

On the Proximity of Markets with Integral Equilibria

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

Sanath Kumar Krishnamurthy
Stanford University
sanathsk@stanford.edu

Abstract

We study Fisher markets that admit equilibria wherein each good is integrally assigned to some agent. While strong existence and computational guarantees are known for equilibria of Fisher markets with additive valuations (Eisenberg and Gale 1959; Orlin 2010), such equilibria, in general, assign goods fractionally to agents. Hence, Fisher markets are not directly applicable in the context of indivisible goods. In this work we show that one can always bypass this hurdle and, up to a bounded change in agents' budgets, obtain markets that admit an integral equilibrium. We refer to such markets as pure markets and show that, for any given Fisher market (with additive valuations), one can efficiently compute a “near-by,” pure market with an accompanying integral equilibrium.

Our work on pure markets leads to novel algorithmic results for fair division of indivisible goods. Prior work in discrete fair division has shown that, under additive valuations, there always exist allocations that simultaneously achieve the seemingly incompatible properties of fairness and efficiency (Caragiannis et al. 2016); here fairness refers to *envy-freeness up to one good* (EF1) and efficiency corresponds to *Pareto efficiency*. However, polynomial-time algorithms are not known for finding such allocations. Considering relaxations of proportionality and EF1, respectively, as our notions of fairness, we show that fair and Pareto efficient allocations can be computed in strongly polynomial time.

1 Introduction

Fisher markets are fundamental models of resource allocation in mathematical economics (Brainard and Scarf 2000). Such markets consist of a set of divisible goods along with a set of buyers who have prespecified budgets and valuations (over all possible bundles of the goods). In this work we focus on the basic setup wherein the valuations of the buyers are additive. In an equilibrium of a Fisher market, goods are assigned prices, each buyer spends its entire budget selecting only those goods that provide maximum value per unit of money spent, and the market clears. The relevance of market equilibria (specifically from a resource-allocation perspective) is substantiated by the first welfare theorem which asserts that such equilibria are always *Pareto efficient* (Mas-Colell, Whinston, and Green 1995, Chapter 16).

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The convex program of Eisenberg and Gale provides a remarkable characterization (and, in conjunction, a proof of existence) of equilibria in Fisher markets with additive valuations: the primal and dual solutions of their convex program correspond to the equilibrium allocations and prices, respectively (Eisenberg and Gale 1959; Nisan et al. 2007). The seminal work of Arrow and Debreu (1954) further shows that equilibria exist under more general market models and convex settings; see, e.g., Mas-Colell et al. (1995). The notable aspect of the Eisenberg-Gale characterization is that—in contrast to the encompassing result of Arrow and Debreu—it provides an efficient method for finding equilibria under additive valuations. Several algorithmic results have been developed recently for computing Fisher market equilibria and, in fact, strongly polynomial-time algorithms are known for the additive case (Orlin 2010; Végh 2012).

Along with efficiency, market equilibria provide strong fairness guarantees. A well-known result of Varian (1974) shows that if in a market all the agents have equal budgets, then any market equilibrium—specifically called competitive equilibrium from equal incomes (CEEI)—leads to an *envy-free* allocation. Envy freeness is a standard solution concept and it deems an (fractional) allocation of the (divisible) goods to be fair if, under it, each agent prefers its own bundle over that of any other agent (Foley 1967).

However, Fisher markets do not yield a representative model in the context of indivisible goods. Such goods correspond to discrete resources (that cannot be fractionally assigned) and naturally occur in several allocation problems, e.g., course assignment (Othman et al. 2010) and inventory pricing (Rotemberg 2011). A market equilibrium, in general, requires a fractional assignment of goods to agents. Hence, one cannot simply consider a market with indivisible goods and expect an equilibrium outcome wherein the goods do not have to be fractionally assigned. In other words, the desirable market properties of efficiency, fairness, and computational tractability are somewhat confined to divisible goods.

Our work shows that one can bypass this hurdle and, up to a bounded change in budgets, always obtain markets that admit *integral equilibria*. Specifically, we will consider markets that admit an equilibrium wherein each good is integrally assigned to some agent. We will refer to such Fisher markets as *pure markets*. Of course, not all mar-

kets are pure.¹ Nevertheless, the present paper shows that for every Fisher market (with additive valuations) there exists a “nearby” market which admits an integral equilibrium. Specifically, we prove that for any given market \mathcal{M} one can construct—with a bounded change in the budgets—a pure market \mathcal{M}' . Here, both the markets have the same set of agents, goods, and valuations, and the absolute change in any agent’s budget is upper bounded by $\|\mathbf{p}\|_\infty$, where \mathbf{p} is the equilibrium price (vector) of \mathcal{M} (Theorem 4 and Theorem 9).

Note that pure markets enable us to treat indivisible goods as divisible ones and apply standard (Fisher market) results, such as the first welfare theorem. The fact that the resulting equilibrium is integral ensures that—independent of the analytic treatment—the final allocation does not require the discrete goods to be fractionally allocated, i.e., it conforms to a legitimate assignment of the given indivisible goods.

Pure Markets for Discrete Fair Division. Our work on pure markets leads to novel algorithmic results for discrete fair division. Specifically, we address fair division of indivisible goods among agents with additive valuations. Note that there are no monetary transfers in this setup, i.e., unlike the market setting, here we do not have budgets or prices.

Classical notions of fairness—e.g., envy-freeness and *proportionality*²—typically address allocation of divisible goods, and are not directly applicable in the discrete setting. For instance, while an envy-free and proportional allocation of divisible goods always exists (Stromquist 1980), such an existential result does not hold when the goods are indivisible.³

To address this issue, in recent years cogent analogues of envy-freeness and proportionality have been proposed for addressing the discrete version of the fair-division problem. A well-studied solution concept in this line of work is *envy-freeness up to one good* (Budish 2011): an (integral) allocation is said to be envy-free up to one good (EF1) iff each agent prefers its own bundle over the bundle of any other agent up to the removal of one good. Along the lines of EF1, a surrogate of proportionality—called *proportionality up to one good*—has also been considered in prior work (Conitzer, Freeman, and Shah 2017). In particular, an allocation is said to be proportional up to one good (PROP1) iff each agent receives its proportional share after the inclusion of one extra good in its bundle.⁴

The work of (Lipton et al. 2004) shows that as long as the valuations of the agents are monotone an EF1 allocation can be computed efficiently. This result is notably general, since it guarantees the existence of EF1 allocations under arbitrary,

combinatorial (monotone) valuations. Caragiannis et al. (2016) established another attractive feature of this solution concept: under additive valuations, there always exists an allocation which is both EF1 and Pareto optimal (PO). Though, polynomial-time algorithms are not known for finding such a fair and efficient allocation—the work of (Barman, Krishnamurthy, and Vaish 2018) provides a pseudopolynomial time algorithm for this problem.

Under additive valuations, an EF1 allocations is also PROP1. Hence, in the additive-valuations context, the result of (Lipton et al. 2004) is also applicable to PROP1. Similarly, via the existence result of (Caragiannis et al. 2016), we get that if the agents’ valuations are additive, then there exists an allocation that is both PROP1 and PO.

We will show that—in contrast to the known pseudopolynomial result for finding EF1 and PO allocations (Barman et al. 2018)—one can compute allocations that are PROP1 and PO in strongly polynomial time (Corollary 10). This result highlights the applicability of our work on pure markets.

We also consider another, natural relaxation of EF1, which we refer to as EF_1^1 : this solution concept requires that any agent i is not envious of any other agent k , up to the inclusion of one good in i ’s bundle and the removal of one good from k ’s bundle. We develop an efficient algorithm for computing allocations of indivisible goods that are simultaneously EF_1^1 and PO (Corollary 11).

It is relevant to note that the work of Barman et al. (2018) can also be considered as one that finds pure markets with limited change in budgets. However, in this sense, the result of Barman et al. (2018) is not stronger than the one obtained in the present paper. That result does provide a stronger fairness guarantee (EF1 and PO in pseudopolynomial time), but one can show that the algorithm of Barman et al. (2018) can lead to larger (than the ones obtained in the present paper) perturbations in the budgets.⁵ Overall, the pure-market existence result obtained in this work (and the accompanying budget-perturbation bound) is not weaker than the one obtained in (Barman, Krishnamurthy, and Vaish 2018). Also, in contrast to that work, the present algorithm runs in strongly polynomial time and is able to address unequal budgets.

Our Techniques: We establish the result for pure markets via a constructive proof. In particular, we develop an efficient algorithm that starts with an equilibrium of the given market and rounds its (fractional) allocation to obtain an integral one. In particular, our algorithm integrally assigns all the goods, which to begin were fractionally assigned. The algorithm does not alter the prices of the goods. We obtain a pure market at the end by setting the new budgets to explicitly satisfy the budget-exhaustion condition with respect to the computed allocation and the unchanged prices. While the algorithm is quite direct, the sequence in which it allocates the goods is fairly relevant. A careful curation ensures that the new budgets are close to the given ones. Notably, in our empirical study (Section 5), it takes less time to execute this rounding than to compute an equilibrium of the given

¹Consider a market of a single good and two agents with equal budgets.

²A division among n agents is said to be *proportionally fair* iff each agent gets a bundle of value at least $1/n$ times her value for the grand bundle of goods.

³If a single indivisible good has to be allocated between two agents, then, under any allocation, the losing agent will be envious and will not achieve proportionality.

⁴In a fair-division problem with n agents, the proportional share of an agent i is defined to be the $1/n$ times the value that i has for the entire set of goods.

⁵The full version of this paper provides an example in which the current algorithm outperforms (in terms of budget perturbations) the one developed in Barman et al. (2018)

Fisher market.

In Section 4 we show that the integral allocation we obtain (via rounding) satisfies notable fairness and efficiency guarantees. Given that fair-division methods are widely used in practice,⁶ efficient and easy-to-implement algorithms—such as the ones developed in this work—have a potential for direct impact.

Additional Related Work: An interesting work of Babaioff et al. (2017) considers markets wherein the indivisibility of goods is explicitly enforced. In particular, in their framework each agent selects its most preferred subset of goods, among all subsets that satisfy the budget constraint. Hence, fractional selection/allocation are ruled out in this setup. For such integral markets, existence of equilibria is not guaranteed. By contrast, we solely focus on pure/fractional markets, wherein equilibria necessarily exist. The key distinction here is that a pure market is a fractional market that happens to admit an integral equilibria. While a pure market is integral in the sense of Babaioff et al. (2017), the indivisibility of goods is not explicitly enforced in this framework.

Babaioff et al. (2017) characterize the existence of equilibria in integral markets with two agents, at most five goods, and generic budgets. On the other hand, this paper establishes that, in the space of Fisher markets, pure (and, hence, integral) markets are dense, up to bounded change in budgets.

2 Notation and Preliminaries

Fisher market is a tuple $\mathcal{M} := \langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$ wherein $[n] = \{1, 2, \dots, n\}$ denotes the set of agents, $[m] = \{1, 2, \dots, m\}$ denotes the set of goods, $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$ denotes the valuation profile, and $\mathbf{e} = (e_1, e_2, \dots, e_n)$ denotes the budget vector. The valuation profile \mathcal{V} specifies the cardinal preferences of each agent $i \in [n]$ over the set of goods $[m]$ via a valuation function $v_i : [0, 1]^m \mapsto \mathbb{R}_{\geq 0}$. For any agent $i \in [n]$, the parameter $e_i \in \mathbb{R}_+$ represents agent i 's budget/endowment.

A *bundle* of goods is a vector $\mathbf{s} = (s_1, s_2, \dots, s_m) \in [0, 1]^m$ in which s_j represents the allocated quantity of the good j . In particular, the value that an agent $i \in [n]$ has for a bundle $\mathbf{s} \in [0, 1]^m$ is denoted as $v_i(\mathbf{s})$. A bundle \mathbf{s} is said to be integral if under it each good is allocated integrally, i.e., for each $j \in [m]$ we have $s_j \in \{0, 1\}$. Note that an integral bundle \mathbf{s} corresponds to the subset of goods $\{j \in [m] \mid s_j = 1\}$. If \mathbf{s} is an integral bundle, we will overload notation and let \mathbf{s} also denote the corresponding subset of goods, i.e., $\mathbf{s} := \{j \in [m] \mid s_j = 1\}$.

Throughout, we will assume that agents have nonnegative and additive valuations, i.e., for each agent $i \in [n]$ and any bundle \mathbf{s} , we have $v_i(\mathbf{s}) := \sum_{j \in [m]} v_{i,j} s_j$, where $v_{i,j} \geq 0$ denotes the value agent i has for good j .

Allocation: An allocation $\mathbf{x} \in [0, 1]^{n \times m}$ refers to a collection of n bundles $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ where $\mathbf{x}_i = (x_{i,1}, x_{i,2}, \dots, x_{i,m}) \in [0, 1]^m$ is the bundle allocated to agent $i \in [n]$. Furthermore, in an allocation at most one unit of each good is allocated, i.e., for all $j \in [m]$, we have

$\sum_{i \in [n]} x_{i,j} \leq 1$. In other words, an allocation corresponds to a fractional allocation of the goods among the agents. We will say that an allocation \mathbf{x} is integral iff its constituent bundles are integral, $\mathbf{x} \in \{0, 1\}^{n \times m}$.

Market outcome and equilibrium: For a Fisher market $\mathcal{M} = \langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$, a *market outcome* is tuple (\mathbf{x}, \mathbf{p}) where $\mathbf{x} \in [0, 1]^{n \times m}$ corresponds to an allocation and the price vector $\mathbf{p} = (p_1, p_2, \dots, p_m)$ associates a price $p_g \in \mathbb{R}_{\geq 0}$ with each good $g \in [m]$.

Given a price vector \mathbf{p} , write MBB_i to denote the set of goods that provide agent i the maximum possible utility per unit of money spent, $\text{MBB}_i := \{g \in [m] \mid v_{i,g}/p_g \geq v_{i,j}/p_j \text{ for all } j \in [m]\}$. MBB_i is called the maximum bang-per-buck set of agent i (under the price vector \mathbf{p}) and, for ease of presentation, we will denote the maximum bang-per-buck ratio by MBB_i as well, i.e., $\text{MBB}_i := \max_{j \in [m]} v_{i,j}/p_j$.

An outcome (\mathbf{x}, \mathbf{p}) is said to an *equilibrium* of a Fisher market $\mathcal{M} = \langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$ iff it satisfies the following conditions:

- *Market clearing:* each good $g \in [m]$ is either priced at zero, $p_g = 0$, or it is completely allocated, $\sum_{i=1}^n x_{i,g} = 1$.
- *Budget exhaustion:* Agents spend their entire budget, i.e., for all $i \in [n]$, the following equality holds $\sum_{g \in [m]} x_{i,g} p_g = \mathbf{x}_i \cdot \mathbf{p} = e_i$.
- *Maximum bang-per-buck allocation:* Each agent $i \in [n]$ spends its budget only on optimal goods, i.e., if $x_{i,g} > 0$ for good $g \in [m]$, then $g \in \text{MBB}_i$.

We will explicitly use the term *integral equilibrium* to refer to a market equilibrium (\mathbf{x}, \mathbf{p}) in which the allocation \mathbf{x} is integral.

Recall that equilibria of markets (with additive valuations) correspond to optimal solutions of the Eisenberg-Gale convex program (Eisenberg and Gale 1959; Nisan et al. 2007). Furthermore, in the additive case, strongly polynomial-time algorithms exist for finding market equilibria (Orlin 2010; Végh 2012).

The first welfare theorem ensures that equilibrium allocations are *Pareto efficient*, i.e., satisfy a standard measure of economic efficiency. Specifically, for an instance $\langle [n], [m], \mathcal{V} \rangle$, an allocation $\mathbf{x} \in [0, 1]^{n \times m}$ is said to be Pareto dominated by another allocation $\mathbf{y} \in [0, 1]^{n \times m}$ if $v_i(\mathbf{y}_i) \geq v_i(\mathbf{x}_i)$, for each agent $i \in [n]$, and $v_k(\mathbf{y}_k) > v_k(\mathbf{x}_k)$ for some agent $k \in [n]$. That is, compared to allocation \mathbf{x} , every agent is better off under \mathbf{y} and at least one agent is strictly better off. An allocation is said to be *Pareto efficient* or *Pareto optimal* (PO) if it is not Pareto dominated by any other allocation.

Definition 1 (Fractionally Pareto Efficient Allocation). *An allocation is said to be fractionally Pareto efficient (fPO) iff it is not Pareto dominated by any fractional allocation $\mathbf{y} \in [0, 1]^{n \times m}$.*

Note that an integral allocation $\mathbf{x} \in \{0, 1\}^{n \times m}$ can be fPO.

Proposition 2 (First Welfare Theorem; Mas-Colell et al. (1995)). *If (\mathbf{x}, \mathbf{p}) is an equilibrium of a Fisher market with*

⁶See, e.g., Spliddit (Goldman and Procaccia 2015): <http://www.spliddit.org/>

additive valuations, then the equilibrium allocation \mathbf{x} is fractionally Pareto efficient (fPO).

Along with efficiency, market equilibria are known to fair. In particular, if in a market all the agents have equal endowments, then any market equilibrium—specifically called competitive equilibrium from equal incomes (CEEI)—leads to an *envy-free* allocation (Varian 1974). Envy freeness is a standard solution concept and it deems an allocation \mathbf{x} to be fair if, under it, each agent prefers its own bundle over that of any other agent: $v_i(\mathbf{x}_i) \geq v_i(\mathbf{x}_k)$ for all $i, k \in [n]$ (Foley 1967). Hence, using Proposition 2 and the result of Varian (Varian 1974), we get that CEEI are both fair and efficient.

However, as observed earlier, equilibrium allocations are not guaranteed to be integral. That is, with indivisible goods, one can not directly apply the market framework and hope to retain the desirable properties of efficiency, fairness, computational tractability, or even universal existence.

Our work shows that interestingly, up to a bounded change in the endowments, one can always bypass this hurdle and obtain integral equilibria. Towards this end, the following notion will be useful.

Definition 3 (Pure Market). *A Fisher market is said to be pure iff it admits an integral equilibrium.*

As mentioned previously, pure markets enable us to treat indivisible goods as divisible ones and apply standard (Fisher market) results, such as the first welfare theorem. The fact that the resulting equilibrium is integral ensures that—independent of the analytic treatment—the final allocation does not require the discrete goods to be fractionally allocated, i.e., it conforms to a legitimate assignment of the given indivisible goods.

Spending Graph: We will use the construct of a spending graph to state and analyze our algorithm. Given a market $\mathcal{M} = \langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$ along with an outcome (\mathbf{x}, \mathbf{p}) , the *spending graph* $G(\mathbf{x}, \mathbf{p})$ is a weighted bipartite graph whose (bipartition) vertex sets correspond to the set of agents $[n]$ and the set of goods $[m]$, respectively. In the spending graph, we have an edge (i, j) between agent i and good j if and only if $x_{i,j} > 0$. The weight of any edge (i, j) in $G(\mathbf{x}, \mathbf{p})$ is the amount that agent i is spending on good j , i.e., weight of edge (i, j) is $x_{i,j}p_j$.

Given a Fisher market \mathcal{M} and an equilibrium (\mathbf{x}, \mathbf{p}) , it is always possible to rearrange the spending so that the spending graph is a forest, i.e., we can, in strongly polynomial time, find an \mathbf{x}' such that $(\mathbf{x}', \mathbf{p})$ is an equilibrium of \mathcal{M} and $G(\mathbf{x}', \mathbf{p})$ is a forest. This fact has been used in computing market equilibrium for markets (Orlin 2010) and for approximating the *Nash social welfare* objective (Cole and Gkatzelis 2015).

3 On the Proximity of Pure Markets

The main result of this section shows that for every Fisher market there always exists a “nearby” market which is pure. Our proof of this result is constructive. In particular, we develop a strongly polynomial-time algorithm (ALG) that, for any given market $\mathcal{M} = \langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$ and its equilibrium (\mathbf{x}, \mathbf{p}) , finds a pure market $\mathcal{M}' = \langle [n], [m], \mathcal{V}, \mathbf{e}' \rangle$ such that

the absolute perturbation in endowments is at most $\|\mathbf{p}\|_\infty$, i.e., $\|\mathbf{e} - \mathbf{e}'\|_\infty \leq \|\mathbf{p}\|_\infty$. ALG, in particular, computes an integral equilibrium (\mathbf{x}', p) of \mathcal{M}' , thereby certifying that it is indeed pure.

Theorem 4 (Main Result). *Given a Fisher market $\mathcal{M} = \langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$ with additive valuations and its equilibrium (\mathbf{x}, \mathbf{p}) , we can find—in strongly polynomial time—a budget vector \mathbf{e}' and an integral allocation \mathbf{x}' such that*

- $(\mathbf{x}', \mathbf{p})$ is an integral equilibrium of the market $\mathcal{M}' = \langle [n], [m], \mathcal{V}, \mathbf{e}' \rangle$.
- The budget vector \mathbf{e}' is close to \mathbf{e} : $\|\mathbf{e}' - \mathbf{e}\|_\infty \leq \|\mathbf{p}\|_\infty$. In addition, $\sum_{i=1}^n e'_i = \sum_{i=1}^n e_i$.

Note that (in contrast to computing an arbitrary equilibrium) finding an integral equilibrium is computationally hard, i.e., determining whether a given Fisher market is pure is an NP-hard problem.⁷ Hence, a notable aspect of ALG is that it efficiently finds an integral equilibrium of the accompanying pure market.

3.1 Our algorithm

Recall that, for any given market \mathcal{M} and its equilibrium (\mathbf{x}, \mathbf{p}) , we can assume, without loss of generality, that the spending graph $G(\mathbf{x}, \mathbf{p})$ is a forest. Our algorithm, ALG, constructs a new (integral) allocation \mathbf{x}' by iteratively assigning goods to agents until all the goods are allocated. In ALG, we initialize G to be the spending forest $G(\mathbf{x}, \mathbf{p})$ and root each tree in G at some agent. Then, we assign child goods to agents $i \in [n]$ with no parents (i.e., to root agents), until adding any more child good to i would violate i ’s original endowment (i.e., budget constraint) e_i . The remaining child goods are then appropriately assigned to grandchildren agents. After each such distribution, we delete this parent agent i and all of its child goods (that have now been allotted). Overall, we repeat this specific method of distributing goods until G is empty.

The integral allocation \mathbf{x}' we construct is a rounding of the allocation \mathbf{x} . In particular, if a good is integrally allocated to agent i under \mathbf{x} , then it will continue to be assigned to i in \mathbf{x}' . Hence, the focus here is to analyze the assignment of goods which are fractionally allocated (i.e., are not integrally allocated) in \mathbf{x} . We will use the term *contested goods* to refer to goods that are fractionally allocated in \mathbf{x} . Note that all the goods considered in the nested while-loops of ALG are contested.

3.2 Proof for Theorem 4

The runtime analysis of ALG is direct:

Proposition 5. *ALG runs in strongly polynomial time.*

In Lemma 7 we will show that the output of ALG, i.e., $(\mathbf{x}', \mathbf{p})$, is an equilibrium of market $\mathcal{M}' = \langle [n], [m], \mathcal{V}, \mathbf{e}' \rangle$. Lemma 8 asserts that the computed endowments \mathbf{e}' are close to given budgets \mathbf{e} . Together, Lemma 7 and Lemma 8 directly imply Theorem 4.

The following supporting claim shows that ALG maintains a useful invariant.

⁷This hardness result can be established via a reduction from the Partition problem.

Algorithm: ALG

Input : A Fisher market $\mathcal{M} = \langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$ with additive valuations and an equilibrium (\mathbf{x}, \mathbf{p}) of \mathcal{M} .

Output : An integral allocation \mathbf{x}' and a budget vector \mathbf{e}' such that $(\mathbf{x}', \mathbf{p})$ is an integral equilibrium of the market $\mathcal{M}' = \langle [n], [m], \mathcal{V}, \mathbf{e}' \rangle$ and $\|\mathbf{e}' - \mathbf{e}\|_\infty \leq \|\mathbf{p}\|_\infty$

- 1 Set $\mathbf{x}' \leftarrow (\emptyset, \emptyset, \dots, \emptyset)$, i.e., for any agent i we initialize $\mathbf{x}'_i \leftarrow \emptyset$
/* We construct \mathbf{x}' by assigning goods to agents until all goods are allocated */
- 2 Initialize G to be the spending forest of (\mathbf{x}, \mathbf{p}) , i.e., $G \leftarrow G(\mathbf{x}, \mathbf{p})$
/* Whenever we allocate a good, we delete the corresponding vertex from G . */
- 3 Root each tree in the forest G at some agent
- 4 Allocate all leaf goods to parent agents
/* That is, for all $j \in [m]$ if $x_{i,j} = 1$ then $x'_i \leftarrow x'_i \cup \{j\}$ and delete j from G . */
- 5 **while** there is an agent i with no parent (i.e., i is a root node) in G **do**
- 6 **while** there is a good g in the neighborhood of i (i.e., edge (i, g) is in G) such that $\mathbf{p}(x'_i \cup \{g\}) \leq e_i$ **do**
- 7 Allocate g to agent i : update $\mathbf{x}'_i \leftarrow \mathbf{x}'_i \cup \{g\}$ and delete g from G .
- 8 **end**
- 9 Allocate every remaining child j of i to any (agent) child k of j and delete j from G . Here, i and k are agents and j is a good
/* That is, before agent i 's deletion, its grandchildren inherit the remaining child goods of i */
- 10 Delete agent i from G .
- 11 **end**
- 12 $\mathbf{e}' \leftarrow (\mathbf{p}(x'_1), \mathbf{p}(x'_2), \dots, \mathbf{p}(x'_n))$

Claim 6. *Throughout the execution of ALG, the graph G is a forest. In addition, the root and leaves of every tree in G correspond to agents (i.e., are agent nodes).*

Proof. The graph G is initialized to be the spending forest, and throughout ALG we only delete vertices from G , without ever adding an edge. Hence, G continues to be a forest.

To establish the property about leaf nodes in G , note that in Step 4 we assign all the leaves which correspond to goods. Therefore, before the while-loop begins, all leaf nodes correspond to agents. If, for contradiction, we assume that a node $j \in [m]$ —which corresponds to a good—becomes a leaf at some point of time, then this must have happened due to the deletion of j 's child node $i \in [n]$ (which corresponds to an agent). However, we delete an agent node i only if it has no parent in G (this is exactly the case in which i is considered in the outer while-loop). This contradicts that fact that ALG would have deleted i , implying that a node j (which corresponds to a good) never becomes a leaf in G .

Finally, note that at the beginning of ALG the root nodes correspond to agents: in Step 3 we explicitly root the trees of G at agent nodes. As before, if we assume, for contradiction, that a good node $j \in [m]$ becomes a root at some point of time, then this must have happened due to the deletion of j 's parent node $i \in [n]$ (which corresponds to an agent). However, we delete an agent node i only after all of i 's child nodes (which includes j) have been assigned (see Step 9).

Therefore, before i 's deletion we would have assigned j to a grandchild of i (who is guaranteed to exist, due to the fact that j is not a leaf node). That is, ALG would have deleted j (from G) before i , contradicting the assumption that j ends up being a root node. Hence, the stated claim follows for the root nodes as well. \square

Lemma 7. *For a given market $\mathcal{M} = \langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$ (with additive valuations) and equilibrium (\mathbf{x}, \mathbf{p}) , let \mathbf{x}' and \mathbf{e}' , respectively, be the allocation and the endowment vector computed by ALG. Then, $(\mathbf{x}', \mathbf{p})$ is an integral equilibrium of the market $\mathcal{M}' = \langle [n], [m], \mathcal{V}, \mathbf{e}' \rangle$.*

Proof. We will first show that ALG ends up allocating every good. For any good $j \in [m]$, consider the iteration in which its parent node $i \in [n]$ is being considered in the outer while-loop, i.e., the loop after which i gets deleted. Note that the parent node i is guaranteed to exist since j is never a root (Claim 6). Furthermore, the algorithm does not terminate till it deletes all the agent nodes from G , hence there necessarily exists a point of time when the agent node i is under consideration.

By construction, good j either gets assigned to i or to a grandchild $k \in [n]$ of node i ; Claim 6 ensures that k exists. Hence, we get that all goods are allocated/deleted from G over the course of the algorithm. Hence, the integral allocation \mathbf{x}' satisfies the market clearing condition.

By construction, the allocation \mathbf{x}' , returned by ALG, is a rounding of the allocation \mathbf{x} . In particular, for every agent $i \in [n]$, the set of goods that i spends on in \mathbf{x}' is a subset of the goods that i spends on in \mathbf{x} , i.e., $\mathbf{x}'_i \subseteq \{j \in [m] \mid x_{i,j} > 0\}$. Therefore, analogous to \mathbf{x} , in \mathbf{x}' agents spend only on maximum bang-per-buck goods, $\mathbf{x}'_i \subseteq \text{MBB}_i$; note that the prices of the goods remain unchanged. Moreover, the budget vector \mathbf{e}' is chosen to satisfy the budget exhaustion condition. Hence $(\mathbf{x}', \mathbf{p})$ is an integral equilibrium of the market \mathcal{M}' . \square

Lemma 8. *For any given market $\mathcal{M} = \langle [n], [m], \mathcal{V}, \mathbf{e} \rangle$ (with additive valuations) and equilibrium (\mathbf{x}, \mathbf{p}) , the budget vector \mathbf{e}' computed by ALG satisfies $\|\mathbf{e}' - \mathbf{e}\|_\infty \leq \|\mathbf{p}\|_\infty$ and $\sum_{i=1}^n e'_i = \sum_{i=1}^n e_i$.*

Proof. In the while-loops of ALG an agent can receive only contested goods: either the parent good and/or its child goods. Agents that have no children in G (at the beginning of the while loops) or are isolated satisfy the endowment bound directly; such an agent i has at most one contested good, its parent \hat{g} , and we have $e_i - p_{\hat{g}} \leq e'_i \leq e_i + p_{\hat{g}}$. Recall that $\mathbf{p}(x_i) = e_i$. Hence, to complete the proof we now need to obtain the endowment bounds for agents that have child nodes.

Note that the child nodes (goods) of an agent i are never deleted before i . The child goods are allocated/deleted only when agent i is selected in the outer while-loop. If an agent i has children, but it does not receive any of its child nodes, then it must be the case that i 's endowment is high enough to not accommodate any child, g . Specifically, we have $\mathbf{p}(x'_i) + p_g > e_i$, i.e., $e'_i \geq e_i - p_g$. Furthermore, in this case, the

only good that i may have received during the execution of the while-loops is its parent good, \hat{g} , hence $e'_i \leq e_i + p_{\hat{g}}$.

The remainder of the analysis addresses agents who have children and receive at least one of their child nodes (goods). For such agents, the condition of the inner while-loop ensures that we never over allocate child nodes, $e'_i = \mathbf{p}(\mathbf{x}'_i) \leq e_i$. We will establish a lower bound for e'_i s by considering different cases based on whether an agent $i \in [n]$ receives all of its child nodes or just some of them. Here, we write $\hat{g} \in [m]$ to denote the parent good of agent i in G .

- If an agent i receives all of its child nodes, then $e'_i = \mathbf{p}(\mathbf{x}'_i) \geq \mathbf{p}(\mathbf{x}_i) - p_{\hat{g}}$; here, the subtracted term, $p_{\hat{g}}$, accounts for the fact that i might not have received its parent good \hat{g} . Hence, in this case we have $e'_i \geq e_i - p_{\hat{g}}$.
- In case agent i does not receive a child good g , from the condition in the inner while-loop, we get $\mathbf{p}(\mathbf{x}'_i) + p_g > e_i$. Otherwise, child g would have been included in \mathbf{x}'_i . Therefore, $e'_i = \mathbf{p}(\mathbf{x}'_i) \geq e_i - p_g$ and we get a lower bound in this case as well.

Overall, the endowments satisfy $\|\mathbf{e}' - \mathbf{e}\|_\infty \leq \|\mathbf{p}\|_\infty$.

Note that ALG does not modify the prices of the goods. Since both the markets \mathcal{M} and \mathcal{M}' have the same equilibrium prices \mathbf{p} , the budget-exhaustion and market-clearing conditions of \mathcal{M} and \mathcal{M}' give us: $\sum_i e'_i = \sum_j p_j = \sum_i e_i$. \square

Remark 1. *The proof of Lemma 8 shows that if $e'_i < e_i$ then there exists a good $g \notin \mathbf{x}'_i$ that was fractionally allocated to i under \mathbf{x} (i.e., $x_{i,g} > 0$) such that $e_i \leq e'_i + p_g$. Note that for such a good g (via the maximum bang-per-buck condition in the definition of an equilibrium) we have $g \in \text{MBB}_i$.*

In addition, the analysis ensure that if $e'_i > e_i$, then there exists a good $\hat{g} \in \mathbf{x}'_i \subseteq \text{MBB}_i$ (specifically, the parent of i) such that $e'_i \leq e_i + p_{\hat{g}}$.

From Proposition 5, Lemma 7, and Lemma 8, we directly obtain Theorem 4.

3.3 An Extension of Theorem 4

The fact that Theorem 4 requires an equilibrium of the given market is not a computational hurdle. The work of (Orlin 2010) provides a strongly polynomial-time algorithm for computing an equilibrium (\mathbf{x}, \mathbf{p}) of a given Fisher market \mathcal{M} . Hence, Theorem 4, along with the result of (Orlin 2010), leads to the following algorithmic result.

Theorem 9. *Given m goods, n agents with additive valuations, $\mathcal{V} = \{v_1, \dots, v_n\}$, and a budget vector \mathbf{e} . In (strongly) polynomial time, we can find a budget vector \mathbf{e}' , an integral allocation \mathbf{x}' , and a price vector \mathbf{p} such that:*

- $(\mathbf{x}', \mathbf{p})$ is an integral equilibrium of the (pure) market $\mathcal{M}' = \langle [n], [m], \mathcal{V}, \mathbf{e}' \rangle$.
- The budget vector \mathbf{e}' is close to \mathbf{e} : $\|\mathbf{e}' - \mathbf{e}\|_\infty \leq \|\mathbf{p}\|_\infty$ and $\sum_{i=1}^n e'_i = \sum_{i=1}^n e_i$.

4 Pure Markets for Discrete Fair Division

The section addresses the problem of fairly dividing m indivisible goods among a set of n agents with nonnegative, additive valuations $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$. We will denote an instance of a fair division problem as a tuple $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$.⁸ Note that for each agent $i \in [n]$ the valuation for a subset of goods $S \subseteq [m]$ satisfies $v_i(S) = \sum_{j \in S} v_{i,j}$, where $v_{i,j} \in \mathbb{R}_+$ is the value that agent i has for good j .

A prominent solution concept in discrete fair division is *envy-freeness up to one good*. Formally, for a fair-division instance $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$, an integral allocation $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n) \in \{0, 1\}^{n \times m}$ is said to be envy-free up to one good (EF1) iff for every pair of agents $i, k \in [n]$ there exists a good $g \in \mathbf{x}_k$ such that $v_i(\mathbf{x}_i) \geq v_i(\mathbf{x}_k \setminus \{g\})$.

Strong existential guarantees are known for EF1, even under combinatorial valuations: the work of (Lipton et al. 2004) shows that as long as the valuations of the agents are monotone an EF1 allocation exists and can be computed efficiently. (Caragiannis et al. 2016) prove that, in the case of additive valuations, this notion of fairness is compatible with (Pareto) efficiency, i.e., there exists an allocation which is both EF1 and Pareto optimal (PO). However, polynomial-time algorithms are not known for finding such allocations—the work of (Barman et al. 2018) provides a pseudopolynomial time algorithm for this problem.

Along the lines of EF1, a surrogate of proportionality—called *proportionality up to one good*—has also been considered in prior work (Conitzer et al. 2017). Formally, an allocation $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ is said to be proportional up to one good (PROP1) iff for every agent $i \in [n]$ there exists a good $g \in [m]$ such that $v_i(\mathbf{x}_i \cup \{g\}) \geq v_i([m])/n$. Write Prop_i to denote the proportional share of agent i , i.e., $\text{Prop}_i := v_i([m])/n$.

Under additive valuations, EF1 allocations are also PROP1. Hence, the result of (Lipton et al. 2004) implies that PROP1 allocations exist when the valuations are additive. Similarly, via (Caragiannis et al. 2016), we get that if the agents' valuations are additive, then there exists an allocation that is both PROP1 and PO.

We will show that—in contrast to the known pseudopolynomial results for finding EF1 and fPO allocations (Barman et al. 2018)—one can compute allocations that are PROP1 and fPO in strongly polynomial time (Corollary 10).⁹ Finding a PROP1 and PO allocation in polynomial time was identified as an open question in (Conitzer et al. 2017), and our algorithmic result for this problem highlights the applicability of Theorem 9.

In addition, we prove a similar result for a natural relaxation of EF1, which we call *envy-free up to addition of a good in the first bundle and removal of another good from the other bundle* (EF₁⁺). Formally, an integral allocation $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ is said to be EF₁⁺ iff for every pair of agents $i, k \in [n]$, there exist goods $g_1 \in [m]$ and $g_2 \in \mathbf{x}_k$,

⁸We do not have budgets or prices in the fair division setup.

⁹Recall that fPO is a stronger solution concept than PO, since it requires that an allocation is not Pareto dominated by any fraction (and, hence, any integral) allocation. On the other hand, PO rules out domination solely by integral allocations.

such that $v_i(\mathbf{x}_i \cup \{g_1\}) \geq v_i(\mathbf{x}_i \setminus \{g_2\})$. Corollary 11 shows that an integral allocation, which is both EF_1^1 and fPO, can be computed efficiently.

Corollary 10. *Given a fair-division instance with indivisible goods and additive valuations, in strongly polynomial time we can compute an integral allocation \mathbf{a} which is both PROP1 (fair) and fPO (efficient).*

Proof. Given a fair-division instance $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$, we construct a Fisher market $\mathcal{M} = \langle [n], [m], \mathcal{V}, \mathbf{e} = \vec{1} \rangle$ by setting the endowment of each agent equal to one. Theorem 9 shows that in strongly polynomial time we can compute an equilibrium (\mathbf{x}, \mathbf{p}) of the market \mathcal{M} and, then, round \mathbf{x} to an integral allocation \mathbf{a} and obtain a budget vector \mathbf{e}' such that (\mathbf{a}, \mathbf{p}) is an integral equilibrium of the market $\mathcal{M}' = \langle [n], [m], \mathcal{V}, \mathbf{e}' \rangle$ and the budget vector \mathbf{e}' is close to $\mathbf{e} = \vec{1}$; in particular, $\|\mathbf{e}' - \vec{1}\|_\infty \leq \|\mathbf{p}\|_\infty$.

Since \mathbf{a} is an equilibrium of the Fisher market \mathcal{M}' , via the first welfare theorem (Proposition 2), we know that \mathbf{a} is fPO. Next we will prove that \mathbf{a} is PROP1 as well.

The conditions that define an equilibrium ensure that for all agents $i \in [n]$ and goods $g \in \mathbf{a}_i$ (i.e., the goods that are allocated to i in \mathbf{a}) we have $\frac{v_{i,g}}{p_g} = \text{MBB}_i := \max_{j' \in [m]} \frac{v_{i,j'}}{p_{j'}}$.¹⁰ The proof of Lemma 8 further provides the guarantee that if $e'_i < e_i$, then there exists a good $g \in \text{MBB}_i$ such that $e'_i \geq e_i - p_g$ (Remark 1). Using these facts we will perform a case analysis to show that allocation \mathbf{a} satisfies the stated fairness guarantee:

- If $\mathbf{p}(\mathbf{a}_i) = e'_i < e_i = 1$, then there exists a good $g \in \text{MBB}_i$ such that $\mathbf{p}(\mathbf{a}_i \cup \{g\}) \geq 1$. Therefore,

$$\begin{aligned} v_i(\mathbf{a}_i \cup \{g\}) &= \text{MBB}_i \mathbf{p}(\mathbf{a}_i \cup \{g\}) \\ &\quad (v_i \text{ is additive and } \mathbf{a}_i \subseteq \text{MBB}_i) \\ &\geq \text{MBB}_i \cdot 1 \quad (\mathbf{p}(\mathbf{a}_i \cup \{g\}) \geq 1) \\ &= \text{MBB}_i \cdot \mathbf{p}([m])/n \\ &\quad (\mathbf{p}([m]) = \sum_i e_i = n) \\ &\geq v_i([m])/n \\ &\quad (\text{MBB}_i p_j \geq v_{i,j} \text{ for all goods } j) \\ &= \text{Prop}_i \end{aligned}$$

- If $\mathbf{p}(\mathbf{a}_i) = e'_i \geq e_i = 1$, then

$$\begin{aligned} v_i(\mathbf{a}_i) &= \text{MBB}_i \mathbf{p}(\mathbf{a}_i) \quad (\mathbf{a}_i \subseteq \text{MBB}_i) \\ &\geq \text{MBB}_i \cdot 1 = \text{MBB}_i \mathbf{p}([m])/n \\ &\geq v_i([m])/n = \text{Prop}_i. \end{aligned}$$

Overall, we get that for any fair-division instance \mathcal{I} , a PROP1 and fPO allocation can be computed in strongly polynomial time. \square

Next, we provide a strongly polynomial-time algorithm for finding integral allocations that are simultaneously EF_1^1 and fPO.

¹⁰Note that the prices of the goods are the same under the two equilibria (\mathbf{x}, \mathbf{p}) and (\mathbf{a}, \mathbf{p}) .

Corollary 11. *Given a fair-division instance with indivisible goods and additive valuations, in strongly polynomial time we can compute an integral allocation \mathbf{a} which is both EF_1^1 and fPO.*

Proof. Given a fair-division instance $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$, we construct a Fisher market $\mathcal{M} = \langle [n], [m], \mathcal{V}, \mathbf{e} = \vec{1} \rangle$ by setting the endowment of each agent equal to one. Theorem 9 shows that in strongly polynomial time we can compute an equilibrium (\mathbf{x}, \mathbf{p}) of the market \mathcal{M} and, then, round \mathbf{x} to an integral allocation \mathbf{a} and obtain a budget vector \mathbf{e}' such that (\mathbf{a}, \mathbf{p}) is an integral equilibrium of the market $\mathcal{M}' = \langle [n], [m], \mathcal{V}, \mathbf{e}' \rangle$ and the budget vector \mathbf{e}' is close to $\mathbf{e} = \vec{1}$; in particular, $\|\mathbf{e}' - \vec{1}\|_\infty \leq \|\mathbf{p}\|_\infty$.

As noted in Remark 1, in this construction, for each agent $i \in [n]$ we have $|e'_i - e_i| \leq p_g$ where g is in fact a good that is fractionally allocated to i under \mathbf{x} , i.e., $x_{i,g} > 0$. Therefore, the following two properties hold

- P1: For each agent $i \in [n]$, there exists a good $g_1 \in \text{MBB}_i$ such that $\mathbf{p}(\mathbf{a}_i \cup \{g_1\}) \geq 1$.
If $\mathbf{p}(\mathbf{a}_i) = e'_i < 1$,¹¹ then this inequality follows from the first part of Remark 1. Otherwise, if $\mathbf{p}(\mathbf{a}_i) = e'_i \geq 1$, then the inequity holds trivially—the prices are nonnegative.
- P2: For each agent $k \in [n]$, there exists a good $g_2 \in \mathbf{a}_k \subseteq \text{MBB}_k$ such that $\mathbf{p}(\mathbf{a}_k \setminus \{g_2\}) \leq 1$.
If $\mathbf{p}(\mathbf{a}_k) = e'_k > 1$, then (as stated in the second part of Remark 1) we have a good $g_2 \in \mathbf{a}_k \subseteq \text{MBB}_k$ such that $\mathbf{p}(\mathbf{a}_k \setminus \{g_2\}) \leq 1$. For the complementary case, $\mathbf{p}(\mathbf{a}_k) = e'_k \leq 1$, this inequality directly holds.

Properties P1 and P2 imply that allocation \mathbf{a} is EF_1^1 (here, for any two agents i and k we select goods g_1 and g_2 as specified in the two properties, respectively):

$$\begin{aligned} v_i(\mathbf{a}_i \cup \{g_1\}) &= \text{MBB}_i \mathbf{p}(\mathbf{a}_i \cup \{g_1\}) \quad (g_1 \in \text{MBB}_i) \\ &\geq \text{MBB}_i \cdot 1 \quad (\text{P1}) \\ &\geq \text{MBB}_i \mathbf{p}(\mathbf{a}_k \setminus \{g_2\}) \quad (\text{P2}) \\ &\geq v_i(\mathbf{a}_k \setminus \{g_2\}) \\ &\quad (\text{MBB}_i p_j \geq v_{i,j} \text{ for all goods } j) \end{aligned} \quad \square$$

5 Some Empirical Results

For an experimental analysis of ALG, we generate random instances of Fisher markets with equal incomes ($\mathbf{e} = \vec{1}$) and uniformly random valuations for the agents. Details of our experimental setup and corresponding empirical results are deferred to a full version of the paper.

In summary, our implementation (complying with Corollaries 10 and 11) always finds an allocation which is PROP1 and EF_1^1 . In fact, for about 96% of the (randomly generated) instances, the implemented method finds an envy-free allocation. This suggests that, in practice, our algorithms outperform the stated theoretical guarantees. In addition, we find that it takes notably less time to execute our rounding method than to compute a market equilibrium (i.e., solve the Eisenberg-Gale program).

¹¹By construction, $e_i = 1$.

Acknowledgment. Siddharth Barman gratefully acknowledges the support of a Ramanujan Fellowship (SERB - SB/S2/RJN-128/2015) and a Pratiksha Trust Young Investigator Award.

References

- Arrow, K. J., and Debreu, G. 1954. Existence of an Equilibrium for a Competitive Economy. *Econometrica: Journal of the Econometric Society* 265–290.
- Babaioff, M.; Nisan, N.; and Talgam-Cohen, I. 2017. Competitive equilibria with indivisible goods and generic budgets. *arXiv preprint arXiv:1703.08150*.
- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 557–574. ACM.
- Brainard, W. C., and Scarf, H. 2000. How to Compute Equilibrium Prices in 1891. Technical report, Cowles Foundation for Research in Economics, Yale University.
- Budish, E. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119(6):1061–1103.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*, 305–322.
- Cole, R., and Gkatzelis, V. 2015. Approximating the nash social welfare with indivisible items. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, 371–380. ACM.
- Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 629–646. ACM.
- Eisenberg, E., and Gale, D. 1959. Consensus of Subjective Probabilities: The Pari-mutuel Method. *The Annals of Mathematical Statistics* 30(1):165–168.
- Foley, D. 1967. Resource Allocation and the Public Sector. *Yale Economic Essays* 45–98.
- Goldman, J., and Procaccia, A. D. 2015. Spliddit: Unleashing Fair Division Algorithms. *ACM SIGecom Exchanges* 13(2):41–46.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, 125–131.
- Mas-Colell, A.; Whinston, M. D.; and Green, J. R. 1995. *Microeconomic Theory*, volume 1. Oxford university press New York.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic Game Theory*. Cambridge University Press.
- Orlin, J. B. 2010. Improved algorithms for computing fisher’s market clearing prices: computing fisher’s market clearing prices. In *Proceedings of the forty-second ACM symposium on Theory of computing*, 291–300. ACM.
- Othman, A.; Sandholm, T.; and Budish, E. 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.
- Rotemberg, J. J. 2011. Fair pricing. *Journal of the European Economic Association* 9(5):952–981.
- Stromquist, W. 1980. How to Cut a Cake Fairly. *The American Mathematical Monthly* 87(8):640–644.
- Varian, H. R. 1974. Equity, Envy, and Efficiency. *Journal of Economic Theory* 9(1):63–91.
- Végh, L. A. 2012. Strongly polynomial algorithm for a class of minimum-cost flow problems with separable convex objectives. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, 27–40. ACM.