Finding Fair and Efficient Allocations

SIDDHARTH BARMAN, Indian Institute of Science, India
SANATH KUMAR KRISHNAMURTHY, Chennai Mathematical Institute, India
ROHIT VAISH, Indian Institute of Science, India

We study the problem of allocating a set of *indivisible* goods among a set of agents in a *fair* and *efficient* manner. An allocation is said to be *fair* if it is envy-free up to one good (EF1), which means that each agent prefers its own bundle over the bundle of any other agent up to the removal of one good. In addition, an allocation is deemed *efficient* if it satisfies Pareto efficiency. While each of these well-studied properties is easy to achieve separately, achieving them together is far from obvious. Recently, Caragiannis et al. [11] established the surprising result that when agents have additive valuations for the goods, there always exists an allocation that simultaneously satisfies these two seemingly incompatible properties. Specifically, they showed that an allocation that maximizes the Nash social welfare objective is both EF1 and Pareto efficient. However, the problem of maximizing Nash social welfare is NP-hard. As a result, this approach does not provide an efficient algorithm for finding a fair and efficient allocation.

In this paper, we bypass this barrier, and develop a pseudopolynomial time algorithm for finding allocations that are EF1 and Pareto efficient; in particular, when the valuations are bounded, our algorithm finds such an allocation in polynomial time. Furthermore, we establish a stronger existence result compared to Caragiannis et al. [11]: For additive valuations, there always exists an allocation that is EF1 and *fractionally* Pareto efficient.

Another key contribution of our work is to show that our algorithm provides a polynomial-time 1.45-approximation to the Nash social welfare objective. This improves upon the best known approximation ratio for this problem (namely, the 2-approximation algorithm of Cole et al. [12]), and also matches the lower bound on the integrality gap of the convex program of Cole et al. [12]. Unlike many of the existing approaches, our algorithm is completely combinatorial, and relies on constructing *integral* Fisher markets wherein specific equilibria are not only efficient, but also fair.

CCS Concepts: • **Theory of computation → Approximation algorithms analysis**; **Market equilibria**; *Algorithmic game theory*; • **Mathematics of computing** → *Combinatorial algorithms*;

Additional Key Words and Phrases: Fair division, Nash Social Welfare, Approximation Algorithms

1 INTRODUCTION

The theory of fair division addresses the fundamental problem of allocating goods or resources among agents in a fair and efficient manner. Such problems arise in many real-world settings such as government auctions, divorce settlements, and border disputes. Starting with the work of Steinhaus [30], there is now a vast literature in economics and mathematics to formally address fair division [7, 8, 26]. Many interesting connections have also been found between fair division and fields such as topology, measure theory, combinatorics, and algorithms [25].

Authors' addresses: Siddharth Barman, Indian Institute of Science, Bangalore, India, 560012, barman@iisc.ac.in; Sanath Kumar Krishnamurthy, Chennai Mathematical Institute, Kelambakkam, India, 603103, sanathkumar9@cmi.ac.in; Rohit Vaish, Indian Institute of Science, Bangalore, India, 560012, rohitv@iisc.ac.in.

ACM EC'18, June 18–22, 2018, Ithaca, NY, USA. ACM ISBN 978-1-4503-5829-3/18/06....$15.00.
https://doi.org/10.1145/3219166.3219176

Much of the prior work in fair division, though, has focused on *divisible* goods, which model resources that can be fractionally allocated (such as land). A standard fairness concept in this setting is *envy-freeness* [17], which requires that each agent prefers its own allocation over that of any other agent. A well-known result of Varian [31] shows that for the divisible setting, there always exists an allocation that is both envy-free (i.e., fair) and *Pareto efficient*. Furthermore, such an allocation can be computed in polynomial time [15, 16]. These results, however, do not extend to the setting of *indivisible* goods, which model discrete resources such as courses at universities [29] or inherited artwork. In fact, many of the classical solution concepts and algorithms that have been developed for divisible goods are not directly applicable to the indivisible setting. For example, an envy-free allocation fails to exist even in the simple setting of a single indivisible good and two agents.

These considerations have motivated recent work in the theoretical computer science and economics communities on developing relevant notions of fairness, along with existence results and algorithms for the problem of fairly allocating indivisible goods [5, 9, 21, 23]. We contribute to this line of work by showing that guarantees analogous to the fundamental result of Varian [31] hold even for indivisible goods in terms of a natural and necessary relaxation of envy-freeness. Specifically, we show that for additive valuations,[1] a fair and efficient allocation always exists, and such an allocation can be computed in (pseudo)-polynomial time.

We consider an allocation of indivisible goods to be fair if it is *envy-free up to one good* (EF1). This notion was defined by Budish [9], and provides a compelling relaxation of the envy-freeness property.[2] An allocation is said to be EF1 if each agent prefers its own bundle over the bundle of any other agent up to the removal of the most valuable good from the other agent's bundle. Although the existence of envy-free allocations is not guaranteed in the context of indivisible goods, an EF1 allocation always exists—even under general, combinatorial valuations—and can be found in polynomial time [23].

With this notion of fairness in hand, it is relevant to ask whether we can achieve *efficiency along with fairness* while allocating indivisible goods.[3] This question was recently studied by Caragiannis et al. [11], who showed a striking result that there is no need to trade efficiency for fairness: For additive valuations, an allocation that maximizes the *Nash social welfare* [20, 27]—defined to be the geometric mean of the agents' valuations—is both fair (EF1) and Pareto efficient. However, maximizing the Nash social welfare (NSW) over integral allocations is an NP-hard problem [28]. (In fact, the problem is known to be APX-hard [22]). Therefore, this existence result does not automatically provide an efficient algorithm for finding a fair and efficient allocation of indivisible goods. Our work bypasses this limitation by providing *a pseudopolynomial time algorithm for finding an EF1 and Pareto efficient allocation* of indivisible goods under additive valuations. In particular, when the valuations are *bounded*, our algorithm finds such an allocation in *polynomial time*. It is worth pointing out that the problem of maximizing NSW remains APX-hard even for bounded valuations [22].

A related problem is that of developing approximation algorithms for NSW maximization. This problem has received considerable attention in recent years [1, 2, 4, 12, 13, 18]. The first constant-factor (specifically, 2.89) approximation for this problem was provided by Cole and Gkatzelis [13]. This approximation factor was subsequently improved to e [2], and most recently to 2 [12]. Similar

[1] *Additivity* means that an agent's valuation for a set of goods is the sum of its valuations for the individual goods in that set.
[2] The notion of EF1 has found practical appeal on the popular fair division website "Spliddit" [19] and in course allocation at Wharton School of Business [10].
[3] Note that fairness, by itself, does not guarantee efficiency; in fact, an EF1 allocation can be highly inefficient (refer to the full version [3] for an example).

approximation guarantees have also been developed for more general market models such as piecewise-linear concave utilities [1], budget additive valuations [18], and multi-unit markets [4].

While the problem of approximating NSW is interesting in its own right, it is relevant to note that an allocation that approximates this objective is, in and of itself, not guaranteed to be EF1 or Pareto efficient (refer to the full version [3] for an example).[4] A second key contribution of our work is to show that our algorithm provides *a polynomial-time* 1.45*-approximation to the* NSW *maximization problem*. Thus, not only does our algorithm improve upon the best-known approximation ratio for this problem (namely, the 2-approximation algorithm of Cole et al. [12]), it is also guaranteed to return a fair and efficient outcome. The following list summarizes our contributions.

Our contributions

- We develop an *algorithm* for computing an EF1 and Pareto efficient allocation for additive valuations. The running time of our algorithm is pseudopolynomial for general integral valuations (Theorem 3.1) and polynomial when the valuations are bounded (Remark 1). In addition, our algorithm can find an approximate EF1 and approximate Pareto efficient allocation in polynomial time even without the bounded valuations assumption (Remark 2).
- We establish a *stronger existence result* compared to Caragiannis et al. [11]: For additive valuations, there always exists an allocation that is EF1 and *fractionally* Pareto efficient (Theorem 3.2). In other words, the problem of finding an EF1 and fractionally Pareto efficient allocation is *total*. An interesting complexity-theoretic implication of this result is that there exists a nondeterministic polynomial time algorithm for finding an EF1 and Pareto efficient allocation (Remark 3). This implication does not directly follow from the existence result of Caragiannis et al. [11], as the problem of verifying whether an arbitrary allocation is Pareto efficient is known to be co-NP-complete [14].
- We show that our algorithm provides a *polynomial-time* 1.45*-approximation* for the Nash social welfare (NSW) maximization problem (Theorem 3.3). This improves upon the best known approximation factor for this problem (namely, the 2-approximation algorithm of Cole et al. [12]), and also matches the lower bound of $e^{1/e} \approx 1.44$ on the integrality gap of the convex program of Cole et al. [12]. An interesting byproduct of our analysis is a novel connection between envy-freeness and NSW: Under identical valuations, an EF1 allocation provides a 1.45-approximation to the maximum NSW (Lemma 3.4).

Our techniques It is known from the fundamental theorems of welfare economics that markets tend toward efficiency. Intuitively, our results are based on establishing a complementary result that markets can be *fair* as well. In particular, we construct a Fisher market along with an underlying equilibrium which is integral (i.e., corresponding to an allocation of the indivisible goods) and EF1. The fact that this allocation is a market equilibrium ensures, via the first welfare theorem, that it is Pareto efficient as well.

More concretely, we start with a Pareto efficient allocation, and iteratively modify the allocation by exchanging goods between the agents. The goal of the exchange step is to locally move toward a fair allocation. Additionally, throughout these exchanges, we maintain a set of prices that ensure that the current allocation corresponds to an equilibrium outcome for the existing market. We stop when the equilibrium of the market (i.e., the allocation at hand) satisfies *price envy-freeness up to one good* (refer to Section 4.1 for a formal definition). Essentially, this property ensures that under the given market prices, the spending of an agent is at least that of any other agent up to the removal of the highest priced good from the other agent's bundle. Requiring the spendings to

[4]We also provide an example in the full version [3] in which *every* rounding of the "spending restricted outcome"—a market equilibrium notion used in the design of approximation algorithms for NSW [1, 12, 13]—violates EF1.

be balanced in this manner implies the desired EF1 property for the corresponding fair division instance; see Section 4 for a detailed description of this construction.

At a conceptual level, our approach differs from the existing approaches in two important ways: First, our algorithm works with an *integral* Fisher market at every step, thereby breaking away from the standard *relax-and-round* paradigm where a fractional market equilibrium is first computed (typically as a solution of some convex program) followed by a rounding step [1, 2, 12, 13, 18]. Second, unlike all existing approaches, our algorithm uses the notion of *price envy-freeness up to one good* as a measure of balanced spending in the Fisher market. To the best of our knowledge, this notion is novel to this work, and might find future use in the design of fair and efficient algorithms for other settings.

2 PRELIMINARIES

2.1 The Fair Division Model

Problem instance An *instance* of the fair division problem is a tuple $\langle [n], [m], \mathcal{V} \rangle$, where $[n] = \{1, 2, \ldots, n\}$ denotes the set of $n \in \mathbb{N}$ *agents*, $[m] = \{1, 2, \ldots, m\}$ denotes the set of $m \in \mathbb{N}$ *goods*, and the *valuation profile* $\mathcal{V} = \{v_1, v_2, \ldots, v_n\}$ specifies the preferences of each agent $i \in [n]$ over the set of goods $[m]$ via a *valuation function* $v_i : 2^{[m]} \to \mathbb{Z}_+ \cup \{0\}$. We will assume throughout that the valuation functions are *additive*, i.e., for each agent $i \in [n]$ and any set of goods $G \subseteq [m]$, $v_i(G) := \sum_{j \in G} v_i(\{j\})$.[5] For simplicity, we will write $v_{i,j}$ instead of $v_i(\{j\})$ for a singleton good $j \in [m]$. Thus, $v_{i,j}$ is non-negative and integral for each agent $i \in [n]$ and each good $j \in [m]$. We will also assume, without loss of generality, that for each good $j \in [m]$, there exists some agent $i \in [n]$ with a nonzero valuation for it, i.e., $v_{i,j} > 0$. Finally, we let $v_{\max} := \max_{i,j} v_{i,j}$.

Allocation An *allocation* $\mathbf{x} \in \{0, 1\}^{n \times m}$ refers to an n-partition $(\mathbf{x}_1, \ldots, \mathbf{x}_n)$ of $[m]$, where $\mathbf{x}_i \subseteq [m]$ is the *bundle* allocated to agent i. We let $\mathcal{X}$ denote the set of all n partitions of $[m]$. Given an allocation $\mathbf{x}$, the valuation of an agent $i \in [n]$ for the bundle $\mathbf{x}_i$ is $v_i(\mathbf{x}_i) = \sum_{j \in \mathbf{x}_i} v_{i,j}$.

Another useful notion is that of a fractional allocation. A *fractional allocation* $\mathbf{x} \in [0, 1]^{n \times m}$ refers to a (possibly) fractional assignment of the goods to the agents such that no more than one unit of each good is allocated, i.e., for all $j \in [m]$, we have $\sum_{i \in [n]} x_{i,j} \leq 1$. We will use the term *allocation* to refer to an integral allocation, and explicitly write *fractional allocation* otherwise.

2.2 Fairness Notions

Envy-freeness and its variants Given an instance $\langle [n], [m], \mathcal{V} \rangle$ and an allocation $\mathbf{x}$, we say that an agent $i \in [n]$ *envies* another agent $k \in [n]$ if i strictly prefers the bundle of k over its own bundle, i.e., $v_i(\mathbf{x}_k) > v_i(\mathbf{x}_i)$. An allocation $\mathbf{x}$ is said to be *envy-free* (EF) if each agent prefers its own bundle over that of any other agent, i.e., for every pair of agents $i, k \in [n]$, we have $v_i(\mathbf{x}_i) \geq v_i(\mathbf{x}_k)$.

An allocation $\mathbf{x}$ is said to be *envy-free up to one good* (EF1) if for every pair of agents $i, k \in [n]$, there exists a good $j \in \mathbf{x}_k$ such that $v_i(\mathbf{x}_i) \geq v_i(\mathbf{x}_k \setminus \{j\})$. Given any $\varepsilon > 0$, an allocation $\mathbf{x}$ is said to be *ε-approximately envy-free up to one good* (ε-EF1) if for every pair of agents $i, k \in [n]$, there exists a good $j \in \mathbf{x}_k$ such that $(1 + \varepsilon) v_i(\mathbf{x}_i) \geq v_i(\mathbf{x}_k \setminus \{j\})$.

Nash social welfare Given an allocation $\mathbf{x}$, write $\mathrm{NSW}(\mathbf{x}) := \left(\prod_{i \in [n]} v_i(\mathbf{x}_i) \right)^{\frac{1}{n}}$ to denote the *Nash social welfare* of $\mathbf{x}$. An allocation $\mathbf{x}^*$ said to be *Nash optimal* if $\mathbf{x}^* \in \arg\max_{\mathbf{x} \in \mathcal{X}} \mathrm{NSW}(\mathbf{x})$.

2.3 Efficiency Notions

Pareto efficiency Given an instance $\langle [n], [m], \mathcal{V} \rangle$ and an allocation $\mathbf{x}$, we say that $\mathbf{x}$ is Pareto dominated by another allocation $\mathbf{y}$ if $v_k(\mathbf{y}_k) \geq v_k(\mathbf{x}_k)$ for every agent $k \in [n]$, and $v_i(\mathbf{y}_i) > v_i(\mathbf{x}_i)$

[5] We will assume that $v_i(\{\emptyset\}) = 0$ for all $i \in [n]$.

for some agent $i \in [n]$. An allocation is said to be *Pareto efficient* or *Pareto optimal* (PO) if it is not Pareto dominated by any other allocation. Similarly, $\mathbf{x}$ is *ε-Pareto efficient* (ε-PO) if it is not ε-Pareto dominated by any other allocation $\mathbf{y}$, i.e., there does not exist an allocation $\mathbf{y}$ such that $v_k(\mathbf{y}_k) \geq (1+\varepsilon)v_k(\mathbf{x}_k)$ for every agent $k \in [n]$ and $v_i(\mathbf{y}_i) > (1+\varepsilon)v_i(\mathbf{x}_i)$ for some agent $i \in [n]$.

Some of our results use a generalization of Pareto efficiency, which we call *fractional Pareto efficiency*. An allocation is said to be fractionally Pareto efficient (fPO) if it not Pareto dominated by any fractional allocation. Thus, a fractionally Pareto efficient allocation is also Pareto efficient, but the converse is not necessarily true (refer to the full version [3] for an example).

3 MAIN RESULTS

This section provides the statements of our three main results: an algorithm for finding an EF1 and PO allocation (Theorem 3.1), an existence result for EF1 and fPO allocation (Theorem 3.2), and an approximation algorithm for Nash social welfare (Theorem 3.3).

Algorithmic Result:

THEOREM 3.1. *Given any fair division instance $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$ with additive valuations, an allocation that is envy-free up to one good (EF1) and Pareto efficient (PO) can be found in $O\,(\mathrm{poly}(m, n, v_{\max}))$ time, where $v_{\max} = \max_{i,j} v_{i,j}$.*

REMARK 1. *Note that when all valuations are* polynomially bounded *(i.e., there exists a polynomial $f(m,n)$ such that for all $i \in [n]$ and $j \in [m]$, $v_{i,j} \leq f(m,n)$), an EF1 and PO allocation can be computed in* polynomial *time. In particular, this is true when all valuations are bounded by a constant. As mentioned earlier in Section 1, the problem of maximizing* NSW *remains APX-hard even for constant valuations [22], and therefore our result circumvents the intractability associated with computing a Nash optimal allocation in order to achieve these two properties.*

REMARK 2. *If we relax the fairness and efficiency requirements in Theorem 3.1 to their approximate analogues, then our algorithm is guaranteed to run in polynomial time. Specifically, our algorithm can find an ε-EF1 and ε-PO allocation in $O\left(\mathrm{poly}(m, n, \frac{1}{\varepsilon}, \ln v_{\max})\right)$ time, where $v_{\max} = \max_{i,j} v_{i,j}$. (Refer to Lemma 5.5 in Section 5.2).*

The proof of Theorem 3.1 is provided in Section 5.

Existence Result:

THEOREM 3.2. *Given any fair division instance with additive valuations, there always exists an allocation that is envy-free up to one good (EF1) and fractionally Pareto efficient (fPO).*

REMARK 3. *Consider the canonical binary relation $\mathcal{R}^{EF1+PO}$ associated with the problem of finding an EF1 and PO allocation, defined as follows: For a fair division instance $\mathcal{I}$ and an allocation $\mathbf{x}$, the relation $\mathcal{R}^{EF1+PO}(\mathcal{I}, \mathbf{x})$ holds if and only if $\mathbf{x}$ is an EF1 and PO allocation of $\mathcal{I}$. It is relevant to note that under standard complexity theoretic assumptions, $\mathcal{R}^{EF1+PO}$ is not in* TFNP.[6] *By contrast, the binary relation $\mathcal{R}^{EF1+fPO}(\mathcal{I}, \mathbf{x})$, which holds if and only if $\mathbf{x}$ is an EF1 and fPO allocation for the instance $\mathcal{I}$, admits efficient verification.[7] Since Theorem 3.2 shows that $\mathcal{R}^{EF1+fPO}$ is total, we get that the binary relation $\mathcal{R}^{EF1+fPO}$ is in* TFNP. *Thus, there exists a nondeterministic polynomial time algorithm for finding an EF1 and fPO (and hence EF1 and PO) allocation.*

The proof of Theorem 3.2 is deferred to the full version of the paper.

[6] It is known that determining whether an arbitrary allocation is PO is co-NP-complete [14]. This fact can be used to show that verifying whether a given allocation is EF1 and PO is also co-NP-complete. Hence, the binary relation $\mathcal{R}^{\mathrm{EF1+PO}}$ cannot be efficiently verified (i.e., it is not in FNP), unless P = NP.

[7] EF1 requires checking $O(n^2)$ inequalities, and fPO can be verified by a linear program (refer to the full version).

Approximating Nash Social Welfare:

Theorem 3.3. *For additive valuations, there exists a polynomial-time* 1.45*-approximation algorithm for the Nash social welfare maximization problem.*

Our proof of Theorem 3.3 draws on the following interesting connection between approximate envy-freeness and Nash social welfare:

Lemma 3.4. *Given a fair division instance with identical and additive valuations, any ε-EF1 allocation provides a $e^{(1+\varepsilon)/e}$-approximation to Nash social welfare.*

The proof of Theorem 3.3 is provided in Section 6.

4 OUR ALGORITHM

This section presents our algorithm. We start with the relevant preliminaries in Section 4.1 that provide the necessary definitions required for describing the algorithm. The pseudocode of the algorithm appears in Section 4.2 along with a brief description.

4.1 Market Terminology

Fisher market The Fisher market is a fundamental model in the economics of resource allocation [6]. It captures the setting where a set of buyers enter the market with prespecified budgets, and use it to buy goods that provide maximum utility per unit of money spent. Specifically, a Fisher market consists of a set $[n] = \{1, 2, \ldots, n\}$ of n *buyers*, a set $[m] = \{1, 2, \ldots, m\}$ of m divisible *goods* (exactly one unit of each good is available), and a *valuation profile* $\mathcal{V} = \{v_1, v_2, \ldots, v_n\}$. Each buyer $i \in [n]$ has an initial *endowment* (or budget) $e_i > 0$. The endowment holds no intrinsic value for a buyer and is only used for buying the goods. We call $\mathbf{e} = (e_1, \ldots, e_n)$ the *endowment vector*, and denote a *market instance* by $\langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$.

A *market outcome* is given by the pair $\langle \mathbf{x}, \mathbf{p} \rangle$, where the *allocation vector* $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_n)$ is a fractional allocation of the m goods, and the *price vector* $\mathbf{p} = (p_1, \ldots, p_m)$ associates a price $p_j \geq 0$ with each good $j \in [m]$. The *spending* of buyer i under the market outcome $\langle \mathbf{x}, \mathbf{p} \rangle$ is given by $\mathbf{p}(\mathbf{x}_i) = \sum_{j=1}^{m} x_{i,j} p_j$. The *valuation* derived by the buyer i under the market outcome $\langle \mathbf{x}, \mathbf{p} \rangle$ is given by $v_i(\mathbf{x}_i) = \sum_{j=1}^{m} x_{i,j} v_{i,j}$.

Given a price vector $\mathbf{p} = (p_1, \ldots, p_m)$, define the *bang per buck* ratio of buyer i for good j as $\alpha_{i,j} := v_{i,j}/p_j$, and its *maximum bang per buck* ratio as $\alpha_i := \max_j \alpha_{i,j}$.[8] Let $\text{MBB}_i := \{j \in [m] : v_{i,j}/p_j = \alpha_i\}$ denote the set of all goods that maximize the bang per buck ratio for buyer i at the price vector $\mathbf{p}$. We call MBB_i the *maximum bang per buck set* (or MBB set) of buyer i at the price vector $\mathbf{p}$.

An outcome $\langle \mathbf{x}, \mathbf{p} \rangle$ is said to be a *Fisher market equilibrium* if it satisfies the following conditions:

- *Market clearing*: Each good is either priced at zero or is completely allocated. That is, for each good $j \in [m]$, either $p_j = 0$ or $\sum_{i=1}^{n} x_{i,j} = 1$.
- *Budget exhaustion*: Buyers spend their endowments completely, i.e., $\mathbf{p}(\mathbf{x}_i) = e_i$ for all $i \in [n]$.
- *Maximum bang per buck allocation*: Each buyer's allocation is a subset of its MBB set. That is, for any buyer $i \in [n]$ and any good $j \in [m]$, $x_{i,j} > 0 \implies j \in \text{MBB}_i$. Stated differently, each buyer only spends on its maximum bang per buck goods. Notice that a consequence of spending only on MBB goods is that each buyer maximizes its utility at the given prices $\mathbf{p}$ under the budget constraints.

We refer the reader to the full version for additional market preliminaries.

[8] If $v_{i,j} = 0$ and $p_j = 0$, then we define $\alpha_{i,j} = 0$.

Proposition 1 (First Welfare Theorem; [24, Chapter 16]). *For a Fisher market with additive valuations, any equilibrium outcome is fractionally Pareto efficient (fPO).*

Price envy-freeness and its variants Several of our results rely on constructing market outcomes with a property called *price envy-freeness*—a notion we consider to be of independent interest. Specifically, let $\mathbf{x}$ be an allocation and let $\mathbf{p}$ be a price vector for a given Fisher market. We say that $\mathbf{x}$ is *price envy-free* (pEF) with respect to $\mathbf{p}$ if for every pair of buyers $i, k \in [n]$, we have $\mathbf{p}(\mathbf{x}_i) \geq \mathbf{p}(\mathbf{x}_k)$.[9] Similarly, $\mathbf{x}$ is said to be *price envy-free up to one good* (pEF1) with respect to $\mathbf{p}$ if for every pair of buyers $i, k \in [n]$, there exists a good $j \in \mathbf{x}_k$ such that $\mathbf{p}(\mathbf{x}_i) \geq \mathbf{p}\left(\mathbf{x}_k \setminus \{j\}\right)$. Finally, given any $\varepsilon > 0$, we say that an allocation $\mathbf{x}$ is *ε-approximately price envy-free up to one good* (ε-pEF1) with respect to $\mathbf{p}$ if for every pair of buyers $i, k \in [n]$, there exists a good $j \in \mathbf{x}_k$ such that $(1+\varepsilon)\mathbf{p}(\mathbf{x}_i) \geq \mathbf{p}(\mathbf{x}_k \setminus \{j\})$.

MBB **graph and alternating paths** The MBB *graph* of a Fisher market instance with a price vector $\mathbf{p}$ is defined as a bipartite graph G whose vertex set consists of the set of agents $[n]$ and the set of goods $[m]$, and there is an edge between an agent $i \in [n]$ and a good $j \in [m]$ if $j \in \text{MBB}_i$ (called an MBB *edge*). Given an allocation $\mathbf{x}$, we can augment the MBB graph by adding *allocation edges*, i.e., an edge between an agent $i \in [n]$ and a good $j \in [m]$ such that $j \in \mathbf{x}_i$. For an augmented MBB graph, we define an *alternating path* $P = (i, j_1, i_1, j_2, i_2, \ldots, i_{\ell-1}, j_\ell, k)$ from agent i to agent k (and involving the agents $i_1, i_2, \ldots, i_{\ell-1}$ and the goods $j_1, j_2, \ldots, j_\ell$) as a series of alternating MBB and allocation edges such that $j_1 \in \text{MBB}_i \cap \mathbf{x}_{i_1}$, $j_2 \in \text{MBB}_{i_1} \cap \mathbf{x}_{i_2}, \ldots, j_\ell \in \text{MBB}_{i_{\ell-1}} \cap \mathbf{x}_k$. If such a path exists, we say that the agent k is *reachable* from agent i via an alternating path. Notice that no agent or good is allowed to repeat in an alternating path. We say that the path P is of *length* 2ℓ since it consists of ℓ MBB edges and ℓ allocation edges.

Hierarchy structure Let G denote the augmented MBB graph for a Fisher market instance with the market outcome $(\mathbf{x}, \mathbf{p})$. Fix a *source* agent $i \in [n]$ in G. Define the *level* of an agent $k \in [n]$ as half the length of the shortest alternating path from i to k (if one exists). The level of the source agent i is defined to be zero. If there is no alternating path from i to some agent k in G (i.e., if k is not reachable from i), then the level of k is set to be n. The *hierarchy* structure $\mathcal{H}_i$ of agent i is defined as a level-wise collection of all agents that are reachable from i, i.e., $\mathcal{H}_i = \{\mathcal{H}_i^0, \mathcal{H}_i^1, \mathcal{H}_i^2, \ldots, \}$, where $\mathcal{H}_i^\ell$ denotes the set of agents that are at level ℓ with respect to the agent i. The full version provides a polynomial time subroutine called BuildHierarchy for constructing the hierarchy.

Given a hierarchy $\mathcal{H}_i$, we will overload the term *alternating path* to refer to a series of alternating MBB and allocation edges *connecting agents at a lower level to those at a higher level.* That is, a path $P = (i, j_1, i_1, j_2, i_2, \ldots, i_{\ell-1}, j_\ell, k)$ involving agents from the hierarchy $\mathcal{H}_i$ is said to be an alternating path if (1) $j_1 \in \text{MBB}_i \cap \mathbf{x}_{i_1}$, $j_2 \in \text{MBB}_{i_1} \cap \mathbf{x}_{i_2}, \ldots, j_\ell \in \text{MBB}_{i_{\ell-1}} \cap \mathbf{x}_k$, and (2) $\text{level}(i) < \text{level}(i_1) < \text{level}(i_2) < \cdots < \text{level}(i_{\ell-1}) < \text{level}(k)$. In particular, an alternating path in a hierarchy cannot have edges between agents at the same level.

Violators and path-violators Given a Fisher market instance and a market outcome $(\mathbf{x}, \mathbf{p})$, an agent $i \in [n]$ with the smallest spending among all the agents is called the *least spender*, i.e., $i \in \arg\min_{k \in [n]} \mathbf{p}(\mathbf{x}_k)$ (ties are broken according to a prespecified ordering over the agents). An agent $k \in [n]$ is said to be a *violator* if for every good $j \in \mathbf{x}_k$, we have that $\mathbf{p}(\mathbf{x}_k \setminus \{j\}) > \mathbf{p}(\mathbf{x}_i)$, where i is the least spender. Similarly, agent $k \in [n]$ is said to be an *ε-violator* if for every good $j \in \mathbf{x}_k$, we have that $\mathbf{p}(\mathbf{x}_k \setminus \{j\}) > (1+\varepsilon)\mathbf{p}(\mathbf{x}_i)$. Notice that an agent can be a violator without being an *ε-violator.* Also notice that if no agent is a violator (ε-violator), then the allocation $\mathbf{x}$ is pEF1 (ε-pEF1) with respect to $\mathbf{p}$.

[9] Equivalently, for every pair of buyers $i, k \in [n]$, we require $\mathbf{p}(\mathbf{x}_i) = \mathbf{p}(\mathbf{x}_k)$.

A closely related notion is that of a *path-violator*. Let i denote the least spender, and let $\mathcal{H}_i$ denote the hierarchy of agent i. An agent $k \in \mathcal{H}_i$ is said to be a *path-violator* with respect to the alternating path $P = (i, j_1, i_1, j_2, i_2, \ldots, i_{\ell-1}, j_\ell, k)$ if $\mathbf{p}(\mathbf{x}_k \setminus \{j_\ell\}) > \mathbf{p}(\mathbf{x}_i)$. Observe that a path-violator (along a path P) need not be a violator, since there can be a good $j \in \mathbf{x}_k$ not on the path P such that $\mathbf{p}(\mathbf{x}_k \setminus \{j\}) \leq \mathbf{p}(\mathbf{x}_i)$. Similarly, an agent $k \in \mathcal{H}_i$ is said to be an *ε-path-violator* with respect to the alternating path $P = (i, j_1, i_1, j_2, i_2, \ldots, i_{\ell-1}, j_\ell, k)$ if $\mathbf{p}(\mathbf{x}_k \setminus \{j_\ell\}) > (1+\varepsilon)\mathbf{p}(\mathbf{x}_i)$.

4.2 Description of the Algorithm

Given any fair division instance $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$ as input and a parameter $\varepsilon > 0$, our algorithm (Algorithm 1), referred to as Alg from here onwards, constructs a market equilibrium $(\mathbf{x}, \mathbf{p})$ with respect to a Fisher market instance $\langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$ (for a suitable endowment vector $\mathbf{e}$). The pair $(\mathbf{x}, \mathbf{p})$ has the following two properties: (1) $\mathbf{x}$ is an integral allocation, and (2) $\mathbf{x}$ is 3ε-pEF1 with respect to $\mathbf{p}$. The second property allows us to show that the allocation $\mathbf{x}$ is 3ε-EF1 for the corresponding fair division instance $\mathcal{I}$ (see Lemma 4.1 below). Furthermore, by the first welfare theorem (Proposition 1), the allocation $\mathbf{x}$ is also guaranteed to be fractionally Pareto efficient (fPO) for the Fisher market instance, and consequently for the fair division instance $\mathcal{I}$.

Lemma 4.1. *Let $\varepsilon \geq 0$, and let $\mathbf{x}$ and $\mathbf{p}$ be an allocation and a price vector respectively for a market instance $\langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$ such that (1) $\mathbf{x}$ is ε-approximately price-envy-free up to one good (ε-pEF1), and (2) $\mathbf{x}_i \subseteq \mathrm{MBB}_i$ for each buyer $i \in [n]$. Then, $\mathbf{x}$ is ε-approximately envy-free up to one good (ε-EF1) for the associated fair division instance $\langle [n], [m], \mathcal{V} \rangle$.*

Proof. Since $\mathbf{x}$ is ε-pEF1 with respect to the price vector $\mathbf{p}$, for any pair of buyers $i, k \in [n]$, there exists a good $j \in \mathbf{x}_k$ such that $(1+\varepsilon)\mathbf{p}(\mathbf{x}_i) \geq \mathbf{p}(\mathbf{x}_k \setminus \{j\})$. Multiplying both sides by the maximum bang per buck ratio α_i of agent i, we get

$$\begin{aligned} & \alpha_i \cdot (1+\varepsilon)\mathbf{p}(\mathbf{x}_i) \geq \alpha_i \cdot \mathbf{p}(\mathbf{x}_k \setminus \{j\}) \\ \implies \quad & (1+\varepsilon)v_i(\mathbf{x}_i) \geq \alpha_i \cdot \mathbf{p}(\mathbf{x}_k \setminus \{j\}) \quad (\text{since } \mathbf{x}_i \subseteq \mathrm{MBB}_i) \\ \implies \quad & (1+\varepsilon)v_i(\mathbf{x}_i) \geq v_i(\mathbf{x}_k \setminus \{j\}), \end{aligned}$$

which is the ε-EF1 guarantee for the allocation $\mathbf{x}$. □

In order to construct the desired Fisher market equilibrium, our algorithm starts with a welfare-maximizing allocation $\mathbf{x}$ and a price vector $\mathbf{p}$ such that $\mathbf{x}$ is fPO and each agent gets a subset of its MBB goods (this is Phase 1 of Alg). If the allocation $\mathbf{x}$ is 3ε-pEF1 with respect to $\mathbf{p}$, then the algorithm terminates with the output $(\mathbf{x}, \mathbf{p})$. Otherwise, the algorithm proceeds to the next phase.

In Phase 2, the algorithm works with the hierarchy of the least spending agent, and performs a series of exchanges (or swaps) of goods between the agents in the hierarchy (without changing the prices). The swaps are aimed at ensuring that at the end of Phase 2, no agent in the hierarchy is ε-pEF1 envied by the least spender. Furthermore, all exchanges in Phase 2 happen only along the MBB edges, thus maintaining at each stage the condition that $\mathbf{x}$ is an equilibrium allocation, and hence, fPO.

If, at the end of Phase 2, the current allocation $\mathbf{x}$ is still not 3ε-pEF1 with respect to the price vector $\mathbf{p}$, the algorithm moves to Phase 3. This phase consists of uniformly raising the prices of the goods owned by the members of the hierarchy. The prices are raised until either the allocation $\mathbf{x}$ becomes 3ε-pEF1 with respect to the new price vector $\mathbf{p}$, or a new agent gets added to the hierarchy. In the latter case, the algorithm goes back to the start of Phase 2.

It is relevant to note that establishing the time complexity of this algorithm is an intricate task; a priori, it is not even clear whether the algorithm terminates. The stated running time bound is in fact obtained via a number of involved arguments which, in particular, rely on analyzing the spending of the agents in different phases.

5 PROOF OF THEOREM 3.1

This section presents the analysis of our algorithm and a proof of Theorem 3.1. Section 5.1 presents the analysis of our algorithm for valuations that satisfy the power-of-$(1 + \varepsilon)$ property. Section 5.2 extends this analysis to general valuations, culminating in the proof of Theorem 3.1.

5.1 Analysis of ALG when the Valuations are power-of-$(1 + \varepsilon)$

In this section, we will analyze ALG under the assumption that all valuations are *power-of-*$(1+\varepsilon)$, i.e., there exists $\varepsilon > 0$ such that for each agent $i \in [n]$ and each good $j \in [m]$, we have $v_{i,j} \in \{0, (1+\varepsilon)^a\}$ for some natural number a (possibly depending on i and j). We will start by defining the notion of a *time step* that will be useful in the subsequent analysis.

Time steps and events The execution of ALG can be described in terms of the following four *events*: (1) *Swap* operation in Phase 2, (2) *Change* in the identity of least spender in Phase 2, (3) *Price-rise* by a factor of α in Phase 3, and (4) *Termination step*. We use the term *time step* (or simply a *step*) to denote the indexing of any execution of ALG, e.g., ALG might perform a swap operation on the first and second time steps, followed by a price-rise in the third time step, and so on. We will use the phrase "at time step t" to denote the state of the algorithm *before* the event at time step t takes place. Notice that each event stated above runs in polynomial time, and therefore it suffices to analyze the running time of ALG in terms of the total number of events (or time steps).

We will now proceed to analyzing the correctness (Lemma 5.1) and the running time (Lemma 5.2) of ALG for power-of-$(1 + \varepsilon)$ valuations.

LEMMA 5.1 (**CORRECTNESS OF ALG FOR POWER-OF–$(1 + \varepsilon)$ INSTANCE**). *Given any power-of-$(1 + \varepsilon)$ instance as input, the allocation returned by ALG is 3ε-approximately envy-free up to one good (3ε-EF1) and fractionally Pareto efficient (fPO).*

PROOF. Let the output of ALG be $(\mathbf{x}, \mathbf{p})$. The fact that $\mathbf{x}$ is fPO follows from the observation that at each step of the algorithm, the allocation of any agent is a subset of its MBB goods, i.e., at each time step, we have $\mathbf{x}_i \subseteq \text{MBB}_i$ for each agent $i \in [n]$. This is certainly true at the end of Phase 1 by way of setting the prices. In Phase 2, each swap operation only happens along an alternating MBB-allocation edge, which maintains the MBB condition. Phase 3 involves raising the prices of the goods owned by the members of the hierarchy $\mathcal{H}_i$ without changing the allocation. We will argue that for each agent $k \in [n]$, if $\mathbf{x}_k \subseteq \text{MBB}_k$ before the price-rise, then the same continues to hold after the price-rise. Indeed, for any agent $k \notin \mathcal{H}_i$, we have $\mathbf{x}_k \cap \mathbf{x}_{\mathcal{H}_i} = \emptyset$. As a result, raising the prices of the goods in $\mathbf{x}_{\mathcal{H}_i}$ does not affect the bang per buck ratio of agent k for the goods in $\mathbf{x}_k$ (and can only reduce its bang per buck ratio for the goods in $\mathbf{x}_{\mathcal{H}_i}$), thus maintaining the above condition. For any agent $k \in \mathcal{H}_i$, we have $\text{MBB}_k \subseteq \mathbf{x}_{\mathcal{H}_i}$ by construction of the hierarchy. Raising the prices of the goods in $\mathbf{x}_{\mathcal{H}_i}$ therefore corresponds to lowering the MBB ratios for the agents in $\mathcal{H}_i$. By choice of α_1, the price-rise stops as soon as a new MBB-edge appears between an agent $k \in \mathcal{H}_i$ and a good $j \notin \mathbf{x}_{\mathcal{H}_i}$. This ensures that the *new* maximum bang per buck ratio for any agent $k \in \mathcal{H}_i$ does not fall below its second highest bang per buck ratio prior to the price-rise, thus guaranteeing $\mathbf{x}_k \subseteq \text{MBB}_k$.

We can now define a Fisher market where each agent's endowment equals its spending under $\mathbf{x}$. Since $(\mathbf{x}, \mathbf{p})$ is an equilibrium for this market, we have that $\mathbf{x}$ is fPO (Proposition 1).

Next, we will argue that $\mathbf{x}$ is 3ε-EF1. Notice that ALG terminates only if either the current outcome $(\mathbf{x}, \mathbf{p})$ is 3ε-pEF1, or when $\alpha = \alpha_2$ (Line 23). In the first case, we get that $\mathbf{x}$ is 3ε-EF1 for the underlying fair division instance (Lemma 4.1). Therefore, we only need to analyze the second case.

ALGORITHM 1: ALG

Input: An instance $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$ such that valuations are power-of–$(1 + \varepsilon)$.
Output: An integral allocation $\mathbf{x}$ and a price vector $\mathbf{p}$.

`// ----------------------------Phase 1: Initialization----------------------------`

1 $\mathbf{x} \leftarrow$ Welfare-maximizing allocation (allocate each good j to the agent $i \in \arg\max_{k \in [n]} v_{k,j}$)
2 $\mathbf{p} \leftarrow$ For each good $j \in [m]$, set $p_j = v_{i,j}$ if $j \in \mathbf{x}_i$.
3 **if** $(\mathbf{x}, \mathbf{p})$ *is* 3ε-pEF1 **then return** $(\mathbf{x}, \mathbf{p})$

`// -----------------Phase 2: Removing price-envy within hierarchy-----------------`

4 $i \leftarrow$ least spender under $(\mathbf{x}, \mathbf{p})$ `/* break ties lexicographically */`
5 $\mathcal{H}_i \leftarrow$ BUILDHIERARCHY$(i, \mathbf{x}, \mathbf{p})$
6 $\ell \leftarrow 1$
7 **while** $\mathcal{H}_i^\ell$ *is non-empty and* $(\mathbf{x}, \mathbf{p})$ *is not* 3ε-pEF1 **do**
8   **if** $h \in \mathcal{H}_i^\ell$ *is an* ε*-path-violator along the alternating path* $P = \{i, j_1, h_1, \ldots, j_{\ell-1}, h_{\ell-1}, j, h\}$ **then**
9         $\mathbf{x}_h \leftarrow \mathbf{x}_h \setminus \{j\}$ and $\mathbf{x}_{h_{\ell-1}} \leftarrow \mathbf{x}_{h_{\ell-1}} \cup \{j\}$ `/* Swap operation */`
10   Repeat Phase 2 starting from Line 4
11   **else**
12         $\ell \leftarrow \ell + 1$
13 **if** $(\mathbf{x}, \mathbf{p})$ *is* 3ε-pEF1 **then**
14   **return** $(\mathbf{x}, \mathbf{p})$
15 **else**
16   Move to Phase 3 starting from Line 17

`// ------------------------------Phase 3: Price-rise------------------------------`

17 $\alpha_1 \leftarrow \min_{h \in \mathcal{H}_i,\, j \in [m] \setminus \mathbf{x}_{\mathcal{H}_i}} \frac{\alpha_h}{v_{h,j}/p_j}$, where α_h is the maximum bang per buck ratio for agent h, and $\mathbf{x}_{\mathcal{H}_i}$ is the set of goods currently owned by members of the hierarchy $\mathcal{H}_i$

`/*` α_1 `corresponds to raising prices until a new agent gets added to the hierarchy */`

18 $\alpha_2 \leftarrow \frac{1}{\mathbf{p}(\mathbf{x}_i)} \max_{k \in [n] \setminus \mathcal{H}_i} \min_{j \in \mathbf{x}_k} \mathbf{p}(\mathbf{x}_k \setminus \{j\})$

`/*` α_2 `corresponds to raising prices until the pEF1 condition is satisfied */`

19 $\alpha_3 \leftarrow (1 + \varepsilon)^s$, where s is the smallest integral power of $(1 + \varepsilon)$ such that $(1 + \varepsilon)^s > \frac{\mathbf{p}(\mathbf{x}_h)}{\mathbf{p}(\mathbf{x}_i)}$; here i is the least spender and $h \in \arg\min_{k \in [n] \setminus \mathcal{H}_i} \mathbf{p}(\mathbf{x}_k)$.

`/*` α_3 `corresponds to raising prices in multiples of` $(1+\varepsilon)$ `until the identity of the least spender changes */`

20 $\alpha \leftarrow \min(\alpha_1, \alpha_2, \alpha_3)$
21 **foreach** *good* $j \in \mathbf{x}_{\mathcal{H}_i}$ **do**
22     $p_j \leftarrow \alpha \cdot p_j$
23 **if** $\alpha = \alpha_2$ **then**
24   **return** $(\mathbf{x}, \mathbf{p})$
25 **else**
26   Repeat Phase 2 starting from Line 4

Let us suppose that the termination happens at time step t, and let $\mathbf{q}$ be the price vector maintained by Alg just *before* the price-rise step that lead to termination. After the time step t, Alg terminates with the allocation $\mathbf{x}$ and price vector $\mathbf{p}$. Since Phase 3 does not change the ownership of the goods, the allocation maintained by Alg just before termination is also $\mathbf{x}$.

Let i be the least spender at time step t, and let $\mathcal{H}_i$ be the hierarchy of agent i. Since Phase 3 only affects the prices of the goods in $\mathbf{x}_{\mathcal{H}_i}$, we have that $\mathbf{p}(\mathbf{x}_k) = \mathbf{q}(\mathbf{x}_k)$ for all $k \in [n] \setminus \mathcal{H}_i$, and $\mathbf{p}(\mathbf{x}_k) = \alpha_2 \mathbf{q}(\mathbf{x}_k)$ for all $k \in \mathcal{H}_i$. Additionally, at the end of (any execution of) Phase 2, no agent in the least spender's hierarchy is an ε-path-violator (and hence is also not an ε-violator). Thus,

$$\begin{aligned} (1+\varepsilon)\mathbf{q}(\mathbf{x}_i) &\geq \max_{k \in \mathcal{H}_i} \min_{j \in \mathbf{x}_k} \mathbf{q}(\mathbf{x}_k \setminus \{j\}) \\ \implies (1+\varepsilon)\mathbf{p}(\mathbf{x}_i) &\geq \max_{k \in \mathcal{H}_i} \min_{j \in \mathbf{x}_k} \mathbf{p}(\mathbf{x}_k \setminus \{j\}). \end{aligned} \tag{1}$$

By definition of α_2, we have the following condition for the agents outside the hierarchy:

$$\mathbf{p}(\mathbf{x}_i) = \alpha_2 \mathbf{q}(\mathbf{x}_i) \geq \max_{k \in [n] \setminus \mathcal{H}_i} \min_{j \in \mathbf{x}_k} \mathbf{q}(\mathbf{x}_k \setminus \{j\}) = \max_{k \in [n] \setminus \mathcal{H}_i} \min_{j \in \mathbf{x}_k} \mathbf{p}(\mathbf{x}_k \setminus \{j\}). \tag{2}$$

Equations (1) and (2) together imply that

$$(1+\varepsilon)\mathbf{p}(\mathbf{x}_i) \geq \max_{k \in [n]} \min_{j \in \mathbf{x}_k} \mathbf{p}(\mathbf{x}_k \setminus \{j\}), \tag{3}$$

which means that the outcome $(\mathbf{x}, \mathbf{p})$ is ε-pEF1 for agent i. If agent i is a least spender under $(\mathbf{x}, \mathbf{p})$ (i.e., i continues to a least spender after the price rise), then $\mathbf{x}$ is ε-pEF1 with respect to $\mathbf{p}$, and the lemma follows. Otherwise, an agent $h \in \arg\min_{k \in [n] \setminus \mathcal{H}_i} \mathbf{q}(\mathbf{x}_k)$ must become the least spender after the final price-rise step. In this case, we have that

$$\begin{aligned} & (1+\varepsilon)\mathbf{q}(\mathbf{x}_h) \geq \alpha_3 \mathbf{q}(\mathbf{x}_i) && \text{(by definition of } \alpha_3\text{)} \\ \implies & (1+\varepsilon)\mathbf{q}(\mathbf{x}_h) \geq \alpha_2 \mathbf{q}(\mathbf{x}_i) && (\alpha = \alpha_2 \implies \alpha_2 \leq \alpha_3) \\ \implies & (1+\varepsilon)\mathbf{p}(\mathbf{x}_h) \geq \mathbf{p}(\mathbf{x}_i) && \text{(since } \mathbf{p}(\mathbf{x}_i) = \alpha_2 \mathbf{q}(\mathbf{x}_i) \text{ and } \mathbf{p}(\mathbf{x}_h) = \mathbf{q}(\mathbf{x}_h)) \\ \implies & (1+\varepsilon)^2 \mathbf{p}(\mathbf{x}_h) \geq \min_{j \in \mathbf{x}_k} \mathbf{p}(\mathbf{x}_k \setminus \{j\}) && \text{for all } k \in [n], \end{aligned}$$

where the last inequality follows from Equation (3). Since $0 < \varepsilon < 1$, we have that $(1+\varepsilon)^2 < 1 + 3\varepsilon$. Thus, the new least spender (agent h) satisfies $(1+3\varepsilon)\mathbf{p}(\mathbf{x}_h) \geq \min_{j \in \mathbf{x}_k} \mathbf{p}(\mathbf{x}_k \setminus \{j\})$ for all $k \in [n]$. This implies that $(\mathbf{x}, \mathbf{p})$ is 3ε-pEF1. The stated claim now follows from Lemma 4.1. □

Lemma 5.2 (**Running time bound for power-of-$(1+\varepsilon)$ instance**). *Given any power-of-$(1+\varepsilon)$ instance as input,* Alg *terminates in time $O\left(\text{poly}(m, n, \frac{1}{\varepsilon}, \ln v_{\max})\right)$, where $v_{\max} = \max_{i,j} v_{i,j}$.*

The proof of Lemma 5.2 appears in the full version of the paper.

5.2 Analysis of Alg for General Valuations: Proof of Theorem 3.1

In this section, we will show that for any given fair division instance $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$ with integral valuations, an allocation that is envy-free up to one good (EF1) and Pareto efficient (PO) can be found in pseudopolynomial time (Theorem 3.1). In particular, for bounded valuations, this result provides a polynomial time algorithm for computing an EF1 and PO allocation (Remark 1).

We will prove Theorem 3.1 by running Alg on a *ε-rounded version* $\mathcal{I}' = \langle [n], [m], \mathcal{V}' \rangle$ of the given instance $\mathcal{I}$ for some parameter $\varepsilon > 0$. The instance $\mathcal{I}'$ is a power-of-$(1+\varepsilon)$ instance[10] constructed by *rounding up* the valuations in $\mathcal{I}$ to the nearest integer power of $(1+\varepsilon)$. From Lemma 5.1, we know that the allocation returned by Alg is 3ε-EF1 and fPO with respect to $\mathcal{I}'$ (for any given

[10]Recall that in a *power-of*-$(1+\varepsilon)$ instance, we have $v_{i,j} \in \{0, (1+\varepsilon)^a\}$ for some $a \in \mathbb{N}$ (possibly depending on i and j).

$\varepsilon > 0$). We will show that for an appropriate choice of ε, the same allocation turns out to be EF1 and PO with respect to the original instance $\mathcal{I}$. In addition, the running time bound in Lemma 5.2 instantiated for this choice of ε will show that Alg runs in pseudopolynomial time.

More formally, the *ε-rounded version* $\mathcal{I}' = \langle [n], [m], \mathcal{V}' \rangle$ of the given instance $\mathcal{I}$ is constructed as follows: For each agent $i \in [n]$ and each good $j \in [m]$, the valuation $v'_{i,j}$ is given by

$$v'_{i,j} := \begin{cases} (1+\varepsilon)^{\lceil \log_{1+\varepsilon} v_{i,j} \rceil} & \text{if } v_{i,j} > 0, \\ 0 & \text{if } v_{i,j} = 0. \end{cases}$$

Notice that $v_{i,j} \leq v'_{i,j} \leq (1+\varepsilon)v_{i,j}$ for each agent i and each good j.

Lemma 5.3 below establishes that for an appropriate choice of ε, an allocation that is fPO for the ε-rounded instance $\mathcal{I}'$ is PO with respect to the original instance $\mathcal{I}$. The proof of Lemma 5.3 appears in the full version of the paper.

Lemma 5.3. *Let $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$ be a fair division instance, and let $\varepsilon \leq \frac{1}{6m^3 v_{\max}^4}$. Then, an allocation $\mathbf{x}$ that is fPO for $\mathcal{I}'$ (the ε-rounded version of $\mathcal{I}$) is PO for the original instance $\mathcal{I}$.*

Lemma 5.4 establishes that for a small enough δ, a δ-EF1 allocation is in fact EF1.

Lemma 5.4. *Let $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$ be a fair division instance, and let $0 < \delta \leq \frac{1}{2m v_{\max}}$. Then, an allocation $\mathbf{x}$ is δ-EF1 for $\mathcal{I}$ if and only if it is EF1 for $\mathcal{I}$.*

Proof. If $\mathbf{x}$ is δ-EF1, we have that for every pair of agents $i, k \in [n]$, there exists a good $j \in \mathbf{x}_k$ such that $(1+\delta)v_i(\mathbf{x}_i) \geq v_i(\mathbf{x}_k \setminus \{j\})$. The bound on δ implies that $v_i(\mathbf{x}_k \setminus \{j\}) - v_i(\mathbf{x}_i) \leq \frac{1}{2}$. Integrality of valuations gives $v_i(\mathbf{x}_k \setminus \{j\}) - v_i(\mathbf{x}_i) \leq 0$, as desired. □

Theorem 3.1. *Given any fair division instance $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$ with additive valuations, an allocation that is envy-free up to one good (EF1) and Pareto efficient (PO) can be found in $O(\mathrm{poly}(m, n, v_{\max}))$ time, where $v_{\max} = \max_{i,j} v_{i,j}$.*

Proof. Let $\mathcal{I}' = \langle [n], [m], \mathcal{V}' \rangle$ be the ε-rounded version of $\mathcal{I}$ with $\varepsilon = \frac{1}{14m^3 v_{\max}^4}$. From Lemmas 5.1 and 5.2, we know that an allocation $\mathbf{x}$ that is 3ε-EF1 and fPO for $\mathcal{I}'$ can be found in $O\left(\mathrm{poly}(m, n, \frac{1}{\varepsilon}, \ln v_{\max})\right)$ time. Under the stated choice of ε, Lemma 5.3 implies that $\mathbf{x}$ must be Pareto efficient (PO) for $\mathcal{I}$. Therefore, we only need to show that $\mathbf{x}$ is EF1 for the instance $\mathcal{I}$.

Since $\mathbf{x}$ is 3ε-EF1 for $\mathcal{I}'$, we have that for every pair of agents $i, k \in [n]$, there exists a good $j \in \mathbf{x}_k$ such that $(1+3\varepsilon)v'_i(\mathbf{x}_i) \geq v'_i(\mathbf{x}_k \setminus \{j\})$. Furthermore, since $\mathcal{I}'$ is a ε-rounded version of $\mathcal{I}$, we have that $v'_{i,j} \leq (1+\varepsilon)v_{i,j}$ for each good $j \in [m]$. Hence, $(1+\varepsilon)(1+3\varepsilon)v_i(\mathbf{x}_i) \geq v'_i(\mathbf{x}_k \setminus \{j\})$. Finally, since the valuations in $\mathcal{I}'$ are a rounded-up version of those in $\mathcal{I}$, we have that $v_{i,j} \leq v'_{i,j}$ for each good $j \in [m]$, and thus $(1+\varepsilon)(1+3\varepsilon)v_i(\mathbf{x}_i) \geq v_i(\mathbf{x}_k \setminus \{j\})$. For $\varepsilon \leq 1$, this expression simplifies to $(1+7\varepsilon)v_i(\mathbf{x}_i) \geq v_i(\mathbf{x}_k \setminus \{j\})$, which means that $\mathbf{x}$ is 7ε-EF1 for the instance $\mathcal{I}$. Instantiating Lemma 5.4 for $\delta = 7\varepsilon$ gives that $\mathbf{x}$ is EF1 for $\mathcal{I}$. □

Lemma 5.5. *Given the ε-rounded version $\mathcal{I}'$ (of the instance $\mathcal{I}$) as input, Alg finds a 7ε-EF1 and ε-PO allocation for $\mathcal{I}$ in $O(\mathrm{poly}(m, n, \frac{1}{\varepsilon}, \ln v_{\max}))$ time.*

Proof. Let $\mathbf{x}$ be the allocation returned by Alg. From Lemma 5.1, we know that $\mathbf{x}$ is 3ε-EF1 and fPO for the ε-rounded instance $\mathcal{I}'$. By an argument similar to the one in the proof of Theorem 3.1, this implies that $\mathbf{x}$ is 7ε-EF1 for the original instance $\mathcal{I}$. The running time guarantee follows from Lemma 5.2. Hence, we only need to show that $\mathbf{x}$ is ε-PO.

Suppose, for contradiction, that $\mathbf{x}$ is ε-Pareto dominated by an allocation $\mathbf{y}$. Thus, $v_k(\mathbf{y}_k) \geq (1+\varepsilon)v_k(\mathbf{x}_k)$ for every agent $k \in [n]$ and $v_i(\mathbf{y}_i) > (1+\varepsilon)v_i(\mathbf{x}_i)$ for some agent $i \in [n]$. By construction of the ε-rounded instance $\mathcal{I}'$, we know that $v_{k,j} \leq v'_{k,j} \leq (1+\varepsilon)v_{k,j}$ for each agent k and each good j. Using the inequality $v'_{k,j} \leq (1+\varepsilon)v_{k,j}$ in a good-by-good manner for the bundle $\mathbf{x}_k$, along with the additivity assumption of valuations in the instance $\mathcal{I}$, we get that $(1+\varepsilon)v_k(\mathbf{x}_k) \geq v'_k(\mathbf{x}_k)$. By a similar application of the inequality $v_{k,j} \leq v'_{k,j}$ for the bundle $\mathbf{y}_k$, we get $v'_k(\mathbf{y}_k) \geq v_k(\mathbf{y}_k)$. Combining these relations gives $v'_k(\mathbf{y}_k) \geq v'_k(\mathbf{x}_k)$ for every agent $k \in [n]$ and $v'_i(\mathbf{y}_i) > v'_i(\mathbf{x}_i)$ for some agent $i \in [n]$. However, this means that the allocation $\mathbf{y}$ Pareto dominates the allocation $\mathbf{x}$ in the instance $\mathcal{I}'$, which is a contradiction since $\mathbf{x}$ is fPO for $\mathcal{I}'$. □

6 NASH SOCIAL WELFARE APPROXIMATION: PROOF OF THEOREM 3.3

This section proves that Alg provides a 1.45 ($\approx e^{1/e}$)-approximation for the Nash social welfare maximization problem in polynomial time. We begin by showing (in Lemma 3.4) that if the valuations of all the agents are identical, then any ε-EF1 allocation provides a $e^{(1+\varepsilon)/e}$-approximation to Nash social welfare. We will then use this result, along with an appropriate choice of ε, to prove the desired approximation bound in Theorem 3.3.

Lemma 3.4. *Given a fair division instance with identical and additive valuations, any ε-EF1 allocation provides a $e^{(1+\varepsilon)/e}$-approximation to Nash social welfare.*

The proof of Lemma 3.4 makes use a structural result stated as Lemma 6.1. A relevant notion used in these results is that of *partially-fractional allocations*, defined as follows: Given a subset $B \subset [m]$ of the set of goods, a partially-fractional allocation (with respect to B) is one where the goods in B have to be integrally allocated and the remaining goods in $[m] \setminus B$ can be fractionally allocated. Formally, a partially-fractional allocation $\mathbf{y} \in [0,1]^{n\times m}$ is such that for each agent $i \in [n]$, we have $\mathbf{y}_{i,j} \in \{0,1\}$ for each $j \in B$, and $\mathbf{y}_{i,j} \in [0,1]$ for each $j \in [m] \setminus B$. We write $\mathcal{F}_B$ to denote the set of all partially-fractional allocations with respect to B.

Lemma 6.1. *Let $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$ be an instance with additive and identical valuation functions (denoted by v for each agent) such that $m \geq n$. Let $B \subset [m]$ be a subset of goods such that $|B| < n$. Then, there is a partially-fractional allocation $\omega = (\omega_1, \ldots, \omega_n) \in \mathcal{F}_B$ that maximizes Nash social welfare (among allocations in $\mathcal{F}_B$) such that*

1. *Each agent gets at most one good from B under ω. That is, for each agent $i \in [n]$, we have that $|\{j \in [m] : \omega_{i,j} > 0\} \cap B| \leq 1$.*
2. *Any agent with strictly-better-than-the-worst allocation under ω gets exactly one integral good (and no fractional good). That is, for any agent $i \in [n]$ such that $v(\omega_i) > \min_k v(\omega_k)$, we have $\omega_{i,j} = 1$ for some $j \in B$, and $\omega_{i,j'} = 0$ for all $j' \neq j$.*

Proof. (of Lemma 6.1) We will consider a partially-fractional allocation $\omega = (\omega_1, \omega_2, \ldots, \omega_n) \in \mathcal{F}_B$ that maximizes the Nash social welfare over $\mathcal{F}_B$, and show that it can be transformed (without decreasing NSW) into an allocation in $\mathcal{F}_B$ that satisfies the stated properties.

Observe that if an agent $i \in [n]$ receives a fractional good under ω (i.e., $\omega_{i,j'} > 0$ for some $j' \notin B$), then its value $v(\omega_i)$ is equal to $\mu := \min_k v(\omega_k)$. This is because if $v(\omega_i) > \mu$, then we can "redistribute" j' between agent i and the least valued agent $\arg\min_k v(\omega_k)$ to obtain another fractional allocation in $\mathcal{F}_B$ with Nash social welfare strictly greater than that of ω. This contradicts the optimality of ω.

Therefore, any agent with value strictly greater than μ can only receive integral goods. We will show that any such agent necessarily receives exactly one integral good, and hence establish property (2). Suppose, for contradiction, that some agent $i \in [n]$ receives distinct goods $a, b \in B$ (i.e., $\omega_{i,a} = \omega_{i,b} = 1$), and $v(\omega_i) > \mu$. Since $|B| < n$, there exists some agent $k \in [n]$ that does not receive any integral good from B. Since agent k only receives fractional goods under ω, it follows from the above argument that $v(\omega_k) = \mu$. We can now assign one of the integral goods (say a) to agent k, and redistribute the fractional goods in ω_k between the agents i and k to obtain a new partially-fractional allocation with strictly greater Nash social welfare than ω. This contradicts the optimality of ω. Therefore, any agent with value strictly greater than μ must receive exactly one integral good under ω.

Next we address property (1). Note that this property already holds for agents with value strictly greater than μ. Hence, we only need to consider agents whose value is exactly equal to μ. Let $i \in [n]$ be an agent that receives two distinct goods $a, b \in B$ such that $v(\omega_i) = \mu$. Since $|B| < n$, there must exist an agent $k \in [n]$ that only receives fractional goods under ω. We can now perform a "swap" by assigning one of the integral goods (say a) to agent k, and fractional goods of total value $v(a)$ in ω_k to agent i to obtain a new partially-fractional allocation in $\mathcal{F}_B$ with the same Nash social welfare as ω. Such a swap is always possible, since $v(a) \leq v(\omega_i) = \mu = v(\omega_k)$. By repeating this process at most n times, we can obtain an NSW maximizer satisfying property (1). □

We will now prove Lemma 3.4.

Proof. (of Lemma 3.4) Let $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$ denote the given instance with identical and additive valuations, and let v denote the valuation function of all the agents. Write $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$ to denote an ε-EF1 allocation of $\mathcal{I}$. Let ℓ denote the value of the least valued bundle in $\mathbf{x}$, i.e., $\ell := \min_{i \in [n]} v(\mathbf{x}_i)$. By reindexing, we have that $v(\mathbf{x}_1) \geq v(\mathbf{x}_2) \geq \ldots \geq v(\mathbf{x}_n) = \ell$.

We will use g_i to denote a largest valued good in the bundle $\mathbf{x}_i$ of agent i, i.e., $g_i \in \arg\max_{g \in \mathbf{x}_i} v(g)$. The fact that $\mathbf{x}$ is an ε-EF1 allocation implies that

$$v(\mathbf{x}_i \setminus \{g_i\}) \leq (1+\varepsilon)\ell \qquad \text{for all } i \in [n]. \tag{4}$$

Define $B := \{g_1, g_2, \ldots, g_{n-1}\}$. Let $\omega = (\omega_1, \ldots, \omega_n) \in \mathcal{F}_B$ be a partially-fractional allocation (with respect to B) that maximizes Nash social welfare among all allocations in $\mathcal{F}_B$. Since $\mathcal{F}_B$ contains all the integral allocations, we have $\text{NSW}(\omega) \geq \text{NSW}(\mathbf{x}^*)$, where $\mathbf{x}^*$ is a Nash optimal (integral) allocation. Hence, to prove the lemma, it suffices to show that $\text{NSW}(\mathbf{x}) \geq \frac{1}{e^{(1+\varepsilon)/e}} \text{NSW}(\omega)$.

Define $\alpha := \min_{k \in [n]} v(\omega_k)/\ell$, and let $H := \{k \in [n] : v(\mathbf{x}_k) > \alpha\ell\}$. We will now consider partially-fractional allocations wherein only (and all) the goods in $\mathbf{x}_H$ have to be allocated integrally, and the remaining goods can be fractionally allocated. Write $\mathcal{F}_{\mathbf{x}_H}$ to denote the set of all such partially-fractional allocations.

The rest of the proof consists of four parts: First, we will construct an allocation $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}_H}$ such that $\text{NSW}(\mathbf{x}') \leq \text{NSW}(\mathbf{x})$. (Doing this will allow us to work with the ratio $\frac{\text{NSW}(\mathbf{x}')}{\text{NSW}(\omega)}$, which is convenient to bound from below.) Second, we will derive a lower bound on $\text{NSW}(\mathbf{x}')$ in terms of the relevant parameters α, ℓ, and n (and two other parameters h and t that we will define shortly). Third, we will derive an upper bound on $\text{NSW}(\omega)$ in terms of the same parameters. Finally, we will derive relationships between these parameters in order to achieve the stated approximation ratio.

- *Constructing the allocation* $\mathbf{x}'$: We start by initializing $\mathbf{x}' \leftarrow \mathbf{x}$. While there exist two agents $i, k \in [n]$ such that $\ell < v(\mathbf{x}'_i) < v(\mathbf{x}'_k) < \alpha\ell$, we transfer goods of value $\Delta := \min\{v(\mathbf{x}'_i) - \ell, \alpha\ell - v(\mathbf{x}'_k)\}$ from $\mathbf{x}'_i$ (the lesser valued bundle) to $\mathbf{x}'_k$ (the larger valued bundle). In particular, this transfer of goods ensures that the Nash social welfare does not increase. Also, this process must terminate because after every iteration of the while-loop, either $v(\mathbf{x}'_i) = \ell$ or $v(\mathbf{x}'_k) = \alpha\ell$, and therefore at least one of these agents does not participate in future iterations of the while-loop. Moreover, we have $\mathbf{x}'_k = \mathbf{x}_k$ for all $k \in H$, as the agents in H do not participate in the transfer. This proves that $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}_H}$.

- *Lower bound for* $\text{NSW}(\mathbf{x}')$: Notice that there can be at most one agent s in the allocation $\mathbf{x}'$ such that $v(\mathbf{x}'_s) \in (\ell, \alpha\ell)$. This is because the while-loop continues to execute if there are two or more such agents. For every other agent $k \in [n] \setminus H$, $v(\mathbf{x}'_k)$ is either ℓ or $\alpha\ell$.
  Let $h := |\{k \in [n] : v(\mathbf{x}_k) > \alpha\ell\}|$ denote the cardinality of the set H (i.e., $h = |H|$). Let $t := |\{k \in [n] : v(\mathbf{x}'_k) \geq \alpha\ell\}|$ denote the number of agents with a valuation at least $\alpha\ell$ in the allocation $\mathbf{x}'$. Thus, there are $(n-t)$ agents with valuation strictly below $\alpha\ell$ in $\mathbf{x}'$. We lower bound the valuations of these agents by ℓ in order to obtain the following relation:
$$\text{NSW}(\mathbf{x}') \geq \left(\prod_{i=1}^{h} v(\mathbf{x}_i) \times (\alpha\ell)^{(t-h)} \times \ell^{(n-t)} \right)^{1/n}. \tag{5}$$
- *Upper bound for* $\text{NSW}(\omega)$: Recall that $\omega = (\omega_1, \ldots, \omega_n) \in \mathcal{F}_B$ is a partially-fractional allocation (with respect to B) that maximizes Nash social welfare, and $|B| < n$. Using Lemma 6.1, we can assume, without loss of generality, that ω has the following two properties: (1) each agent $i \in [n-1]$ gets the good g_i under ω (this can be ensured via reindexing since the valuations are identical), and (2) if $v(\omega_i) > \min_{k \in [n]} v(\omega_k)$ for any $i \in [n]$, then agent i gets exactly one integral good under ω (and no fractional good).
  We will now argue that $v(\omega_k) \leq v(\mathbf{x}_k)$ for all $k \in H$. Suppose, for contradiction, that there exists an agent $k \in H$ such that $v(\omega_k) > v(\mathbf{x}_k)$. By definition of H, $v(\mathbf{x}_k) > \alpha\ell$ for all $k \in H$, and therefore $v(\omega_k) > \alpha\ell$. We also know that $\min_{k \in [n]} v(\omega_k) = \alpha\ell$, and therefore, by property (2), agent k must get exactly one integral good g_k under ω (and no fractional good). However, since $g_k \subseteq \mathbf{x}_k$, this contradicts the condition $v(\omega_k) > v(\mathbf{x}_k)$.
  By a similar reasoning, we can show that $v(\omega_k) = \alpha\ell$ for all $k \in [n] \setminus H$. Indeed, if $v(\omega_k) > \alpha\ell = \min_a v(\omega_a)$ for some $k \in [n] \setminus H$, then by property (2), agent k must get exactly one integral good g_k under ω (and no fractional good). This would imply that $v(\mathbf{x}_k) \geq v(g_k) = v(\omega_k) > \alpha\ell$, which contradicts the fact that $k \in [n] \setminus H$.
  These observations imply the following upper bound on the Nash social welfare of ω:
$$\text{NSW}(\omega) \leq \left(\prod_{i=1}^{h} v(\mathbf{x}_i) \times (\alpha\ell)^{(n-h)} \right)^{1/n}. \tag{6}$$
- *Deriving relationship between the parameters*: From Equation (4) and from the fact that $\mathbf{x}'_k = \mathbf{x}_k$ for all $k \in H$, we have $v(\mathbf{x}'_k \setminus \{g_k\}) = v(\mathbf{x}_k \setminus \{g_k\}) \leq (1+\varepsilon)\ell$ for all $k \in H$. Hence, we can upper bound the value of all goods in $[m]$ excluding the h goods in the set $\bigcup_{k \in H}\{g_k\}$, as follows:
$$\sum_{i \in H} v(\mathbf{x}'_i \setminus \{g_i\}) + \sum_{i \in [n] \setminus H} v(\mathbf{x}'_i) \leq (1+\varepsilon)h\ell + \alpha\ell(t-h+1) + \ell(n-t-1). \tag{7}$$
  Next, we will derive a lower bound for this total value by considering ω. Note that all the goods in the set $\bigcup_{k \in H}\{g_k\}$ are integrally allocated under ω. Hence, at least $(n-h)$ agents do not receive any good from the set $\bigcup_{k \in H}\{g_k\}$. The cumulative value derived by these agents under ω is at least $\alpha\ell(n-h)$, since $\min_k v(\omega_k) = \alpha\ell$. Using Equation (7) and the fact that $h \leq t$, we get
$$\alpha\ell(n-h) \leq (1+\varepsilon)t\ell + \alpha\ell(t-h+1) + \ell(n-t-1).$$
  Simplification gives that $t \geq \frac{(\alpha-1)(n-1)}{\alpha+\varepsilon}$, which can be further simplified to obtain
$$(n-t) \leq \frac{n(1+\varepsilon)+\alpha-1}{\alpha+\varepsilon} \leq \frac{n(1+\varepsilon)}{\alpha} + \frac{\alpha-1}{\alpha} \leq \frac{n(1+\varepsilon)}{\alpha} + 1.$$

Recall that $\mathrm{NSW}(\mathbf{x}) \geq \mathrm{NSW}(\mathbf{x}')$. Using the above relation with Equations (5) and (6) gives

$$\frac{\mathrm{NSW}(\mathbf{x})}{\mathrm{NSW}(\omega)} \geq \frac{\mathrm{NSW}(\mathbf{x}')}{\mathrm{NSW}(\omega)} \geq \frac{\left(\prod_{i=1}^{h} v(\mathbf{x}_i) \times (\alpha\ell)^{(t-h)} \times \ell^{n-t}\right)^{1/n}}{\left(\prod_{i=1}^{h} v(\mathbf{x}_i) \times (\alpha\ell)^{(n-h)}\right)^{1/n}} = \alpha^{-\frac{n-t}{n}} \geq \alpha^{-\frac{1+\varepsilon}{\alpha}-\frac{1}{n}}. \tag{8}$$

The $-1/n$ term in the exponent of α can be neglected via a scaling argument, as follows: Construct (for analysis only) a *scaled-up* instance $\mathcal{I}'$ consisting of $c \geq 1$ copies of the instance $\mathcal{I}$. For any allocation $\mathbf{y}$ that is ε-EF1 for $\mathcal{I}$, the allocation $\mathbf{y}' = (\mathbf{y}, \mathbf{y}, \ldots, \mathbf{y})$ is ε-EF1 for $\mathcal{I}'$. Write $n', \alpha', \ell', \omega'$ to denote the analogues of n, α, ℓ, ω in $\mathcal{I}'$. Also, let $\tilde{\omega}$ denote the fractional allocation $(\omega, \omega, \ldots, \omega)$ in $\mathcal{I}'$. It is easy to see that $n' = cn$, $\alpha' = \alpha$, and $\ell' = \ell$. Moreover,

$$\frac{\mathrm{NSW}(\mathbf{y})}{\mathrm{NSW}(\omega)} = \frac{\mathrm{NSW}(\mathbf{y}')}{\mathrm{NSW}(\tilde{\omega})} \geq \frac{\mathrm{NSW}(\mathbf{y}')}{\mathrm{NSW}(\omega')} \geq \alpha^{-\frac{1+\varepsilon}{\alpha}-\frac{1}{cn}},$$

where the first term is for the instance $\mathcal{I}$, and the remaining terms are for the instance $\mathcal{I}'$. In addition, the relation $\mathrm{NSW}(\omega') \geq \mathrm{NSW}(\tilde{\omega})$ follows from the optimality of ω' for $\mathcal{I}'$. By choosing a sufficiently large value of c, the term $-1/cn$ in the exponent can be made arbitrarily small. Therefore, the lower bound in Equation (8) is (arbitrarily close to) $\alpha^{-\frac{1+\varepsilon}{\alpha}}$. Finally, notice that the function $z^{-\frac{1+\varepsilon}{z}}$ with $z \geq 0$ is minimized at $z = e$. This gives a lower bound of $e^{-(1+\varepsilon)/e}$, as desired. □

We will now proceed to the proof of Theorem 3.3. The proof relies on transforming a general fair division instance into one with identical valuations, and showing that the Nash social welfare of the allocation returned by Alg is preserved in this transformation.

Theorem 3.3. *For additive valuations, there exists a polynomial-time* 1.45*-approximation algorithm for the Nash social welfare maximization problem.*

Proof. For a given instance $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$, let $\mathbf{z}$ and $\mathbf{q}$ denote the allocation and the price vector respectively that are returned by Alg when provided as input the ε-rounded version of $\mathcal{I}$ (the parameter ε is set to a small constant). Let $\alpha_i \neq 0$ denote the maximum bang per buck ratio of agent i with respect to $\mathbf{q}$. Construct a scaled instance $\mathcal{I}^{\mathrm{sc}} = \langle [n], [m], \mathcal{V}^{\mathrm{sc}} \rangle$ such that $v^{\mathrm{sc}}_{i,j} = \frac{1}{\alpha_i} v_{i,j}$ for all i and j.[11] Then, for any allocation $\mathbf{y}$, $\mathrm{NSW}(\mathbf{y})$ in $\mathcal{I}^{\mathrm{sc}}$ is $\frac{1}{(\prod_{i=1}^{n} \alpha_i)^{1/n}}$ times $\mathrm{NSW}(\mathbf{y})$ in the original instance $\mathcal{I}$. Therefore, in order to obtain the desired approximation guarantee, it suffices to show that $\mathbf{z}$ achieves an approximation factor of 1.45 in the scaled instance $\mathcal{I}^{\mathrm{sc}}$.

Let ω denote a Nash optimal (integral) allocation in the original instance $\mathcal{I}$. By the above argument, ω is Nash optimal in the scaled instance $\mathcal{I}^{\mathrm{sc}}$ as well. Additionally, for each agent i in $\mathcal{I}^{\mathrm{sc}}$, we have that $v^{\mathrm{sc}}_{i,j} = q_j$ for all $j \in \mathrm{MBB}_i$, and $v^{\mathrm{sc}}_{i,j} < q_j$ for all $j \notin \mathrm{MBB}_i$. Therefore, for any agent i, we have $v^{\mathrm{sc}}(\mathbf{z}_i) = \mathbf{q}(\mathbf{z}_i)$ (since, from Lemma 5.1, we have that $\mathbf{z}_i \subseteq \mathrm{MBB}_i$), and $v^{\mathrm{sc}}_i(\omega_i) \leq \mathbf{q}(\omega_i)$. Consequently, the Nash social welfare of the computed allocation $\mathbf{z}$ and the optimal allocation ω satisfy the following relations in the scaled instance $\mathcal{I}^{\mathrm{sc}}$:

$$\left(\prod_{i=1}^{n} v^{\mathrm{sc}}_i(\mathbf{z}_i)\right)^{1/n} = \left(\prod_{i=1}^{n} \mathbf{q}(\mathbf{z}_i)\right)^{1/n} \tag{9}$$

and

$$\left(\prod_{i=1}^{n} v^{\mathrm{sc}}_i(\omega_i)\right)^{1/n} \leq \left(\prod_{i=1}^{n} \mathbf{q}(\omega_i)\right)^{1/n}. \tag{10}$$

[11] A similar scaling was used by Cole and Gkatzelis [13] in their analysis of NSW approximation.

We will further transform the valuations in $\mathcal{I}^{\text{sc}}$ to obtain an instance $\mathcal{I}^{\text{id}} = \langle [n], [m], \mathcal{V}^{\text{id}} \rangle$ with *identical* valuations. Specifically, we set $v^{\text{id}}_{i,j} = q_j$ for all i and j. We know from Lemma 5.1 that $\mathbf{z}$ is 3ε-pEF1 with respect to $\mathbf{q}$ in the original instance $\mathcal{I}$, and that $\mathbf{z}_i \subseteq \text{MBB}_i$ for each agent $i \in [n]$. It then follows that z is 3ε-EF1 in the identical valuations instance $\mathcal{I}^{\text{id}}$. Furthermore, when $\varepsilon = \frac{1}{300}$, the allocation z is $\frac{1}{100}$–EF1 in $\mathcal{I}^{\text{id}}$, and therefore from Lemma 3.4, we have that

$$\begin{aligned}
\left(\prod_{i=1}^{n} \mathbf{q}(\mathbf{z}_i)\right)^{1/n} &\geq e^{-(1+0.01)/e} \max_{\mathbf{y}\in\mathcal{X}} \left(\prod_{i=1}^{n} \mathbf{q}(\mathbf{y}_i)\right)^{1/n} \\
&\geq \frac{1}{1.45} \max_{\mathbf{y}\in\mathcal{X}} \left(\prod_{i=1}^{n} \mathbf{q}(\mathbf{y}_i)\right)^{1/n} \\
&\geq \frac{1}{1.45} \left(\prod_{i=1}^{n} \mathbf{q}(\omega_i)\right)^{1/n} \\
&\geq \frac{1}{1.45} \left(\prod_{i=1}^{n} v_i^{\text{sc}}(\omega_i)\right)^{1/n} \qquad \text{(using Equation (10))}.
\end{aligned}$$

The previous inequality and Equation (9) together give us an approximation factor of 1.45 under the valuation profile $\mathcal{V}^{\text{sc}}$:

$$\left(\prod_{i=1}^{n} v_i^{\text{sc}}(\mathbf{z}_i)\right)^{1/n} \geq \frac{1}{1.45} \left(\prod_{i=1}^{n} v_i^{\text{sc}}(\omega_i)\right)^{1/n},$$

which provides a similar approximation guarantee for the original instance $\mathcal{I}$. Finally, observe that the allocation z can be computed in polynomial time for the above choice of ε (Lemma 5.5). This completes the proof of Theorem 3.3. □

7 CONCLUDING REMARKS

We studied the problem of finding a fair and efficient allocation of indivisible goods. Our work provided a framework based on integral Fisher markets and an (approximate) price envy-freeness condition resulting in a pseudopolynomial algorithm for finding an EF1 and PO allocation, and a polynomial time 1.45-approximation algorithm for Nash social welfare. Determining whether there exists a (strongly) polynomial time algorithm for the problem of finding an EF1 and PO allocation remains an interesting direction for future work. Extensions of our results to more general classes of valuations (e.g., submodular) will also be interesting.

ACKNOWLEDGMENTS

SB was supported in part by a Ramanujan Fellowship (SERB - SB/S2/RJN-128/2015). The second author gratefully acknowledges the support of a Google travel grant.

REFERENCES

[1] Nima Anari, Shayan Oveis Gharan, Tung Mai, and Vijay V Vazirani. 2018. Nash Social Welfare for Indivisible Items under Separable, Piecewise-Linear Concave Utilities. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2274–2290.

[2] Nima Anari, Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. 2017. Nash Social Welfare, Matrix Permanent, and Stable Polynomials. In *Proceedings of the 8th Conference on Innovations in Theoretical Computer Science (ITCS)*.

[3] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2017. Finding Fair and Efficient Allocations. *arXiv preprint arXiv:1707.04731* (2017). https://arxiv.org/abs/1707.04731

[4] Xiaohui Bei, Jugal Garg, Martin Hoefer, and Kurt Mehlhorn. 2017. Earning Limits in Fisher Markets with Spending-Constraint Utilities. In *Proceedings of the International Symposium on Algorithmic Game Theory (SAGT)*. 67–79.
[5] Sylvain Bouveret and Michel Lemaître. 2016. Characterizing Conflicts in Fair Division of Indivisible Goods Using a Scale of Criteria. *Autonomous Agents and Multi-Agent Systems* 30, 2 (2016), 259–290.
[6] William C Brainard and Herbert Scarf. 2000. *How to Compute Equilibrium Prices in 1891.* Technical Report. Cowles Foundation for Research in Economics, Yale University.
[7] Steven J Brams and Alan D Taylor. 1996. *Fair Division: From Cake-cutting to Dispute Resolution.* Cambridge University Press.
[8] Felix Brandt, Vincent Conitzer, Ulle Endriss, Ariel D Procaccia, and Jérôme Lang. 2016. *Handbook of Computational Social Choice.* Cambridge University Press.
[9] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
[10] Eric Budish, Gérard P Cachon, Judd B Kessler, and Abraham Othman. 2016. Course Match: A Large-scale Implementation of Approximate Competitive Equilibrium from Equal Incomes for Combinatorial Allocation. *Operations Research* 65, 2 (2016), 314–336.
[11] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*. 305–322.
[12] Richard Cole, Nikhil Devanur, Vasilis Gkatzelis, Kamal Jain, Tung Mai, Vijay V Vazirani, and Sadra Yazdanbod. 2017. Convex Program Duality, Fisher Markets, and Nash Social Welfare. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*. 459–460.
[13] Richard Cole and Vasilis Gkatzelis. 2015. Approximating the Nash Social Welfare with Indivisible Items. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC)*. 371–380.
[14] Bart de Keijzer, Sylvain Bouveret, Tomas Klos, and Yingqian Zhang. 2009. On the Complexity of Efficiency and Envy-freeness in Fair Division of Indivisible Goods with Additive Preferences. In *International Conference on Algorithmic Decision Theory (ADT)*. 98–110.
[15] Nikhil R Devanur, Christos H Papadimitriou, Amin Saberi, and Vijay V Vazirani. 2008. Market Equilibrium via a Primal–Dual Algorithm for a Convex Program. *Journal of the ACM (JACM)* 55, 5 (2008), 22.
[16] Edmund Eisenberg and David Gale. 1959. Consensus of Subjective Probabilities: The Pari-mutuel Method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168.
[17] Duncan Foley. 1967. Resource Allocation and the Public Sector. *Yale Economic Essays* (1967), 45–98.
[18] Jugal Garg, Martin Hoefer, and Kurt Mehlhorn. 2018. Approximating the Nash Social Welfare with Budget-Additive Valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2326–2340.
[19] Jonathan Goldman and Ariel D Procaccia. 2015. Spliddit: Unleashing Fair Division Algorithms. *ACM SIGecom Exchanges* 13, 2 (2015), 41–46.
[20] Mamoru Kaneko and Kenjiro Nakamura. 1979. The Nash Social Welfare Function. *Econometrica: Journal of the Econometric Society* (1979), 423–435.
[21] David Kurokawa, Ariel D Procaccia, and Junxing Wang. 2018. Fair Enough: Guaranteeing Approximate Maximin Shares. *Journal of the ACM (JACM)* 65, 2 (2018), 8.
[22] Euiwoong Lee. 2017. APX-hardness of Maximizing Nash Social Welfare with Indivisible Items. *Inform. Process. Lett.* 122 (2017), 17–20.
[23] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*. 125–131.
[24] Andreu Mas-Colell, Michael Dennis Whinston, and Jerry R Green. 1995. *Microeconomic Theory.* Vol. 1. Oxford university press New York.
[25] Jiri Matousek. 2008. *Using the Borsuk-Ulam theorem: Lectures on Topological Methods in Combinatorics and Geometry.* Springer Science & Business Media.
[26] Hervé Moulin. 2004. *Fair Division and Collective Welfare.* MIT Press.
[27] John F Nash Jr. 1950. The Bargaining Problem. *Econometrica: Journal of the Econometric Society* (1950), 155–162.
[28] Nhan-Tam Nguyen, Trung Thanh Nguyen, Magnus Roos, and Jörg Rothe. 2014. Computational Complexity and Approximability of Social Welfare Optimization in Multiagent Resource Allocation. *Autonomous Agents and Multi-agent Systems* 28, 2 (2014), 256–289.
[29] Abraham Othman, Tuomas Sandholm, and Eric Budish. 2010. Finding Approximate Competitive Equilibria: Efficient and Fair Course Allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 873–880.
[30] Hugo Steinhaus. 1948. The Problem of Fair Division. *Econometrica* 16 (1948), 101–104.
[31] Hal R Varian. 1974. Equity, Envy, and Efficiency. *Journal of Economic Theory* 9, 1 (1974), 63–91.