proof. Recall that allocation \mathcal{W} returned by Feige's algorithm in Line 5 satisfies $\text{SW}(\mathcal{W}) \geq \frac{1}{2} \text{OPT}$. Let the while loop in ALG-EF1-HIGH terminate with partial allocation \mathcal{P} ; that is, $\mathcal{P} = \mathcal{P}^T$, where T is the number of iterations of the while loop. Lemma 3 ensures that the while loop terminates, and therefore, \mathcal{P} is well-defined. We want to show that

$$\sqrt{n} + 6\sqrt{n} \cdot \text{SW}(\mathcal{P}) \geq \text{SW}(\mathcal{W}) \tag{2}$$

By Lemma 2, \mathcal{P} is EF1. Also, note that the algorithm of Lipton et al. [30] used in Line 17 can extend a partial EF1 allocation \mathcal{P} into a complete EF1 allocation \mathcal{A} with $\text{SW}(\mathcal{A}) \geq \text{SW}(\mathcal{P})$. Hence, Equation (2), along with the fact that $\text{SW}(\mathcal{W}) \geq \frac{1}{2} \text{OPT}$, completes the proof of the stated claim.

Let us now prove that Equation (2) holds. First, due to Lemma 2, each P_i is a connected bundle in the line graph L . Hence, removing the goods allocated in \mathcal{P} creates at most $n + 1$ connected components, i.e., $|\mathcal{U}(\mathcal{P})| \leq n + 1$. Let $\mathcal{T}(\mathcal{P}) = \{P_j : j \in [n]\} \cup \mathcal{U}(\mathcal{P})$ be the collection of assigned and unassigned connected bundles in L created by \mathcal{P} . Hence, $|\mathcal{T}(\mathcal{P})| \leq 2n + 1$. For each agent $i \in [n]$, let \mathcal{Q}_i denote the bundles

in $\mathcal{T}(\mathcal{P})$ that intersect with W_i , i.e., $\mathcal{Q}_i := \{S \in \mathcal{T}(\mathcal{P}) : S \cap W_i \neq \emptyset\}$. Write $q_i := |\mathcal{Q}_i|$. For instance, in Figure 1, $\mathcal{Q}_1 = \{U_1, P_i\}$ and $q_1 = 2$.

Note that $\mathcal{T}(\mathcal{P})$ forms a partition of L into at most $2n + 1$ connected components, while \mathcal{W} forms a partition of L into at most n connected components. This, along with the observation that L is a line graph, can be used to bound the total number of intersections between the two. Specifically, we can move a marker from left to right, and keep track of the bundle from \mathcal{W} and the bundle from $\mathcal{T}(\mathcal{P})$ that contain the current good (and thus intersect). Every time the marker encounters a *new* intersection, it must have entered either a new bundle from \mathcal{W} or a new bundle from $\mathcal{T}(\mathcal{P})$. Thus, we have

$$\sum_{i \in [n]} q_i \leq (2n + 1) + n = 3n + 1. \quad (3)$$

Next, we partition the set of agents $[n]$ based on their q_i value. Let $H := \{i \in [n] : q_i > 3\sqrt{n}\}$ and $H^c := \{i \in [n] : q_i \leq 3\sqrt{n}\}$. Equation (3) immediately implies that $|H| \leq \sqrt{n}$. This, along with the observation that the valuations are scaled, gives us

$$\sum_{i \in H} v_i(W_i) \leq |H| \leq \sqrt{n}. \quad (4)$$

Next, we bound $v_i(W_i)$ for $i \in H^c$. By the definition of \mathcal{Q}_i , we have $v_i(W_i) = \sum_{S \in \mathcal{Q}_i} v_i(S \cap W_i)$. We break the summation into two parts depending on whether $S \in \mathcal{U}(\mathcal{P})$ or $S \in \mathcal{P}$.

1. If $S \in \mathcal{U}(\mathcal{P})$, note that from Lemma 3, we have $v_i(S \cap W_i) \leq v_i(S) \leq v_i(P_i) \leq 2v_i(P_i)$.
2. If $S \in \mathcal{P}$, then $S = P_j$ for some agent $j \in [n]$. Let us write $\hat{P}_j = P_j \cap W_i$. By the definition of \mathcal{Q}_i , $\hat{P}_j \neq \emptyset$. Because \mathcal{P} is EF1 (Lemma 2) and v_i is monotonic, agent i does not envy the bundle P_j (and thus the bundle \hat{P}_j) up to one good in it. Thus, there exists a good $\hat{g} \in \hat{P}_j$ such that $v_i(P_i) \geq v_i(\hat{P}_j \setminus \{\hat{g}\}) \geq v_i(\hat{P}_j) - v_i(\hat{g})$, where the last transition uses subadditivity of v_i . Further, $\hat{g} \in \hat{P}_j \subseteq W_i$. The initialization in Line 8 and Lemma 2 ensure that $v_i(\hat{g}) \leq v_i(P_i^0) \leq v_i(P_i)$. Combining with the previous step, we have $v_i(S \cap W_i) = v_i(\hat{P}_j) \leq 2v_i(P_i)$.

Thus, for $i \in H^c$, we have $v_i(W_i) \leq |\mathcal{Q}_i| \cdot 2v_i(P_i) \leq 3\sqrt{n} \cdot 2v_i(P_i)$. Thus, with Equation (4), we get

$$\text{SW}(\mathcal{W}) \leq \sqrt{n} + 6\sqrt{n} \cdot \sum_{i \in H^c} v_i(P_i) \leq \sqrt{n} + 6\sqrt{n} \cdot \text{SW}(\mathcal{P}),$$

which is the desired Equation (2), and the proof is complete. \square

3.3 Proof of Theorem 1

We are now ready to restate and prove Theorem 1. The upper bounds are established using Lemmas 1 and 4. As mentioned earlier, the lower bound for scaled valuations is already obtained by Bei et al. [10].

Theorem 1. *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.*

Proof. Let us first prove the upper bounds. 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.

We now prove the lower bounds. For scaled valuations, the desired $\Omega(\sqrt{n})$ bound is proved by Bei et al. [10]. For unscaled valuations, consider an instance with n agents and $m = n$ goods, and the following additive valuations: the first agent has value n for each good, while the other agents have value $1/n$ for each good. The optimal social welfare is achieved by allocating all goods to the first agents, yielding $\text{OPT} = n^2$. In contrast, that any EF1 allocation gives exactly one good to each agent, thus achieving welfare less than $n + 1$. This implies the desired $\Omega(n)$ bound. \square

Before concluding this section, we recall that for additive valuations, EF1 logically 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 construction from the proof of Theorem 1 for unscaled valuations and the construction due to Bei et al. [10] for scaled valuations can be modified slightly for Prop1 as well. For completeness, we provide these modified constructions in the appendix.

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. In this section, we limit our attention to additive valuations. Our main result, stated next, 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 begin by proving the upper bound, which is based on Algorithms 3 and 4. Recall that $\mathcal{W}^* = (W_1^*, \dots, W_n^*) \in \arg \max_{\mathcal{A} \in \Pi_n([m])} \text{SW}(\mathcal{A})$ denotes 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 $\text{SW}(\mathcal{A}) \geq \frac{1}{3n} \sum_{i=1}^n v_i([m])$. In fact, Algorithm 3 computes one such allocation efficiently.

The implication of this result is twofold. First, for both scaled and unscaled additive valuations, this establishes that the price of $1/2$ -MMS is $O(n)$ because, trivially, $\text{OPT} \leq \sum_{i=1}^n v_i([m])$. For unscaled valuations, this is precisely the bound we seek. For scaled valuations, we need to improve this to $O(\sqrt{n})$. To that end, note that for scaled valuations, $v_i([m]) = 1$ for each agent $i \in [n]$, so $\text{SW}(\mathcal{A}) \geq 1/3 = \Omega(1)$. Thus, at least when $\text{OPT} = O(\sqrt{n})$, this already establishes the desired upper bound of $O(\sqrt{n})$ on the price of $1/2$ -MMS, allowing us to limit our attention to instances with $\text{OPT} = \Omega(\sqrt{n})$, as we do in Section 4.2.

Algorithm 3 ALG-MMS-ABS

Input: A fair division instance $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$ 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 provides the desired welfare guarantee as well.

We note that there is another minor difference. In Line 6, we compute an arbitrary Prop1 allocation of the remaining instance. While Amanatidis et al. [1] use a specific method — greedy round-robin — that is guaranteed to return a Prop1 allocation, they note that their proof works for any Prop1 allocation. We state our algorithm with this generality in order to avoid introducing how greedy round-robin works. There are other methods that can be used to compute a Prop1 allocation [8; 17; 31].

To establish the desired welfare bound, we first establish a lower bound on the agents' values for the goods that remain at various iterations of the algorithm. For the given instance $\langle [n], [m], \{v_i\}_{i \in [n]} \rangle$, we let t denote the number of iterations of the while loop in ALG-MMS-ABS. We re-index the agents such that agents 1 through t receive, in order, a good (a singleton bundle) in this while loop. That is, each agent $i \in [t]$ is assigned a good in the i^{th} iteration of the loop. The remaining $n - t$ agents receive a bundle through the Prop1 allocation computed in Line 6.

Lemma 5. *Consider a fair division instance $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$ with the re-indexing of the agents mentioned above, and t denoting the number of agents assigned a good in the while loop of ALG-MMS-ABS. Let $\mathcal{B} = (B_1, B_2, \dots, B_n)$ denote the allocation returned by ALG-MMS-ABS. Then, we have*

$$\sum_{k=1}^{i-1} \frac{1}{n-k+1} \cdot v_k(B_k) + \frac{1}{n-i+1} \cdot v_i([m] \setminus \cup_{k=1}^{i-1} B_k) \geq \frac{1}{n} \cdot v_i([m]) \quad \text{for all } i \in [t], \text{ and}$$
$$\sum_{k=1}^t \frac{1}{n-k+1} \cdot v_k(B_k) + \frac{1}{n-t} \cdot v_i([m] \setminus \cup_{k=1}^t B_k) \geq \frac{1}{n} \cdot v_i([m]) \quad \text{for all } i \in [n] \setminus [t].$$

Proof. For each agent $i \in [n]$, we establish the following inequality for every index $j \in \{0, \dots, \min\{i-1, t\}\}$, via an induction on j . Then, instantiating $j = \min\{i-1, t\}$ yields the desired result.

$$\sum_{k=1}^j \frac{1}{n-k+1} \cdot v_k(B_k) + \frac{1}{n-j} \cdot v_i([m] \setminus \cup_{k=1}^j B_k) \geq \frac{1}{n} \cdot v_i([m]). \quad (5)$$

Fix an arbitrary agent $i \in [n]$. The base case of $j = 0$ is trivial, as it reduces to the inequality $\frac{1}{n} \cdot v_i([m]) \geq \frac{1}{n} \cdot v_i([m])$. For the induction hypothesis, suppose that Equation (5) holds with the index set to $j - 1$. That

is, we have

$$\sum_{k=1}^{j-1} \frac{1}{n-k+1} \cdot v_k(B_k) + \frac{1}{n-j+1} \cdot v_i \left([m] \setminus \cup_{k=1}^{j-1} B_k \right) \geq \frac{1}{n} \cdot v_i([m]). \quad (6)$$

To complete the proof via induction, we need to extend this inequality to index j . Note that $j \leq \min \{i-1, t\} \leq t$. Hence, agent j is assigned a singleton bundle $B_j := \{g_j\}$, and this assignment is made before i receives a bundle. Consider two cases about the value that agent i places on the good g_j .

Case 1: $v_i(B_j) < \frac{1}{n-j+1} v_i \left([m] \setminus \cup_{k=1}^{j-1} B_k \right)$. Note that before agent j was assigned B_j , there were $n-j+1$ agents left. In this case, agent i 's value for B_j is below the ‘‘average value’’ of the remaining goods divided into $n-j+1$ bundles. Hence, the ‘‘average’’ goes up after assigning B_j to agent j , i.e.,

$$\frac{1}{n-j} v_i \left([m] \setminus \cup_{k=1}^j B_k \right) \geq \frac{1}{n-j+1} v_i \left([m] \setminus \cup_{k=1}^{j-1} B_k \right). \quad (7)$$

To see this formally, let us rearrange the inequality of this case as

$$(n-j+1) \cdot v_i(B_j) < v_i \left([m] \setminus \cup_{k=1}^{j-1} B_k \right),$$

or equivalently,

$$(n-j) \cdot v_i(B_j) < v_i \left([m] \setminus \cup_{k=1}^j B_k \right).$$

Adding $(n-j) \cdot v_i \left([m] \setminus \cup_{k=1}^j B_k \right)$ on both sides and dividing both sides by $(n-j) \cdot (n-j+1)$ yields Equation (7).

Combining this with the induction hypothesis, we get that in Equation (6), changing the index from $j-1$ to j adds $\frac{1}{n-j+1} \cdot v_j(B_j) \geq 0$ to the first term in LHS, and due to Equation (7), the second term also weakly increases. Hence, the LHS weakly increases, and by the induction hypothesis, stays at least $\frac{1}{n} \cdot v_i([m])$.

Case 2: $v_i(B_j) \geq \frac{1}{n-j+1} v_i \left([m] \setminus \cup_{k=1}^{j-1} B_k \right)$. This condition implies that the pair (i, g_j) was considered in Line 3 of the algorithm in iteration j of the while loop. The argmax selection criterion of this step implies that $v_j(B_j) \geq v_i(B_j)$. Now, we have

$$\begin{aligned} & \sum_{k=1}^j \frac{1}{n-k+1} \cdot v_k(B_k) + \frac{1}{n-j} \cdot v_i \left([m] \setminus \cup_{k=1}^j B_k \right) \\ & \geq \sum_{k=1}^j \frac{1}{n-k+1} \cdot v_k(B_k) + \frac{1}{n-j+1} \cdot v_i \left([m] \setminus \cup_{k=1}^j B_k \right) \\ & = \sum_{k=1}^{j-1} \frac{1}{n-k+1} \cdot v_k(B_k) + \frac{1}{n-j+1} \cdot v_j(B_j) + \frac{1}{n-j+1} \cdot v_i \left([m] \setminus \cup_{k=1}^j B_k \right) \\ & \geq \sum_{k=1}^{j-1} \frac{1}{n-k+1} \cdot v_k(B_k) + \frac{1}{n-j+1} \cdot v_i(B_j) + \frac{1}{n-j+1} \cdot v_i \left([m] \setminus \cup_{k=1}^j B_k \right) \\ & = \sum_{k=1}^{j-1} \frac{1}{n-k+1} \cdot v_k(B_k) + \frac{1}{n-j+1} \cdot v_i \left([m] \setminus \cup_{k=1}^{j-1} B_k \right) \\ & \geq \frac{1}{n} \cdot v_i([m]), \end{aligned}$$

where the third transition follows due to $v_j(B_j) \geq v_i(B_j)$, and the last inequality follows due to the induction hypothesis (Equation (6)).

Setting $j = \min\{i - 1, t\}$ for each agent i in Equation (5) completes the proof. \square

We now use Lemma 5 to derive the desired guarantees for ALG-MMS-ABS (Algorithm 3).

Lemma 6. *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 $\text{SW}(\mathcal{B}) \geq \frac{1}{3n} \sum_{i=1}^n v_i([m])$.*

Proof. The fact that the algorithm runs efficiently and computes a $1/2$ -MMS allocation is due to Amanatidis et al. [1], as argued above. It remains to show that the stated welfare bound holds.

Recall that t denotes the number of iteration of the while loop in ALG-MMS-ABS, and we re-indexed the agents so that each agent $i \in [t]$ is assigned a singleton bundle $B_i = \{g_i\}$ in iteration i of the while loop.

For an agent $i \in [t]$, the selection in Line 3 of the algorithm in i^{th} iteration of the while loop ensures that $v_i(B_i) \geq \frac{1}{2(n-i+1)} \cdot v_i([m] \setminus \cup_{k=1}^{i-1} B_k)$. Substituting this into Lemma 5, we have

$$\sum_{k=1}^{i-1} \frac{1}{n-k+1} \cdot v_k(B_k) + 2 \cdot v_i(B_i) \geq \frac{1}{n} \cdot v_i([m]) \quad \text{for all } i \in [t]. \quad (8)$$

Similarly, fix an agent $i \in [n] \setminus [t]$. Note that at the termination of the while loop, we have $A = [n] \setminus [t]$, $G = [m] \setminus \cup_{k=1}^t B_k$ and $v_i(g) < \frac{1}{2(n-t)} v_i(G)$ for each $g \in G$. Because Line 6 computes a Prop1 allocation of G to agents in A , we have that for some $g \in G$, $v_i(B_i) + v_i(g) \geq \frac{1}{n-t} v_i(G)$. The inequality $v_i(g) < \frac{1}{2(n-t)} v_i(G)$ yields $v_i(B_i) > \frac{1}{2(n-t)} v_i(G)$. Substituting this into Lemma 5, we have

$$\sum_{k=1}^t \frac{1}{n-k+1} \cdot v_k(B_k) + 2 \cdot v_i(B_i) \geq \frac{1}{n} \cdot v_i([m]) \quad \text{for all } i \in [n] \setminus [t]. \quad (9)$$

Summing Equations (8) and (9) across all agents yields the desired welfare bound:

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n v_i([m]) &\leq \sum_{i=1}^t \left(2 \cdot v_i(B_i) + \sum_{k=1}^{i-1} \frac{1}{n-k+1} \cdot v_k(B_k) \right) + \sum_{i=t+1}^n \left(2 \cdot v_i(B_i) + \sum_{k=1}^t \frac{1}{n-k+1} \cdot v_k(B_k) \right) \\ &= 2 \cdot \sum_{i=1}^n v_i(B_i) + \sum_{i=1}^t \sum_{k=1}^{i-1} \frac{1}{n-k+1} \cdot v_k(B_k) + \sum_{i=t+1}^n \sum_{k=1}^t \frac{1}{n-k+1} \cdot v_k(B_k) \\ &= 2 \cdot \sum_{i=1}^n v_i(B_i) + \sum_{k=1}^t \sum_{i=k+1}^n \frac{1}{n-k+1} \cdot v_k(B_k) \\ &= 2 \cdot \sum_{i=1}^n v_i(B_i) + \sum_{k=1}^t \frac{n-k}{n-k+1} \cdot v_k(B_k) \\ &\leq 2 \cdot \sum_{i=1}^n v_i(B_i) + \sum_{k=1}^n v_k(B_k) = 3 \cdot \sum_{i=1}^n v_i(B_i). \end{aligned}$$

This completes the proof of the lemma. \square

4.2 The Case of High Optimal Welfare

As argued in Section 4.1, Lemma 6 allows us to restrict our attention to scaled additive valuations in which the optimal social welfare $\text{OPT} = \Omega(\sqrt{n})$. Similarly 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 [38]. For a fixed $\epsilon \in (0, 1)$, this PTAS can compute an estimate $Z_i \in [(1 - \epsilon)\text{MMS}_i, \text{MMS}_i]$ 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 introduced here is solely for computational purposes. To derive our main existential result about the price of $\frac{1}{2}$ -MMS, we can simply pass $Z_i = \text{MMS}_i$ to ALG-MMS-HIGH, and it would return an exact $\frac{1}{2}$ -MMS allocation with the desired welfare guarantee.

Lemma 7. *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 $Z_i \in [(1 - \epsilon)\text{MMS}_i, \text{MMS}_i]$ for each agent i , ALG-MMS-HIGH (Algorithm 4) efficiently computes a $(\frac{1}{2} - \epsilon)$ -MMS allocation \mathcal{B} with the property that $3\sqrt{n} \cdot \text{SW}(\mathcal{B}) + 4\sqrt{n} \geq \text{OPT}$; here, OPT denotes the optimal social welfare in \mathcal{I} .*

The intuition behind the algorithm is as follows. Our goal is to assign each agent i a bundle of value at least $\frac{1}{2} \cdot Z_i \geq (\frac{1}{2} - \epsilon) \cdot \text{MMS}_i$, 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 .

Line 5 and 9 handle easy cases, when an agent i either has $\text{MMS}_i = 0$, 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 (proved below) 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.

We partition the set of agents into three sets— Γ_{MMS} , Γ_{single} , and Γ_{hard} —where agent i is placed into a set depending on how the estimate Z_i of the maximin share MMS_i relates to $v_i(W_i^*)$.

$$\begin{aligned} \Gamma_{\text{MMS}} &:= \left\{ i \in [n] : Z_i \geq \frac{2}{3\sqrt{n}}v_i(W_i^*) \right\}, \\ \Gamma_{\text{single}} &:= \left\{ i \in [n] : Z_i < \frac{2}{3\sqrt{n}}v_i(W_i^*), \text{ and there exists } g \in W_i^* \text{ such that } v_i(g) \geq \frac{1}{3\sqrt{n}}v_i(W_i^*) \right\}, \\ \Gamma_{\text{hard}} &:= \left\{ i \in [n] : Z_i < \frac{2}{3\sqrt{n}}v_i(W_i^*), \text{ and for all } g \in W_i^* \text{ we have } v_i(g) < \frac{1}{3\sqrt{n}}v_i(W_i^*) \right\}. \end{aligned}$$

For an agent $i \in \Gamma_{\text{single}}$, giving her a single good from $\arg \max_{g \in W_i^*} v_i(g)$ gives her value at least $\frac{1}{3\sqrt{n}}v_i(W_i^*) > \frac{1}{2} \cdot Z_i$. Thus, ALG-MMS-HIGH begins by giving each such agent i her most valuable good

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 a .

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_i = 0 iff agent i has positive value for less than n goods, which can be checked efficiently.}
 - 6: Let $\Gamma_{\text{single}} = \left\{ i \in [n] : Z_i < \frac{2}{3\sqrt{n}} v_i(W_i^*), \text{ and there exists } g \in W_i^* \text{ such that } v_i(g) \geq \frac{1}{3\sqrt{n}} v_i(W_i^*) \right\}$
 - 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)
-

from W_i^* and adding her directly to P . Next, for an agent $i \in \Gamma_{\text{MMS}}$, giving her value at least $\frac{1}{2} \cdot Z_i$ also guarantees that her value is at least $\frac{1}{3\sqrt{n}} v_i(W_i^*)$; thus, ALG-MMS-HIGH solely aims to give such agents $\frac{1}{2} \cdot Z_i$ value. The situation is more tricky for agents in Γ_{hard} .

Before we begin to analyze ALG-MMS-HIGH, we need the following lemma. This builds upon a similar result by Amanatidis et al. [1], who show that giving a subset of agents a single good each and removing these agents and their assigned goods from the instance cannot reduce the maximin share of any remaining agent. In our case, we need to consider assigning bundles that either consist of a single good or are of limited worth to an agent ℓ , and we are interested in the value that agent ℓ must have for the remaining goods. This lemma will be useful later in analyzing ALG-MMS-HIGH. Let us first introduce some useful notation. For

an agent $i \in [n]$, an integer parameter $k \in \mathbb{N}$, and a subset of goods $G \subseteq [m]$, define

$$\text{MMS}_i^k(G) := \max_{(P_1, \dots, P_k) \in \Pi_k(G)} \min_{j \in [k]} v_i(P_j).$$

This is the highest value that agent i can guarantee by splitting G into k bundles and receiving the worst bundle. Note that $\text{MMS}_i = \text{MMS}_i^n([m])$.

Lemma 8. *Let $\mathcal{I} = \langle [n], [m], \{v_i\}_{i \in [n]} \rangle$ be a fair division instance with additive valuations. Let $\ell \in [n]$ be an agent. Let $\{B_a\}_{a \in S}$ be a collection of bundles assigned to a subset of agents $S \subseteq [n] \setminus \{\ell\}$ with the property that, for all $a \in S$, we either have $|B_a| = 1$ or $v_\ell(B_a) \leq \text{MMS}_\ell$. Then, the value of agent ℓ for the unassigned goods satisfies*

$$v_\ell([m] \setminus \cup_{a \in S} B_a) \geq (n - |S|) \cdot \text{MMS}_\ell.$$

Proof. Amanatidis et al. [1] show that removing a single good along with a single agent from the instance does not decrease the maximin share of any remaining agent; in particular, for any integer $k \in \mathbb{N}$ and subset of goods $G \subseteq [m]$, the following inequality holds for all $g \in G$ and agents $\ell \in [n]$:

$$\text{MMS}_\ell^k(G) \leq \text{MMS}_\ell^{k-1}(G \setminus \{g\}) \quad (10)$$

Among the set of agents S considered in the lemma statement, write $S' \subseteq S$ to denote the ones that receive exactly one good, i.e., $S' := \{a \in S : |B_a| = 1\}$. A repeated application of Equation (10) gives us

$$\text{MMS}_\ell^{n-|S'|}([m] \setminus \cup_{a \in S'} B_a) \geq \text{MMS}_\ell^n([m]) = \text{MMS}_\ell.$$

Furthermore, using the fact that the valuation v_ℓ is additive, we get

$$v_\ell([m] \setminus \cup_{a \in S'} B_a) \geq (n - |S'|) \cdot \text{MMS}_\ell^{n-|S'|}([m] \setminus \cup_{a \in S'} B_a) \geq (n - |S'|) \cdot \text{MMS}_\ell. \quad (11)$$

For each remaining agent $a \in S \setminus S'$, the assumption in the lemma's statement ensures that $v_\ell(B_a) \leq \text{MMS}_\ell$. Therefore,

$$\begin{aligned} v_\ell([m] \setminus \cup_{a \in S} B_a) &= v_\ell([m] \setminus \cup_{a \in S'} B_a) - \sum_{a \in S \setminus S'} v_\ell(B_a) && \text{(since } v_\ell \text{ is additive)} \\ &\geq (n - |S'|) \cdot \text{MMS}_\ell - \sum_{a \in S \setminus S'} v_\ell(B_a) && \text{(via Equation (11))} \\ &\geq (n - |S'|) \cdot \text{MMS}_\ell - |S \setminus S'| \text{MMS}_\ell \\ &= (n - |S|) \cdot \text{MMS}_\ell, \end{aligned}$$

which is the desired inequality. \square

We now begin our analysis of ALG-MMS-HIGH. First, we show that certain invariants hold throughout the execution of the algorithm.

Lemma 9. *The following hold at any stage during the execution of ALG-MMS-HIGH*

1. $P \cap T = \emptyset$.
2. If $i \in P$, $v_i(B_i) \geq \frac{1}{3\sqrt{n}} v_i(W_i^*)$ and B_i is never updated afterwards.
3. If $i \in P \cup T$, $v_i(B_i) \geq \frac{1}{2} \cdot Z_i$, and i is never removed from $P \cup T$ afterwards.

4. $T \subseteq \Gamma_{\text{hard}}$

Proof. The first two invariants are easy to verify by considering at the pseudocode of the algorithm. For the third invariant, note that once an agent i is added to $P \cup T$, it is never removed (but can move from T to P).

Next we show that when $i \in P \cup T$, we must have $v_i(B_i) \geq \frac{1}{2} \cdot Z_i$. Let us begin by showing that when i is *first* added to $P \cup T$, we have $v_i(B_i) \geq \frac{1}{2} \cdot Z_i$. This can happen in one of Lines 9, 13, or 28. In each of these cases, one can verify that the condition for selection of i ensures that $v_i(B_i) \geq \frac{1}{2} \cdot Z_i$ at the time it is included in $P \cup T$. We also need to consider Line 24, in which the bundle of an agent i , already added to T , can be updated. Note that, for such an agent i , we cannot have $i \in \Gamma_{\text{single}}$, as agents in Γ_{single} are immediately added to P in Line 16. Similarly, we also cannot have $i \in \Gamma_{\text{MMS}}$ because in that case, $\frac{1}{2} \cdot Z_i \geq \frac{1}{3\sqrt{n}} \cdot v_i(W_i^*)$; hence, when i was first added to $P \cup T$, it would have been added to P directly. Therefore, for an agent i considered in Line 24 we must have $i \in \Gamma_{\text{hard}}$. However, in that case, the selection criterion in Line 24 implies that $v_i(B_i) \geq \frac{1}{3\sqrt{n}} \cdot v_i(W_i^*) \geq \frac{1}{2} \cdot Z_i$, as desired.

Finally, we show that $T \subseteq \Gamma_{\text{hard}}$ throughout the execution of the algorithm. This is because each agent $i \in \Gamma_{\text{single}}$ is added to P at Line 9, and since $T \cap P = \emptyset$, we have that any agent $i \in \Gamma_{\text{single}}$ cannot be in T . Similarly for agents in Γ_{MMS} , recall that whenever an agent $i \in \Gamma_{\text{MMS}}$ is added to $P \cup T$ she must have received a bundle worth at least $\frac{1}{2} \cdot Z_i$ to her. Since $\frac{1}{2} \cdot Z_i \geq \frac{1}{3\sqrt{n}} \cdot v_i(W_i^*)$ for all $i \in \Gamma_{\text{MMS}}$, agent i would be added included directly into P at this point. Therefore, any agent that is placed in T upon receiving a bundle must belong to the set Γ_{hard} , i.e., the containment $T \subseteq \Gamma_{\text{hard}}$ holds throughout the execution of ALG-MMS-HIGH. \square

Next, we show that the algorithm terminates in polynomial time, and certain desired properties hold at its termination. As we will see later, these properties are key to establishing the desired welfare guarantee of the algorithm.

Lemma 10. ALG-MMS-HIGH terminates in polynomial time, and at its termination, the following hold.

1. $P \cup T = [n]$.
2. $|T| \leq 4\sqrt{n}$.

Proof. Termination: To argue that the algorithm terminates, we must show that the while loop (Line 20) terminates. This is easy to see because in each iteration of the loop, a new agent is added to $P \cup T$, and by the third invariant in Lemma 9, it is never removed from $P \cup T$ once added. It is easy to verify that the algorithm runs in polynomial time.

Let \bar{P} and \bar{T} denote the values of sets P and T at the termination of the algorithm, respectively. We want to show that $\bar{P} \cup \bar{T} = [n]$ and $|\bar{T}| \leq 4\sqrt{n}$.

$\bar{P} \cup \bar{T} = [n]$: Because $P \cup T$ grows monotonically (Lemma 9), we only need to show that each agent is included in $P \cup T$ at some point. Suppose, towards a contradiction, that agent i is never added to $P \cup T$. Since i is not added to $P \cup T$ during the first while loop in Line 13, it must be the case that

$$v_i(g) < \frac{1}{2} \cdot Z_i, \forall g \in R. \quad (12)$$

Otherwise, the while loop would not have terminated with goods in R left unassigned.

We want to show that whenever an agent $j \neq i$ is assigned a nonempty bundle B_j during the execution of the algorithm, B_j satisfies the condition of Lemma 8, i.e., either $|B_j| = 1$ or $v_i(B_j) \leq \text{MMS}_i$. We prove this by induction on the progress of the algorithm. Note that bundles assigned during Lines 9 and 13 are singleton, and therefore trivially satisfy the former condition.

Consider a bundle K is assigned to an agent in the t^{th} iteration of the for loop in Line 17, where $g_t \in R$ (that is, g_t was unassigned at the start of the loop, and was thus added to K). If $g_t \notin K$, then the swap in

Line 28 must have been executed. In this case, K is the bundle that some agent was previously assigned, and the induction hypothesis implies that K satisfies the condition of Lemma 8. If $g_t \in K$, then we consider two cases. If $K = \{g_t\}$, then the former condition trivially holds. If $|K| \geq 2$, then note that $K \setminus \{g_t\}$ was not assigned in the previous iteration of the for loop. Since we had $i \in [n] \setminus (P \cup T)$, it must have been the case that $v_i(K \setminus \{g_t\}) < \frac{1}{2} \cdot Z_i$. Combining this with Equation (12) and using the additivity of v_i , we have $v_i(K) \leq Z_i \leq \text{MMS}_i$. Hence, the latter condition is satisfied.

Consider the set of unassigned goods X at the end of the for loop in Line 24. We have established that at this stage, for each $j \in P \cup T$, either $|B_j| = 1$ or $v_i(B_j) \leq \text{MMS}_i$. Using Lemma 8, we have $v_i(X) \geq (n - |P \cup T|) \cdot \text{MMS}_i \geq \text{MMS}_i > 0$, where the penultimate inequality follows because $|P \cup T| \leq n - 1$ as $i \notin P \cup T$, and the final inequality follows because all agents i with $\text{MMS}_i = 0$ are added to $P \cup T$ in Line 5 and then never removed from it (Lemma 9). Note that $v_i(X) > 0$ implies $X \neq \emptyset$. Then, in the last iteration of the for loop, we must have had $K = X$. This equality follows from the observation that as soon as any good gets unassigned it is included in the set K (Line 23). Hence, for every iteration count t , the set K contains all unassigned goods from $\{g_1, \dots, g_t\}$. Furthermore, note that at termination $K = X$ was not assigned to any remaining agent. However, since $i \in [n] \setminus (P \cup T)$ and $v_i(X) \geq \text{MMS}_i$, agent i should have been assigned $K = X$ in Line 28, which is the desired contradiction.

$|\bar{T}| \leq 4\sqrt{n}$: Lemma 9 ensures that the containment $T \subseteq \Gamma_{\text{hard}}$ holds throughout the execution of the algorithm. Hence, in particular, we have $\bar{T} \subseteq \Gamma_{\text{hard}}$.

Using this property and a charging argument, we will establish the inequality $|\bar{T}| \leq \sqrt{n}$. We know from $\bar{P} \cup \bar{T} = [n]$ and monotonic growth of $P \cup T$ that each agent is added to $P \cup T$ exactly once during the execution of the algorithm. Define a partition (C_1, \dots, C_n) of the set of agents $[n]$ as follows. For each agent j , consider the bundle B_j that she is assigned the first time she is added to $P \cup T$ (i.e., in one of Lines 9, 13, or 28, but not in Line 24), and add j to C_i , where agent i is such that W_i^* contains the highest indexed good in B_j , according to the indexing in Line 2. That is, the set C_i populates agents j whose initial bundle B_j ends in W_i^* . Our goal is to show that for each $i \in \bar{T}$, $|C_i| \geq \sqrt{n} - 1$, which (for $n \geq 2$) would imply $|\bar{T}| \leq 4\sqrt{n}$.

To see this, fix $i \in \bar{T}$, and consider a partition of C_i into C_i^1 and C_i^2 , where C_i^1 contains agents in C_i who were first added to $P \cup T$ in either Line 9 or Line 13, and C_i^2 contains agents in C_i who were first added to $P \cup T$ in Line 28. If $|C_i^1| \geq \sqrt{n}$, then trivially $|C_i| \geq \sqrt{n}$, and we are done. Thus, assume $|C_i^1| < \sqrt{n}$.

Note that C_i^1 consists of all agents who were assigned a good from W_i^* before the start of the for loop in Line 17. Since $i \in \bar{T} \subseteq \Gamma_{\text{hard}}$, $v_i(g) < \frac{1}{3\sqrt{n}} \cdot v_i(W_i^*)$ for each $g \in W_i^*$. Since there were less than \sqrt{n} such assignments, we have that for the goods R left unassigned before the for loop,

$$v_i(R \cap W_i^*) \geq v_i(W_i^*) - \sqrt{n} \cdot \frac{1}{3\sqrt{n}} \cdot v_i(W_i^*) \geq \frac{2}{3} \cdot v_i(W_i^*).$$

Let $W_i^* = \{g_t : t_1 \leq t \leq t_2\}$ (the indexing in Line 2 ensures that W_i^* has consecutively indexed goods). We show that during iterations t_1 through t_2 of the for loop, at least \sqrt{n} agents will be added to C_i . Note that because agent i is never added to P , the if condition in Line 22 must never be executed during these iterations. Hence, any bundle K assigned in Line 28 during iterations t_1 through t_2 must be worth at most $\frac{2}{3\sqrt{n}} \cdot v_i(W_i^*)$ to agent i : if agent i is in the set T , then the non-execution of the if-clause in Line 22 gives us $v_i(K) < \frac{1}{3\sqrt{n}} \cdot v_i(W_i^*)$. Otherwise, even if $i \in \Gamma_{\text{hard}}$ is not in T , then the removal of the highest-index good in K must reduce its value below $\frac{1}{2} \cdot Z_i \leq \frac{1}{3\sqrt{n}} v_i(W_i^*)$. The definition of Γ_{hard} ensures this highest-index good is of value at most $\frac{1}{3\sqrt{n}} v_i(W_i^*)$, and hence, $v_i(K) < \frac{2}{3\sqrt{n}} \cdot v_i(W_i^*)$.

Since $v_i(R \cap W_i^*) \geq \frac{2}{3} \cdot v_i(W_i^*)$, at least $\sqrt{n} - 1$ such assignments must be made during iterations t_1 through t_2 , where we subtract 1 for the last batch of goods from $W_i^* \cap R$ (again, worth at most $\frac{2}{3\sqrt{n}} \cdot v_i(W_i^*)$) that may be assigned after iteration t_2 . Hence, we have $|C_i| \geq \sqrt{n} - 1$ in this case as well. This completes the charging argument, and thus the proof that $|\bar{T}| \leq 4\sqrt{n}$. \square

We are now ready to establish the desired guarantees of ALG-MMS-HIGH from Lemma 7.

Proof of Lemma 7. The fact that ALG-MMS-HIGH terminates in polynomial time is proved in Lemma 10. We need to show that it returns a $(\frac{1}{2} - \epsilon)$ -MMS allocation with the desired welfare guarantee.

As before, let \bar{P} and \bar{T} denote the values of P and T at the termination of the algorithm, respectively. Lemma 10 gives us $\bar{P} \cup \bar{T} = [n]$, and Lemma 9 ensures that for each $i \in \bar{P} \cup \bar{T} = [n]$, $v_i(B_i) \geq \frac{1}{2} \cdot Z_i \geq \frac{1}{2} \cdot (1 - \epsilon) \cdot \text{MMS}_i \geq (\frac{1}{2} - \epsilon) \cdot \text{MMS}_i$. Thus, the allocation returned by ALG-MMS-HIGH is $(\frac{1}{2} - \epsilon)$ -MMS.

For the welfare guarantee, note that for each $i \in \bar{P}$, Lemma 9 shows that $v_i(B_i) \geq \frac{1}{3\sqrt{n}} \cdot v_i(W_i^*)$. And Lemma 10 shows that $|\bar{T}| \leq 4\sqrt{n}$. Combining them, we have

$$\text{OPT} = \sum_{i \in [n]} v_i(W_i^*) = \sum_{i \in \bar{P}} v_i(W_i^*) + \sum_{i \in \bar{T}} v_i(W_i^*) \leq 3\sqrt{n} \sum_{i \in \bar{P}} v_i(B_i) + |\bar{T}| \leq 3\sqrt{n} \sum_{i \in [n]} v_i(B_i) + 4\sqrt{n},$$

where the third transition holds because the valuations are scaled. \square

4.3 Proof of Theorem 2

Using the properties of ALG-MMS-ABS and ALG-MMS-HIGH from Sections 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. [16]. We restate and prove the theorem below.

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

Proof. Unscaled valuations: Lemma 6 shows that ALG-MMS-ABS returns a $1/2$ -MMS allocation \mathcal{B} with social welfare $\text{SW}(\mathcal{B}) \geq \frac{1}{3n} \sum_{i=1}^n v_i([m])$. Because the optimal social welfare is trivially upper bounded as $\text{OPT} \leq \sum_{i=1}^n v_i([m])$, we have that the price of $1/2$ -MMS is at most $3n = O(n)$.

To show that this is tight, consider a fair division instance with n agents with the following (unscaled) additive valuations over n goods. Fix $\epsilon \in (0, 1)$. Let $v_1(j) = 1$ for each $j \in [n]$. For each $i \in [n] \setminus \{1\}$, $v_i(j) = \epsilon$ for each $j \in [m]$. Thus, the goods are identical for all agents. It is easy to see that $\text{MMS}_1 = 1$ while $\text{MMS}_i = \epsilon$ for each $i \in [n] \setminus \{1\}$. Hence, the only $1/2$ -MMS allocation is to give each agent a single good. This generates social welfare $1 + (n - 1)\epsilon$, whereas the optimal social welfare is n , obtained by assigning all the goods to agent 1. The ratio $n/(1 + (n - 1)\epsilon)$ converges to n as $\epsilon \rightarrow 0$. Hence, the price of $1/2$ -MMS is at least $n = \Omega(n)$.

Scaled valuations: Again, from Lemma 6, we know that ALG-MMS-ABS (with the input $Z_i = \text{MMS}_i$ for each agent i) returns a $1/2$ -MMS allocation \mathcal{B} with $\text{SW}(\mathcal{B}) \geq \frac{1}{3n} \sum_{i=1}^n v_i([m]) = \frac{1}{3}$ for scaled valuations. Hence, if $\text{OPT} \leq 5\sqrt{n}$, then the $1/2$ -MMS allocation returned by ALG-MMS-ABS gives a $15\sqrt{n}$ approximation to the optimal social welfare.

Otherwise, suppose $\text{OPT} \geq 5\sqrt{n}$. Then by Lemma 7, we know that ALG-MMS-HIGH returns an allocation with social welfare $\text{SW}(\mathcal{B}) \geq \frac{\text{OPT} - 4\sqrt{n}}{3\sqrt{n}} \geq \frac{\text{OPT} - 4/5 \text{OPT}}{3\sqrt{n}} = \frac{\text{OPT}}{15\sqrt{n}}$. Hence, in this case, the $1/2$ -MMS allocation returned by ALG-MMS-HIGH provides a $15\sqrt{n}$ approximation to OPT.

We have established that in either case, the price of $1/2$ -MMS is at most $15\sqrt{n} = O(\sqrt{n})$. For the lower bound, we use a construction from Caragiannis et al. [16], which later appeared in several other papers on the price of fairness. The instance has n agents and goods, with the following scaled additive valuations:

- For $i \leq \lfloor \sqrt{n} \rfloor$, agent i has value $1/\lfloor \sqrt{n} \rfloor$ for the $\lfloor \sqrt{n} \rfloor$ goods $[(i - 1)\lfloor \sqrt{n} \rfloor + 1$ to $i\lfloor \sqrt{n} \rfloor]$, and value 0 for the remaining goods. Call such an agent a *high* agent.
- For $i > \lfloor \sqrt{n} \rfloor$, agent i has value $1/n$ for all n goods. Call such an agent a *low* agent.

It is easy to check that for each low agent i , $MMS_i = 1/n$. Thus, any $1/2$ -MMS allocation must give each such agent at least one good. The social welfare of any such allocation is at most $(n - \lfloor \sqrt{n} \rfloor) \cdot \frac{1}{n} + \lfloor \sqrt{n} \rfloor \cdot \frac{1}{\lfloor \sqrt{n} \rfloor} \leq 2$, where the second term in the summation is due to the fact that each of the remaining (at most) $\lfloor \sqrt{n} \rfloor$ goods which may be assigned to high agents can contribute at most $1/\lfloor \sqrt{n} \rfloor$ to the social welfare. In contrast, the social welfare generated by assigning each high agent the $\lfloor \sqrt{n} \rfloor$ distinct goods she values positively is at least $\lfloor \sqrt{n} \rfloor \cdot 1 = \lfloor \sqrt{n} \rfloor$. Hence, the price of $1/2$ -MMS is at least $\lfloor \sqrt{n} \rfloor/2 = \Omega(\sqrt{n})$. \square

5 Discussion

In this work, we focus on the allocation of indivisible goods to n agents. We show that the price of EF1 is $O(\sqrt{n})$ when agent valuations are subadditive (and $\Omega(\sqrt{n})$ even if they are additive), whereas the price of $1/2$ -MMS is $\Theta(\sqrt{n})$ when agent valuations are additive.

An immediate future direction is to analyze the price of $1/2$ -MMS under more general valuation classes. But even for additive valuations, it would be interesting to consider other (combinations of) fairness and efficiency guarantees. The following are just a few interesting examples.

1. EF1+PO (Pareto optimality), and EF1+fPO (fractional Pareto optimality): fPO implies PO, and both of these allocations are guaranteed to exist [9; 17].
2. α -MMS for $\alpha > 1/2$: Even $3/4$ -MMS allocations are guaranteed to exist [24; 25].
3. α -GMMS (groupwise maximin share): GMMS implies MMS, and $1/2$ -GMMS allocations are guaranteed to exist [7].
4. EFX (envy-freeness up to any good): EFX is a strengthening of EF1 [17]. While it is an open question whether EFX allocations always exist with more than three agents (see [18] for the case of three agents), one can still define its price by excluding instances (if any) where such allocations do not exist.

Another direction would be to analyze the price of fairness when the items are *chores* [6], or a mixture of goods and chores [5]. More broadly, while formal definitions of fairness and their corresponding prices are relatively well-understood within fair division theory, they are much less explored within other paradigms of social choice theory such as voting or matching. Developing a comprehensive understanding of what fairness means and what its price is requires significant future work.

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 Transactions on 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: International Conference on Web and Internet Economics, pp. 57–70, Springer (2019)
- [4] Aumann, Y., Dombb, Y.: The efficiency of fair division with connected pieces. *ACM Transactions on Economics and Computation (TEAC)* **3**(4), 1–16 (2015)
- [5] Aziz, H., Caragiannis, I., Igarashi, A., Walsh, T.: Fair allocation of indivisible goods and chores. In: IJCAI, pp. 53–59 (2019)
- [6] Aziz, H., Rauchecker, G., Schryen, G., Walsh, T.: Algorithms for max-min share fair allocation of indivisible chores. In: Thirty-First AAAI Conference on Artificial Intelligence (2017)
- [7] Barman, S., Biswas, A., Krishnamurthy, S.K., Narahari, Y.: Groupwise maximin fair allocation of indivisible goods. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)
- [8] Barman, S., Krishnamurthy, S.K.: On the proximity of markets with integral equilibria. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, pp. 1748–1755 (2019)
- [9] Barman, S., Krishnamurthy, S.K., Vaish, R.: Finding fair and efficient allocations. In: Proceedings of the 2018 ACM Conference on Economics and Computation, pp. 557–574 (2018)
- [10] 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, August 10-16, 2019, pp. 81–87 (2019)
- [11] Bertsimas, D., Farias, V.F., Trichakis, N.: The price of fairness. *Operations Research* **59**(1), 17–31 (2011)
- [12] Bogomolnaia, A., Moulin, H.: Random matching under dichotomous preferences. *Econometrica* **72**, 257–279 (2004)
- [13] Bollobás, B.: *Modern Graph Theory*, Graduate Texts in Mathematics, vol. 184. Springer (2002)
- [14] Budish, E.: The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* **119**(6), 1061–1103 (2011)
- [15] 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)
- [16] Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M.: The efficiency of fair division. *Theory of Computing Systems* **50**(4), 589–610 (2012)
- [17] Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)* **7**(3), 1–32 (2019)
- [18] Chaudhury, B.R., Garg, J., Mehlhorn, K.: EFX exists for three agents. In: EC '20, pp. 1–19, ACM (2020)
- [19] Dickerson, J.P., Procaccia, A.D., Sandholm, T.: Price of fairness in kidney exchange. In: Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), pp. 1013–1020 (2014)

- [20] Dobzinski, S., Nisan, N., Schapira, M.: Approximation algorithms for combinatorial auctions with complement-free bidders. *Mathematics of Operations Research* **35**(1), 1–13 (2010)
- [21] Dutta, B., Peters, H., Sen, A.: Strategy-proof cardinal decision schemes. *Social Choice and Welfare* **28**(1), 163–179 (2007)
- [22] Feige, U.: On maximizing welfare when utility functions are subadditive. *SIAM Journal on Computing* **39**(1), 122–142 (2009)
- [23] Foley, D.: Resource allocation and the public sector. *Yale Economics Essays* **7**, 45–98 (1967)
- [24] Garg, J., Taki, S.: An improved approximation algorithm for maximin shares. In: *EC '20*, pp. 379–380, ACM (2020)
- [25] 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)
- [26] Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: *Annual Symposium on Theoretical Aspects of Computer Science*, pp. 404–413, Springer (1999)
- [27] Kurokawa, D., Procaccia, A.D., Shah, N.: Leximin allocations in the real world. *ACM Transactions on Economics and Computation (TEAC)* **6**(3-4), 1–24 (2018)
- [28] 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)
- [29] Lehmann, B., Lehmann, D., Nisan, N.: Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* **55**(2), 270–296 (2006)
- [30] 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)
- [31] 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)
- [32] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A.: A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635* (2019)
- [33] 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)
- [34] Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.): *Algorithmic Game Theory*. Cambridge University Press (2007)
- [35] 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)
- [36] Rawls, J.: *A Theory of Justice*. Harvard University Press (1971)
- [37] Suksompong, W.: Fairly allocating contiguous blocks of indivisible items. *Discrete Applied Mathematics* **260**, 227–236 (2019)
- [38] Woeginger, G.J.: A polynomial-time approximation scheme for maximizing the minimum machine completion time. *Operations Research Letters* **20**(4), 149–154 (1997)

Appendix

A Price of Fairness with Supermodular Valuations

Let us consider two more classes of agent valuations.

- Superadditive: $v_i(S \cup T) \geq v_i(S) + v_i(T)$ for each agent $i \in [n]$ and all disjoint subsets $S, T \subseteq [m]$.
- Supermodular: $v_i(S \cup T) + v_i(S \cap T) \geq v_i(S) + v_i(T)$ for each agent $i \in [n]$ and all subsets $S, T \subseteq [m]$.

Clearly, if agent valuations are supermodular, then they are also superadditive. Given that we establish the price of EF1 for subadditive valuations (Theorem 1), one might wonder what the price of EF1 is for the complementary class of superadditive valuations. We show that even if the valuations are supermodular and identical across agents, the price of EF1 is unbounded in terms of n and m .

Example 1. Suppose there are n agents with the following identical valuation function v over $m = n$ goods. For all $S \subseteq [m]$,

$$v(S) = \begin{cases} \epsilon |S|, & \text{if } |S| \leq 1, \\ \epsilon + (|S| - 1) \cdot L, & \text{if } |S| \geq 1. \end{cases}$$

Here, ϵ is very small, and $L = (1 - \epsilon)/(m - 1)$. It can be checked that the valuation function is supermodular, and that $v([m]) = 1$ (and hence valuations are scaled).

The optimal allocation assigns all goods to a single agent, for a social welfare of 1. However in an EF1 allocation, each agent must receive one good, and hence the maximum social welfare that can be achieved by any EF1 allocation is $n\epsilon$. Hence, the price of EF1 is at least $1/(n\epsilon)$. Since ϵ can be arbitrarily small, the price of EF1 is unbounded in terms of n and m even for identical supermodular valuations.

B Lower Bounds for the Price of Proportionality Up To One Good

In this section, we show how to modify the construction from the proof of Theorem 1 and that due to Bei et al. [10] to derive that the price of Prop1 is $\Omega(n)$ for unscaled additive valuations and $\Omega(\sqrt{n})$ for scaled additive valuations. In each case, the key modification is to have $n + 1$ goods instead of n goods.

Example 2 (Unscaled Additive Valuations). Consider an instance with n agents and $m = n + 1$ goods. Suppose the agents have the following additive valuations over the goods: the first agent has value $n + 1$ for each good and every other agent has value $1/(n + 1)$ for each good. The optimal social welfare is achieved by allocating all goods to the first agent, yielding $\text{OPT} = (n + 1)^2$. In contrast, it is easy to see that any Prop1 allocation must give at least one good to each of agents $2, \dots, n$. Hence, any Prop1 allocation has welfare less than $2n + 3$. This implies the desired $\Omega(n)$ bound on the price of Prop1 under unscaled additive valuations.

Example 3 (Scaled Additive Valuations). Consider an instance with n agents and $m = n + 1$ goods. For the ease of exposition, suppose that \sqrt{n} is an integer. For $k \in [\sqrt{n}]$, define $T_k = \{(k - 1)\sqrt{n} + 1, \dots, k\sqrt{n}\}$. Note that $T_1, \dots, T_{\sqrt{n}}$ partition goods $1, \dots, n$.

Suppose the agents have the following scaled additive valuations.

- For $i \leq \sqrt{n}$, agent i has value $1/\sqrt{n}$ for each of \sqrt{n} goods in T_i , and value 0 for every other good.
- For $i > \sqrt{n}$, agent i has value $1/(n + 1)$ for each of $n + 1$ goods.

The optimal social welfare is achieved by partitioning the first n goods among the first \sqrt{n} agents, and giving the last good to a later agent, yielding $\text{OPT} = \sqrt{n} + \frac{1}{n+1} = \Omega(\sqrt{n})$.

In contrast, any Prop1 allocation must give agent i at least one good for each $i > \sqrt{n}$. Thus, any Prop1 allocation has welfare at most $(\sqrt{n} + 1) \cdot \frac{1}{\sqrt{n}} + (n - \sqrt{n}) \cdot \frac{1}{n+1} = O(1)$. This implies the desired $\Omega(\sqrt{n})$ bound on the price of Prop1 under scaled additive valuations.