

Infinite-Dimensional Fisher Markets: Equilibrium, Duality and Optimization

Yuan Gao, Christian Kroer

Columbia University, Department of Industrial Engineering and Operations Research
gao.yuan@columbia.edu, christian.kroer@columbia.edu

Abstract

This paper considers a linear Fisher market with n buyers and a continuum of items. In order to compute market equilibria, we introduce (infinite-dimensional) convex programs over Banach spaces, thereby generalizing the Eisenberg-Gale convex program and its dual. Regarding the new convex programs, we establish existence of optimal solutions, KKT conditions, as well as strong duality. All these properties are established via non-standard arguments, which circumvent the limitations of duality theory in optimization over infinite-dimensional vector spaces. Furthermore, we show that there exists a pure equilibrium allocation, i.e., a division of the item space. Similar to the finite-dimensional case, a market equilibrium under the infinite-dimensional Fisher market is Pareto optimal, envy-free and proportional. We also show how to obtain the (a.e. unique) equilibrium prices and a pure equilibrium allocation from the (unique) equilibrium utility prices. When the item space is the unit interval $[0, 1]$ and buyers have piecewise linear utilities, we show that approximate equilibrium prices can be computed in polynomial time. This is achieved by solving a finite-dimensional convex program using the ellipsoid method. To this end, we give non-trivial and efficient subgradient and separation oracles. For general buyer valuations, we propose computing market equilibrium using stochastic dual averaging, which finds approximate equilibrium prices with high probability.

Introduction

¹ Market equilibrium (ME) is a classical concept from economics, where the goal is to find an allocation of a set of items to a set of buyers, as well as corresponding prices, such that the market clears. One of the simplest equilibrium models is the (finite-dimensional) linear *Fisher market*. A Fisher market consists of a set of n buyers and m divisible items, where the utility for a buyer is linear in their allocation. Each buyer i has a budget B_i and valuation v_{ij} for each item j . A ME consists of an allocation (of items to buyers) and prices (of items) such that (i) each buyer receives a bundle of items that maximizes their utility subject to their budget constraints, and (ii) the market clears (all items such that $p_j > 0$ are exactly allocated). In spite of

its simplicity, this model has several applications. Perhaps the most well-known application is in the *competitive equilibrium from equal incomes* (CEEI), where m items must be fairly divided among n agents. By giving each agent one unit of faux currency, the allocation from the resulting ME can be used as the fair division. This approach guarantees several fairness desiderata. Linear Fisher markets also have applications in large-scale ad markets (Conitzer et al. 2018, 2019) and fair recommender systems (Kroer et al. 2019; Kroer and Peysakhovich 2019).

For finite-dimensional linear Fisher markets, the Eisenberg-Gale convex program computes a market equilibrium via its optimal solution and Lagrange multipliers (Eisenberg and Gale 1959; Eisenberg 1961; Jain and Vazirani 2010; Nisan et al. 2007; Cole et al. 2017). However, in settings like Internet ad markets and fair recommender systems, the number of items is often huge (Kroer et al. 2019; Kroer and Peysakhovich 2019; Balseiro, Besbes, and Weintraub 2015), if not infinite or even uncountable. For example, each item can be characterized by a set of features, where features come from a compact set in a Euclidean space. This motivates our study on infinite-dimensional Fisher markets and ME for a *continuum* of items.

A problem closely related to our infinite-dimensional Fisher-market setting is the *cake-cutting* or *fair division* problem. There, the goal is to efficiently partition a “cake” – often modeled as a compact measurable space, or simply the unit interval $[0, 1]$ – among n agents so that certain fairness and efficiency properties are satisfied (Weller 1985; Brams and Taylor 1996; Cohler et al. 2011; Procaccia 2013; Cohler et al. 2011; Brams et al. 2012; Chen et al. 2013; Aziz and Ye 2014; Aziz and Mackenzie 2016; Legut 2017, 2020). See (Procaccia 2016) for a survey for the various problem setups, algorithms and complexity results. Weller (1985) shows the existence of a fair allocation, that is, a measurable division of a measurable space satisfying weak Pareto optimality and envy freeness. As will be seen shortly, when all buyers have the same budget, our definition of a *pure* ME, i.e., where the allocation consist of indicator functions of a.e.-disjoint measurable sets, is in fact equivalent to this notion of fair division (Weller 1985). A subtly different notion is considered in (Cohler et al. 2011; Chen et al. 2013): there, Pareto optimality is w.r.t. the envy-free divisions only. In addition, we also give an explicit characterization of the unique equi-

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹An extended version is available at <https://arxiv.org/abs/2010.03025>.

librium prices based on a pure equilibrium allocation under arbitrary budgets, generalizing the result of Weller (1985) which only hold for buyers with the same budgets.

Under piecewise constant valuations over the cake $[0, 1]$, the equivalence of fair division and market equilibrium in certain setups has been discovered and utilized in the design of cake-cutting algorithms (Brams et al. 2012; Aziz and Ye 2014). Here, we extend this connection to arbitrary valuations in the L_1 function space: we propose Eisenberg-Gale-type convex programs that characterize *all* ME (and hence all fair divisions). As a concrete example, we show that we can efficiently compute approximate equilibrium utility prices (of the buyers) and equilibrium prices (of the items) under piecewise linear valuations.

Summary of contributions. First, we consider an infinite-dimensional Fisher market with n buyers and a continuum of items Θ and generalize the concept of a market equilibrium to this setting. We then give two infinite-dimensional convex programs over Banach spaces of measurable functions on Θ ((\mathcal{P}_{EG}) and (\mathcal{D}_{EG})), generalizing the EG convex program and its dual under finite dimensions. For the new convex programs, we first establish existence of optimal solutions (Lemma 1 and 2). Due to the lack of a compatible constraint qualification, general duality theory does not apply to these convex programs. Instead, we establish various duality properties directly through nonstandard arguments (Lemma 3). Based on these duality properties and the existence of a minimizer in the “primal” convex program (\mathcal{P}_{EG}) , we show that an allocation and prices pair is a ME *if and only if* they are optimal solutions of our convex programs (Theorem 1). Furthermore, since (\mathcal{P}_{EG}) has a pure optimal solution (i.e., buyers get disjoint measurable item subsets), there exists a pure equilibrium allocation, i.e., a *division* (module zero-value item) of the item space. Finally, we show that a ME under the infinite-dimensional Fisher market satisfies (budget-weighted) proportionality, Pareto optimality and envy-freeness. Our results on the existence of ME and its fairness properties can be viewed as generalizations of those in (Weller 1985), in which every buyer (agent) has the same budget. When the item space is the unit interval $[0, 1]$ and buyers have piecewise linear utilities, we show that ϵ -approximate equilibrium prices can be computed in time polynomial in the market size and $\log \frac{1}{\epsilon}$ (Theorem 5). This is achieved by solving a reformulation of our convex programs where the infinitely many items only occur in the objective function, and showing that the ellipsoid method can be applied. To this end, we give nontrivial, polynomial-time first-order and separation oracles for the seemingly intractable objective function. Finally, for more general buyer valuations, we propose using the stochastic dual averaging algorithm (SDA) on the same reformulated convex program to compute approximate equilibrium prices and establish convergence guarantees (Theorem 6).

Infinite-Dimensional Fisher Markets

Measure-theoretic preliminaries. First, we introduce the measure-theoretic concepts that we will need. The following paragraph can be skimmed and referred back to later. The items will be represented by Θ , a compact subset of \mathbb{R}^d .

Denote the Lebesgue measure on \mathbb{R}^d as μ . Let \mathcal{M} be the set of real-valued (Borel) measurable functions on Θ . Since Θ is compact, it is (Borel) measurable and $\mu(\Theta) < \infty$. Functions that are equal a.e. on Θ form an equivalence class, which are treated as the same function. In fact, any function f in this equivalence class give the same linear functional $g \mapsto \int fg d\mu$. The suffix a.e. will be omitted unless the emphasis is necessary. For any $S \subseteq \mathcal{M}$, denote $S_+ = \{f \in S : f \geq 0\}$. For $f \in L^1(\Theta)$ and $g \in L^\infty(\Theta)$, denote $\langle f, g \rangle = \int_\Theta fg d\mu$. This notation aligns with the usual notation for bilinear form, i.e., applying a linear functional to a function. Since $L^\infty(\Theta)$ is the dual space of $L^1(\Theta)$, the integration $\int_\Theta fg d\mu$ is well-defined and is finite. Let $\mathbf{1}$ be the constant function taking value 1 on Θ . For any measurable set $A \subseteq \Theta$, $\mathbf{1}_A$ denotes the $\{0, 1\}$ -indicator function of A . For $q \in [1, \infty]$, let $L_q(\Theta)$ be the Banach space of L_q (integrable) functions on Θ with the usual L_q norm, i.e, for $f \in L^q(\Theta)$,

$$\|f\| = \begin{cases} \int_\Theta |f|^q d\mu & \text{if } q < \infty, \\ \inf\{M > 0 : |f| < M \text{ a.e.}\} & \text{if } q = \infty. \end{cases}$$

Any $\tau \in L^1(\Theta)_+$ can also be viewed as a measure on Θ via $\mu_\tau(A) := \int_A \tau d\mu$ for any measurable $A \subseteq \Theta$. Here, τ is in fact the Radon–Nikodym derivative of μ_τ w.r.t. μ . We will denote $\mu_\tau(A)$ simply as $\tau(A)$ for a measurable set $A \subseteq \Theta$ whenever there is no confusion. In this work, unless otherwise stated, any measure m used or constructed is absolutely continuous w.r.t. the Lebesgue measure μ and hence *atomless*. In other words, for any measurable set $A \subseteq \Theta$ such that $m(A) > 0$ and any $0 < c < m(A)$, there exists a measurable subset $B \subseteq A$ such that $m(B) = c$. Two measurable sets $A, B \subseteq \Theta$ are said to be *a.e.-disjoint* if $\mu(A \cap B) = 0$. We use equations and inequalities involving measurable functions to denote the corresponding (measurable) *preimages* in Θ . For example, $\{f \leq 0\} := \{\theta \in \Theta : f(\theta) \leq 0\}$ and $\{f \leq g\} := \{\theta \in \Theta : f(\theta) \leq g(\theta)\}$.

Fisher market. Here, we formally describe the infinite-dimensional Fisher market setup that we use throughout our work. There are n buyers and an item space Θ , which is a compact subset of \mathbb{R}^d . Each buyer has a valuation over the item space $v_i \in L^1(\Theta)_+$ (nonnegative L^1 functions on Θ). The items’ prices $p \in L^1(\Theta)_+$ live in the same space as valuations. An allocation of items to a buyer i is denoted by $x_i \in L^\infty(\Theta)_+$. Use $x = (x_1, \dots, x_n) \in (L^\infty(\Theta)_+)^n$ to denote the aggregate allocation. An allocation x is said to be a *pure* allocation (or a pure solution, when viewed as variables of a convex program) if for all i , $x_i = \mathbf{1}_{\Theta_i}$ for *a.e.-disjoint* measurable sets $\Theta_i \subseteq \Theta$ (where leftover is possible, i.e., $\Theta \setminus (\cup_i \Theta_i) \neq \emptyset$). When x is a pure allocation (solution), we also denote x as $\{\Theta_i\}$. An allocation is *mixed* if it is not pure, or equivalently, the set $\{0 < x_i < 1\} \subseteq \Theta$ has positive measure for some i . Each buyer has a budget $B_i > 0$ and all items have unit supply, i.e., x is supply-feasible if $\sum_i x_i \leq \mathbf{1}$. Without loss of generality, we also assume that $v_i(\Theta) = \|v_i\| > 0$ for all i (otherwise buyer i can be removed). Given prices $p \in L^1(\Theta)_+$, the *demand set* of buyer i is the set of utility-maximizing allocations subject to its budget constraint:

$$D_i(p) = \arg \max \{ \langle v_i, x_i \rangle : x \in L^\infty(\Theta)_+, \langle p, x_i \rangle \leq B_i \}.$$

Generalizing its finite-dimensional counterpart (Eisenberg and Gale 1959; Eisenberg 1961; Jain and Vazirani 2007, 2010; Nisan et al. 2007), a *market equilibrium* is defined as a pair $(x^*, p^*) \in (L^\infty(\Theta)_+)^n \times L^1(\Theta)_+$ satisfying the following.

- Buyer optimality: for every $i \in [n]$, $x_i^* \in D_i(p^*)$.
- Market clearance (up to zero-price items): $\sum_i x_i^* \leq \mathbf{1}$ and $\langle p^*, \mathbf{1} - \sum_i x_i^* \rangle = 0$.

We say that $x^* \in (L^\infty(\Theta)_+)^n$ is an equilibrium allocation if (x^*, p^*) is a ME for some $p^* \in L^1(\Theta)_+$. A pair (x^*, p^*) is called a pure ME if it is a ME and x^* is a pure allocation. From the definition of market equilibrium, we can assume the following normalizations w.l.o.g. First, $v_i(\Theta) = \|v_i\| = 1$ for all i , since $D_i(p)$ is invariant under scaling of v_i . Second, $\|B\|_1 = 1$ since if (x^*, p^*) is a ME under $B = (B_i)$, then $(x^*, p^*/\|B\|_1)$ is a ME under normalized budgets $(B_i/\|B\|_1)$. Finally, The total supply of all items is $\mu(\Theta) = \|\mathbf{1}\| = 1$ (by either scaling the item space Θ via $\theta \mapsto \alpha\theta$ for some constant α or scaling the measure μ), since this scales all prices $\langle p, x_i \rangle$ and utilities $\langle v_i, x_i \rangle$ by the same constant.

Equilibrium and Duality

Due to intrinsic limitations of general infinite-dimensional convex optimization duality theory, in this case, we cannot start with a convex program and then derive its Lagrange dual (the reason will be explained in more detail later). Instead, we directly *propose* two infinite-dimensional convex programs, and then proceed to show from first principles that they exhibit optimal solutions and a strong-duality-like relationship. First, we propose a generalization of the (finite-dimensional) Eisenberg-Gale convex program (Eisenberg 1961; Nisan et al. 2007):

$$z^* = \sup_{x \in (L^\infty(\Theta)_+)^n} \sum_i B_i \log \langle v_i, x_i \rangle \text{ s.t. } \sum_i x_i \leq \mathbf{1}. \quad (\mathcal{P}_{EG})$$

Motivated by the dual of the finite-dimensional EG convex program (Cole et al. 2017, Lemma 3), we also consider the following convex program:

$$w^* = \inf_{p \in L^1(\Theta)_+, \beta \in \mathbb{R}_+^n} \left[\langle p, \mathbf{1} \rangle - \sum_i B_i \log \beta_i \right] \quad (\mathcal{D}_{EG})$$

s.t. $p \geq \beta_i v_i$ a.e., $\forall i$.

Remark. If we view (\mathcal{D}_{EG}) as the primal, then it can be shown that its Lagrange dual is (\mathcal{P}_{EG}) and weak duality follows (see, e.g., (Ponstein 2004, §3)). However, we cannot conclude strong duality, or even primal or dual optimum attainment, since $L^1(\Theta)_+$ has an empty interior (Luenberger 1997, §8.8 Problem 1) and hence Slater's condition does not hold. If we choose $L^1(\Theta)_+ = L^\infty(\Theta)$ instead of $L^1(\Theta)$ for the space of allocations x_i (i.e., the underlying Banach space of (\mathcal{P}_{EG})), then (\mathcal{D}_{EG}) , with $p \in L^\infty(\Theta)_+$ instead of $L^1(\Theta)_+$, does satisfy Slater's condition (Luenberger 1997, §8.8 Problem 2). However, its dual is (\mathcal{P}_{EG}) but with the nonnegative cone $L^\infty(\Theta)_+$ (in which each x_i lies) replaced

by the (much larger) cone $\{g \in L^\infty(\Theta)^* : \langle f, g \rangle \geq 0, \forall f \in L^\infty(\Theta)_+\} \subseteq L^\infty(\Theta)^*$. In this case, not every bounded linear functional $g \in L^\infty(\Theta)$ can be represented by a measurable function \tilde{g} such that $\langle f, g \rangle = \int \tilde{g} f d\mu$ (see, e.g., (Day 1973)). Therefore, we still cannot conclude that (\mathcal{P}_{EG}) has an optimal solution in $(L^1(\Theta)_+)^n$ satisfying strong duality. Similar dilemmas occur when (\mathcal{P}_{EG}) is viewed as the primal instead.

Nevertheless, through derivations based on first principles, we can establish optimum attainment of the convex programs, weak duality, necessary and sufficient conditions for optimality (strong duality). First, we show that the optima of (\mathcal{P}_{EG}) is attained. All proofs can be found in the extended version (see footnote 1).

Lemma 1 *The supremum z^* of (\mathcal{P}_{EG}) is attained via a pure optimal solution x^* , that is, $x^* = (x_i^*)$ and $x_i^* = \mathbf{1}_{\Theta_i}$ for a.e.-disjoint measurable subsets $\Theta_i \subseteq \Theta$.*

Unlike the finite-dimensional case, the feasible region of (\mathcal{P}_{EG}) here, although being closed and bounded in the Banach space $L^\infty(\Theta)$, is not compact (since an infinite sequence of feasible $x^{(k)}$ without a converging subsequence can be easily constructed). However, optimum attainment still holds thanks to the fact that the set of feasible utilities

$$U = \left\{ u \in \mathbb{R}_+^n : \sum_i x_i \leq \mathbf{1}, x_i \in L^\infty(\Theta)_+, i \in [n] \right\}$$

is convex and compact.

Next, we show optimum attainment for (\mathcal{D}_{EG}) by reformulating it into a finite-dimensional convex program in β . For a fixed $\beta > 0$, setting $p = \max_i \beta_i v_i$ clearly minimizes the objective of (\mathcal{D}_{EG}) subject to its constraints. Since $\beta \geq 0$, $v_i \in L^1(\Theta)_+$, we have (where $\|f\| = \int_\Theta f d\mu$ is the L^1 -norm)

$$0 \leq \max_i \beta_i v_i \leq \|\beta\|_1 \sum_i v_i,$$

where the right-hand side is L^1 -integrable since each v_i is. Hence, $\max_i \beta_i v_i \in L^1(\Theta)_+$ as well. Thus, we can eliminate p in (\mathcal{D}_{EG}) and reformulate it into following finite-dimensional convex program:

$$\inf_{\beta \in \mathbb{R}_+^n} \left[\left\langle \max_i \beta_i v_i, \mathbf{1} \right\rangle - \sum_i B_i \log \beta_i \right]. \quad (1)$$

Lemma 2 *The infimum of (1) is attained via a unique minimizer $\beta^* > 0$. The optimal solution (p^*, β^*) of (\mathcal{D}_{EG}) has a unique β^* and satisfies $p^* = \max_i \beta_i^* v_i$ a.e.*

Later, we will see that, for piecewise linear v_i , the finite-dimensional convex program (1) exhibits efficient first-order (subgradient) oracles and therefore can be solved efficiently using well-known optimization algorithms.

Due to the lack of general duality results in infinite dimensions, we first establish weak duality and KKT conditions (necessary and sufficient for optimality) in the following lemma. These conditions parallel those in nonlinear optimization in Euclidean spaces (see, e.g., (Nocedal and Wright 2006, §12.3) and (Bertsekas 1999, §3.3.1)).

Lemma 3 *Let $C = \|B\|_1 - \sum_i B_i \log B_i$. We have*

- (a) *Weak duality:* $C + z^* \leq w^*$.
(b) *KKT conditions:* For x^* feasible to (\mathcal{P}_{EG}) and (p^*, β^*) feasible to (\mathcal{D}_{EG}) , they are both optimal (i.e., attaining the optima z^* and w^* respectively) if and only if

$$\left\langle p^*, \mathbf{1} - \sum_i x_i^* \right\rangle = 0, \quad (2)$$

$$\langle v_i, x_i^* \rangle = u_i^* := \frac{B_i}{\beta_i^*}, \quad \forall i, \quad (3)$$

$$\langle p^* - \beta_i^* v_i, x_i^* \rangle = 0, \quad \forall i. \quad (4)$$

Thus, we see that in spite of the general difficulties with duality theory in infinite dimensions, we have shown that (\mathcal{P}_{EG}) and (\mathcal{D}_{EG}) behave like duals of each other: strong duality holds, and KKT conditions hold if and only if a pair of feasible solutions are both optimal (see, e.g., (Nisan et al. 2007, §5.2) for the finite-dimensional counterparts). Using Lemma 3, we can establish the equivalence of market equilibrium and optimality w.r.t. the convex programs.

Theorem 1 *Assume x^* and (p^*, β^*) are optimal solutions of (\mathcal{P}_{EG}) and (\mathcal{D}_{EG}) , respectively. Then (x^*, p^*) is a ME, $\langle p^*, x_i^* \rangle = B_i$ for all i and the equilibrium utility of buyer i is $u_i^* = \langle v_i, x_i^* \rangle = \frac{B_i}{\beta_i^*}$. Conversely, if (x^*, p^*) is a ME, then x^* is an optimal solution of (\mathcal{P}_{EG}) and (p^*, β^*) , where $\beta_i^* := \frac{B_i}{\langle v_i, x_i^* \rangle}$, is an optimal solution of (\mathcal{D}_{EG}) .*

We list some direct consequences of the results we have obtained so far. Below is a direct consequence of Theorem 1 and Part (a) of Lemma 3 on the structural properties of a market equilibrium.

Corollary 1 *Let (x^*, p^*) be a ME. Then, x^* and (p^*, β^*) , where $\beta_i^* := \frac{B_i}{\langle v_i, x_i^* \rangle}$, satisfy (2)-(4). In particular, (4) shows that a buyer's equilibrium allocation x_i^* must be zero a.e. outside its "winning" set of items $\{p^* = \beta_i^* v_i\}$.*

Remark. The equilibrium β^* , or equivalently, the second part of the unique optimal solution (p^*, β^*) of (\mathcal{D}_{EG}) , is often known as the (equilibrium) *utility price*, that is, $\beta_i^* = \frac{B_i}{u_i^*}$ is the price each buyer i pays for a unit of utility. The above corollary shows that, at equilibrium, each buyer i only gets items where its $\beta_i^* v_i$ is the maximum among all buyers, that is, where $p^* = \beta_i^* v_i$. In other words, buyer i only pays for items with the lowest price per unit utility, or equivalently, the most utility per unit price. Since $p^* \geq \beta_i^* v_i$, under prices p^* , buyer i must pay at least β_i^* for each unit of utility. From a pure optimal solution of (\mathcal{P}_{EG}) , we can construct the (a.e.-unique) optimal solution of (\mathcal{D}_{EG}) . In particular, such a construction ensures feasibility to (\mathcal{D}_{EG}) .

Corollary 2 *Let $\{\Theta_i\}$ be a pure optimal solution of (\mathcal{P}_{EG}) , $u_i^* = v_i(\Theta_i)$ and $\beta_i^* = \frac{B_i}{u_i^*}$.*

- (a) *On each Θ_i , $\beta_i^* v_i \geq \beta_j^* v_j$ a.e. for all $j \neq i$.*
(b) *Let $p^* := \max_i \beta_i^* v_i$. Then, $p^*(A) = \sum_i \beta_i^* v_i(A \cap \Theta_i)$ for any measurable set $A \subseteq \Theta$.*
(c) *The constructed (p^*, β^*) is an optimal solution of (\mathcal{D}_{EG}) and satisfies (2)-(4).*

Given a pure allocation, we can also verify whether it is an equilibrium allocation using the following corollary.

Corollary 3 *A pure allocation $\{\Theta_i\}$ is an equilibrium allocation (with equilibrium prices p^*) if and only if the following conditions hold with $\beta_i^* := \frac{B_i}{v_i(\Theta_i)}$ and $p^* := \max_i \beta_i^* v_i$.*

1. *Prices of items in Θ_i are given by $\beta_i^* v_i$: $p^* = \beta_i^* v_i$ on each Θ_i , $i \in [n]$.*
2. *Prices of leftover is zero: $p^*(\Theta \setminus (\cup_i \Theta_i)) = 0$.*

Fairness and efficiency properties of ME. Let $x \in (L^\infty(\Theta)_+)^n$, $\sum_i x_i \leq \mathbf{1}$ be an allocation. It is (strongly) *Pareto optimal* if there does not exist $\tilde{x} \in (L^\infty(\Theta)_+)^n$, $\sum_i \tilde{x}_i \leq \mathbf{1}$ such that $\langle v_i, \tilde{x}_i \rangle \geq \langle v_i, x_i \rangle$ for all i and the inequality is strict for at least one i (Cohler et al. 2011). It is *envy-free* (in a budget-weighted sense) if

$$\frac{1}{B_i} \langle v_i, x_i \rangle \geq \frac{1}{B_j} \langle v_i, x_j \rangle$$

for any $j \neq i$ (Nisan et al. 2007; Kroer et al. 2019). When all $B_i = 1$, this is sometimes referred to as being "equitable" (Weller 1985). It is *proportional* if $\langle v_i, x_i \rangle \geq \frac{B_i}{\|B\|_1} v_i(\Theta)$ for all i , that is, each buyer gets at least the utility of its *proportional share* allocation, $x^{\text{PS}} := \frac{B_i}{\|B\|_1} \mathbf{1}$. Similar to the finite-dimensional case (Jain and Vazirani 2010; Nisan and Ronen 2001), market equilibria in infinite-dimensional Fisher markets also exhibit these properties.

Theorem 2 *Let (x^*, p^*) be a ME. Then, x^* is Pareto optimal, envy-free and proportional.*

ME as generalized fair division. By Corollary 3 and Theorem 1, we can see that a pure ME $\{\Theta_i\}$ under uniform budgets ($B_i = 1/n$) is a fair division in the sense of Weller (1985), that is, a Pareto optimal and envy-free division (into a.e.-disjoint measurable subsets) of Θ . Furthermore, (Weller 1985, §3) shows that, there exist equilibrium prices p^* such that

- $p^*(\Theta_i) = 1/n$ for all i .
- $v_i(\Theta_i) \geq v_i(A)$ for any measurable set $A \subseteq \Theta$ such that $p^*(A) \leq 1/n$.
- For any measurable set $A \subseteq \Theta$, $p^*(A) = \frac{1}{n} \sum_i \frac{v_i(A \cap \Theta_i)}{v_i(\Theta_i)}$.

Utilizing our results, when $B_i = 1/n$, and $\{\Theta_i\}$ is a pure ME, the first property above is a special case of $\langle p^*, x_i^* \rangle = B_i$ in Theorem 1 (with $x_i^* = \mathbf{1}_{\Theta_i}$); the second property can be easily derived from the ME property $x_i^* \in D_i(p^*)$; the third property is a special case of Part (b) in Corollary 2, since $\beta_i^* = \frac{B_i}{u_i^*} = \frac{1}{n} \cdot \frac{1}{v_i(\Theta_i)}$. Hence, ME under a continuum of items can be viewed as generalized fair division, while our results extend those of Weller (1985).

Bounds on equilibrium quantities. We can establish upper and lower bounds on equilibrium quantities. These bounds will be useful in subsequent convergence analysis of stochastic optimization. Similar bounds hold in the finite-dimensional case (Gao and Kroer 2020). Recall that we assume $v_i(\Theta) = 1$ for all i and $\|B\|_1 = 1$ w.l.o.g.

Lemma 4 *For any ME (x^*, p^*) , we have $p^*(\Theta) = 1$. Furthermore, $B_i \leq u_i^* = \langle v_i, x_i^* \rangle \leq 1$ and hence $\underline{\beta}_i := B_i \leq \beta_i^* := \frac{B_i}{u_i^*} \leq \bar{\beta}_i := 1$ for all i .*

Efficient Optimization of (1)

In the rest of the paper, unless otherwise stated, we always use x^* or $\{\Theta_i\}$ to denote a pure equilibrium allocation. We also use β^* to denote the unique optimal solution of (1) (the equilibrium utility prices) and p^* the a.e. unique equilibrium prices which satisfy $p^* = \max_i \beta_i^* v_i$ and (2)-(4) together with x^* (Lemma 3 and Theorem 1).

The convex program (1) is finite-dimensional and has a real-valued, convex and continuous objective function (Lemma 2). By Lemma 4, we can also add the constraint $\beta \in [\beta, \bar{\beta}]$ without affecting the optimal solution. This makes the “dual” (1) more computationally tractable than its “primal” (\mathcal{P}_{EG}).

Ellipsoid method for piecewise linear v_i .² Assume that each v_i is K_i -piecewise linear (possibly discontinuous). There are in total $K = \sum_i K_i$ pieces. We show that, for piecewise linear (p.w.l.) v_i over $\Theta = [0, 1]$, we can compute a solution $\tilde{\beta}$ such that $\|\tilde{\beta} - \beta^*\| \leq \epsilon$ (all norms for finite-dimensional vectors are Euclidean 2-norms unless otherwise specified) in time polynomial in $\log \frac{1}{\epsilon}$, n and $K = \sum_i K_i$. This is achieved via solving (1) using the ellipsoid method. Consider the following generic convex program (Ben-Tal and Nemirovski 2019, §4.1.4):

$$f^* := \min_x f(x) \text{ s.t. } x \in X \quad (5)$$

where f is convex and continuous (and hence subdifferentiable) on a convex compact $X \subseteq \mathbb{R}^n$. Assume we have access to the following oracles:

- The *separation* oracle \mathcal{S} : given any $x \in \mathbb{R}^n$, either report $x \in \text{int } X$ or return a $g \neq 0$ (representing a separating hyperplane) such that $\langle g, x \rangle \geq \langle g, y \rangle$ for any $y \in X$.
- The *first-order* or *subgradient* oracle \mathcal{G} : given $x \in \text{int } X$ (the interior of X), return a subgradient $f'(x)$ of f at x , that is, $f(y) \geq f(x) + \langle f'(x), y - x \rangle$ for any y .

The time complexity of the ellipsoid method is as follows.

Theorem 3 (Ben-Tal and Nemirovski 2019, Theorem 4.1.2) *Let $V = \max_{x \in X} f(x) - f^*$, $R = \sup_{x \in X} \|x\|$, and $r > 0$ be the radius of a Euclidean ball contained in X . For any $\epsilon > 0$, it is possible to find an ϵ -solution x_ϵ (i.e., $f(x_\epsilon) \leq f^* + \epsilon$) with no more than $N(\epsilon)$ calls to \mathcal{S} and \mathcal{G} , followed by no more than $O(1)n^2N(\epsilon)$ arithmetic operations to process the answer of the oracles, where $N(\epsilon) = O(1)n^2 \log(2 + \frac{VR}{\epsilon r})$.*

In order to make use of the ellipsoid method for (\mathcal{D}_{EG}) for p.w.l. v_i , we need to derive efficient oracles \mathcal{S} and \mathcal{G} . To this end, we need some elementary lemmas regarding p.w.l. linear functions.

Lemma 5 *For any $\beta \in \mathbb{R}_+^n$, the function $\theta \mapsto \max_i \beta_i v_i(\theta)$ is piecewise linear with at most $n(K - n + 1)$ pieces.*

Lemma 6 *Suppose $f_i(\theta) = c_i\theta + d_i \geq 0$, for all $\theta \in [l, u] \subseteq [0, 1]$, $i \in [n]$. Then, $h_n(\theta) = \max_i f_i(\theta)$ is piecewise linear on $[l, u]$ with at most n pieces. Furthermore, the breakpoints*

²In the extended version (see the previous footnote), we give a compact finite-dimensional convex conic reformulation of (\mathcal{P}_{EG}) which can be solved using off-the-shelf optimization software.

of h_n , $l = a_0 < a_1 < \dots < a_{n'} = u$ ($n' \leq n$) can be found in $O(n^2)$ time.

Denote $\phi(\beta) = \langle \max_i \beta_i v_i, \mathbf{1} \rangle$, which can be easily seen to be finite, convex and continuous on \mathbb{R}_+^n . Hence, it is subdifferentiable on \mathbb{R}_{++}^n (Ben-Tal and Nemirovski 2019, Proposition C.6.5). First, we show that, if all v_i are linear on a common interval and zero otherwise, a subgradient of $\phi(\beta)$ can be constructed in $O(n^2)$ time. This utilizes the additivity (in terms of integration or expectation) property of subgradients, as formalized in the following lemma. Here, $\Theta \subseteq \mathbb{R}^d$ can be a general compact set and $\mathbf{e}^{(i)}$ is the i th unit vector in \mathbb{R}^d .

Lemma 7 *Let $f(\beta, \theta) = \max_i \beta_i v_i(\theta)$. For any $\theta \in \Theta$, a subgradient of $f(\cdot, \theta)$ at β is $g(\beta, \theta) = v_{i^*}(\theta) \mathbf{e}^{(i^*)}$, where $i^* \in \arg \max_i \beta_i v_i(\theta)$ (taking the smallest index if there is a tie). Hence, a subgradient of $\phi(\beta)$ is $\phi'(\beta) = \int_{\Theta} g(\beta, \theta) d\theta = \mu(\Theta) \cdot \mathbb{E}_{\theta} g(\beta, \theta)$, where the expectation is over $\theta \sim \text{Unif}(\Theta)$.*

Using Lemma 7 and the p.w.l. structure of v_i , we have the following for computing a subgradient of ϕ .

Lemma 8 *For each i , assume that $v_i(\theta) = c_i\theta + d_i \geq 0$ on an interval $[l, u] \subseteq [0, 1]$.*

- The function $\theta \mapsto \max_i \beta_i v_i(\theta)$ has at most n linear pieces on $[l, u]$, with breakpoints $l = a_0 < a_1 < \dots < a_{n'} = \dots = a_n = u$, $n' \leq n$ (depending on β).
- We can construct $\phi'(\beta) \in \partial\phi(\beta)$ for any $\beta > 0$ as follows: the i th component of $\phi'(\beta)$ is

$$\sum_{k \in [n'] : i_k^* = i} \left(\frac{c_{i_k^*}}{2} (a_k^2 - a_{k-1}^2) + d_{i_k^*} (a_k - a_{k-1}) \right),$$

where i_k^* is the (unique) winner (with the smallest index among ties) on $[a_{k-1}, a_k]$.

- The above construction of $\phi'(\beta)$ takes $O(n^2)$ time.

When v_i are K_i -piecewise linear on $[0, 1]$, using Lemma 8, we can compute a subgradient $\phi'(\beta)$ by summing up the above construction over the intervals given by the breakpoints of all v_i , and there are at most K such intervals.

Theorem 4 *For any $\beta > 0$, a subgradient $\phi'(\beta)$ can be computed in $O(n^2K)$ time.*

Combining the above results, we have the following overall time complexity. Again, we assume that $v_i(\theta) = 1$ and $\|B\|_1 = 1$ (w.l.o.g.). For general v_i and B_i that do not satisfy this, we can normalize them in $O(nK)$ time.

Theorem 5 *Let $\Theta = [0, 1]$, $v_i(\theta) = 1$ for all i , $\|B\|_1 = 1$ and $\epsilon > 0$. A solution $\tilde{\beta}$ such that $\|\tilde{\beta} - \beta^*\| \leq \epsilon$ can be computed in $O\left(n^4 K \log \frac{n \cdot \max_i B_i}{\epsilon \cdot \min_i B_i}\right)$ time, which is $O\left(n^4 K \log \frac{n}{\epsilon}\right)$ when $B_i = 1/n$ for all i .*

The ellipsoid method can be applied to (1) more generally than for the case of p.w.l. v_i . As long as we can compute $\phi'(\beta)$ in time polynomial in nK , it finds a solution β that is ϵ -close to β^* in time $\log \frac{1}{\epsilon}$ via the same ellipsoid method framework. By Lemma 7, since a “pointwise” subgradient $g(\beta, \theta)$ of $f(\beta, \theta)$ is much easier to compute, as long as the

Algorithm 1: Stochastic dual averaging (SDA)

Initialize: Choose $\beta^1 \in \text{dom } \Psi$ and $\bar{g}^0 = 0$
for $t = 1, 2, \dots$ **do**
 Sample $\theta_t \sim \mathcal{D}$ and compute $g_t \in \partial_\beta f(\beta, \theta_t)$
 $\bar{g}^t = \frac{t-1}{t} \bar{g}^{t-1} + \frac{1}{t} g_t$
 $\beta^{t+1} = \arg \min_\beta \{ \langle \bar{g}^t, \beta \rangle + \Psi(\beta) \}$ (*)

integral $\int_\Theta g(\beta, \theta) d\theta$ can be evaluated efficiently, we can compute a “full” subgradient $\phi'(\beta)$ efficiently.

Stochastic optimization for general Θ and v_i . When a full subgradient $\phi'(\beta)$ is difficult to compute, we can still utilize the expectation characterization in Lemma 7 to use a stochastic optimization algorithm to solve (1). The problem structure is particularly suitable for the *stochastic dual averaging* (SDA) algorithm (Xiao 2010; Nesterov 2009). It solves problems of the following form:

$$\min_\beta \mathbb{E}_\theta f(\beta, \theta) + \Psi(\beta), \quad (6)$$

where Ψ is a strongly convex regularization function such that $\text{dom } \Psi = \{\beta : \Psi(\beta) < \infty\}$ is closed. Let $\theta \sim \mathcal{D}$ be a random variable with distribution \mathcal{D} and $f(\cdot, \theta)$ be convex and subdifferentiable on $\text{dom } \Psi$ for all $\theta \in \Theta$. The algorithm works as follows (Xiao 2010, Algorithm 1) (where we assume $\mu(\Theta) = 1$ w.l.o.g.).

To solve (1), we set $f(\beta, \theta) = \max_i \beta_i v_i(\theta)$ and $\mathcal{D} = \text{Unif}(\Theta)$. By Lemma 7, we can choose $g^t = g(\beta, \theta_t) \in \partial_\beta f(\beta, \theta_t)$. Let $\Psi(\beta) = -\sum_i B_i \log \beta_i$ if $\beta \in [\beta, \bar{\beta}]$ and $= \infty$ o.w. (that is, $\text{dom } \Psi = [\beta, \bar{\beta}]$). Given these specifications, in Algorithm 1, the step (*) yields a simple, explicit update: at iteration t , compute $\beta_i^{t+1} = \Pi_{[\beta_i, \bar{\beta}_i]} \left(\frac{B_i}{\bar{g}_i^t} \right)$ for all i , where $\Pi_{[a,b]}(c) = \min\{\max\{a, c\}, b\}$ is the projection onto a closed interval. It can be derived easily from its first-order optimality condition. Using the convergence results in (Xiao 2010) for strongly convex Ψ , we can show that the uniform average of all β^t generated by SDA converges to β^* both in mean square error (MSE) and with high probability, with mild finiteness assumptions on v_i .

Theorem 6 Assume $v_i \in L_2(\Theta)$, that is, $\langle v_i^2, \mathbf{1} \rangle = \mathbb{E}_\theta [v_i(\theta)^2] < \infty$ for all i . Let $G^2 := \mathbb{E}_\theta [\max_i v_i(\theta)^2] < \infty$ and $\sigma = \min_i B_i > 0$. Let $\tilde{\beta}^t := \frac{1}{t} \sum_{\tau=1}^t \beta^\tau$. Then,

$$\mathbb{E} \|\tilde{\beta}^t - \beta^*\|^2 \leq \frac{6(1+\log t) + \frac{1}{2}(\log t)^2}{t} \times \frac{G^2}{\sigma^2}.$$

Next, further assume that $v_i \leq G$ a.e. for all i . Then, for any $\delta > 0$, with probability at least $1 - 4\delta \log t$, we have $\|\tilde{\beta}^t - \beta^*\|^2 \leq \frac{2M_t}{\sigma}$, where

$$M_t = \frac{\Delta_t}{t} + \frac{4G}{t} \sqrt{\frac{\Delta_t \log(1/\delta)}{\sigma}} + \max \left\{ \frac{16G^2}{\sigma}, 6V \right\} \frac{\log(1/\delta)}{t},$$

$$\Delta_t = \frac{G^2}{2\sigma} (6 + \log t) \text{ and } V = \frac{2n}{\min_i B_i}.$$

Remark In the above theorem, the bound on $\mathbb{E} \|\tilde{\beta}^t - \beta^*\|^2$ (MSE) is of order $O\left(\frac{(\log t)^2}{t}\right)$, where the constant degrades upon buyer heterogeneity, i.e., a smaller $\min_i B_i$ leads to a larger bound (recall that $\|B\|_1 = 1$ and therefore $\min_i B_i \leq \frac{1}{n}$). For the second half regarding $\|\tilde{\beta}^t - \beta^*\|^2$, substituting

$\delta = \frac{1}{t^\alpha}$ ($\alpha \geq 1$) yields a bound of order $O\left(\frac{\log t}{t}\right)$ (also depending inversely on $\min_i B_i$), with probability at least $1 - \frac{4 \log t}{t^\alpha}$. In addition, the added assumptions $\mathbb{E}_\theta [\max_i v_i^2] < \infty$ and $v_i \leq G$ a.e. for all i are always satisfied as long as they are (a.e.) bounded (e.g., p.w.l. functions).

Deterministic optimization using $\phi'(\beta)$. When $\phi'(\beta)$ can be computed, such as when v_i are piecewise linear on $\Theta = [0, 1]$ (Theorem 4), in Algorithm 1, we can replace g^t with a full subgradient $\phi'(\beta^t)$. Then, $\|\beta^t - \beta^*\|^2$ is deterministic and bounded by the same right hand side as the first half of Theorem 6. In this case, if $v_i \leq G$ a.e., then it can be easily verified that $\|\phi'(\beta)\|^2 \leq nG^2 < \infty$ for all $\beta > 0$. Then, we can also use a projected subgradient descent method that achieves $\|\hat{\beta}^t - \beta^*\|^2 = O(n/t)$, where $\hat{\beta}^t$ is a weighted average of β^1, \dots, β^t (see, e.g., (Lacoste-Julien, Schmidt, and Bach 2012) and (Bubeck 2015, Theorem 3.9)).

Approximate equilibrium prices. Suppose we have obtained an approximate solution $\tilde{\beta}$ such that $\|\tilde{\beta} - \beta^*\| \leq \epsilon$. Define $\tilde{p} = \max_i \tilde{\beta}_i v_i \in L_1(\Theta)_+$, which satisfies

$$\begin{aligned} \|\tilde{p} - p^*\| &= \int_\Theta \left| \max_i \tilde{\beta}_i v_i(\theta) - \max_i \beta_i^* v_i(\theta) \right| d\theta \\ &\leq \|\tilde{\beta} - \beta^*\|_\infty \sum_i \|v_i\| \leq n\epsilon. \end{aligned}$$

Recall that for any $\tilde{\beta}$, the prices \tilde{p} defined above is a p.w.l. function with at most $n(K - n + 1)$ pieces (Lemma 5). By Lemma 6, finding its p.w.l. representation (breakpoints and linear coefficients on each piece) takes $O(n^2 K)$ time. Therefore, under the same time complexity as in Theorem 5 (where the additional factor n is inside log and is absorbed into the constant), we can compute an approximate equilibrium prices \tilde{p} such that $\|\tilde{p} - p^*\| \leq \epsilon$. Furthermore, under prices \tilde{p} , an equilibrium allocation x^* may (slightly) violate buyers' budget constraints:

$$\langle \tilde{p}, x^* \rangle = \langle p^*, x^* \rangle + \langle \tilde{p} - p^*, x^* \rangle \leq B_i + \|\tilde{p} - p^*\| = B_i + n\epsilon,$$

where the inequality uses Theorem 1 (budget of buyer i depleted, i.e., $\langle p^*, x_i^* \rangle = 1$) and $x_i^* \leq \mathbf{1}$ (x^* is feasible w.r.t. item supplies). Hence, consider the allocation $\tilde{x}_i = \frac{B_i}{B_i + n\epsilon} x_i^*$ for all i . This allocation clearly satisfies $\sum_i \tilde{x}_i \leq \mathbf{1}$ and, for each i , its budget constraint is satisfied: $\langle \tilde{p}, \tilde{x}_i \rangle = \frac{B_i}{B_i + n\epsilon} \langle \tilde{p}, x^* \rangle \leq B_i$. The utility of buyer i from \tilde{x}_i is $\tilde{u}_i = \langle v_i, \tilde{x}_i \rangle = \frac{B_i}{B_i + n\epsilon} u_i^*$, which is close to the equilibrium utility u_i^* as long as $n\epsilon \ll B_i$.

Construction of an Equilibrium Allocation

Throughout this section, same as before, β^* denotes the optimal solution of (1), $p^* = \max_i \beta_i^* v_i$ is the equilibrium prices and $u_i^* = \frac{B_i}{\beta_i^*}$ is the equilibrium utility. Given β^* , we can construct a pure equilibrium allocation x^* explicitly if we allow arbitrary division of measurable subsets. This is possible only if v_i , and hence p^* , are atomless.

Theorem 7 For any $S \subseteq [n]$, define

$$\Theta_S = \begin{cases} \{p^* = \beta_i^* v_i > \beta_\ell^* v_\ell, \forall i \in S, \ell \notin S\} & \text{if } S \neq \emptyset, \\ \{p^* = \beta_i^* v_i = 0, \forall i \in [n]\} & \text{if } S = \emptyset. \end{cases}$$

- $\{\Theta_S\}_{S \subseteq [n]}$ is a measurable partition of Θ .
- There exists $b = (b_{i,S}) \geq 0$, $i \in [n]$, $S \subseteq [n]$ such that (i) $b_{i,S} = 0$ if $i \notin S$, (ii) $\sum_i b_{i,S} = p^*(\Theta_S)$ for all S and (iii) $\sum_S b_{i,S} = B_i$ for all i .
- We can partition each Θ_S into measurable subsets $\Theta_{i,S}$, $i \in S$ such that $p^*(\Theta_{i,S}) = b_{i,S}$ (meanwhile, $\Theta_{i,S} := \emptyset$ if $i \notin S$). Let $\Theta_i = \cup_S \Theta_{i,S}$ (the leftover Θ_\emptyset can be assigned arbitrarily). The resulting pure allocation $\{\Theta_i\}$ is an optimal solution of (\mathcal{P}_{EG}) and hence, by Theorem 1, a pure equilibrium allocation.

The above construction of x^* requires constructing Θ_S , $S \subseteq [n]$ via partitioning measurable sets and solving a large linear system of size 2^n . In some cases, an equilibrium allocation (or even a pure one) can be constructed easily.

Finite item space Θ . In this case, simply solve a system of linear equations $\sum_{\theta \in \Theta} v_i(\theta)x_i(\theta) = u_i^*$, $i \in [n]$ and $\sum_i x_i(\theta) \leq 1$, $\theta \in \Theta$ to obtain an equilibrium allocation. It has a solution since an equilibrium allocation exists.

P.w.l. valuations. For the case of p.w.l. v_i , p^* can have at most $n(K - n + 1)$ linear pieces with breakpoints $0 = a_0 < a_1 < \dots < a_N = 1$, where $N \leq n(K - n + 1)$. These breakpoints consist of (i) the “static” breakpoints of all v_i and (ii) the intersections of the linear pieces of $\beta_i^* v_i$ between these breakpoints. Each interval between two breakpoints, on which every v_i is linear, is further divided into at most n small subintervals (Lemma 6), leading to the subintervals $[a_{k-1}, a_k]$, $k \in [N]$ above. On each $[a_{k-1}, a_k]$, for $i \neq j$, the line segments $\beta_i^* v_i$ and $\beta_j^* v_j$ either separate completely or overlapping completely, except possibly at endpoints. For $k \in [N]$, let $I_k \neq \emptyset$ be the set of “winners” on $[a_{k-1}, a_k]$, that is, for $i \in I_k$ and $j \notin I_k$, $p^* = \beta_i^* v_i > \beta_j^* v_j$ on (a_{k-1}, a_k) . In the language of Theorem 7, we only need to consider $\Theta_{i,k}$, $i \in [n]$, $k \in [N]$ such that $i \in I_k$. This is because x_i^* must *not* allocate anything on $[a_{k-1}, a_k]$ such that $p^* > \beta_i^* v_i$ (i.e., $x_i^* = 0$ on $[a_{k-1}, a_k]$ if $i \notin I_k$) (see (4) in Lemma 3). Therefore, we only need to solve for $b = (b_{i,k}) \in \mathbb{R}_+^{n \times N}$. Meanwhile, $p^*([a_{k-1}, a_k]) = \beta_i^* v_i([a_{k-1}, a_k])$ (for some $i \in I_k$) can be easily computed, since v_i is linear on $[a_{k-1}, a_k]$. Splitting each $\Theta_k := [a_{k-1}, a_k]$ into $\Theta_{i,k}$, $i \in I_k$ according to each winner’s $b_{i,k}$ is also trivial: since $p^* = \beta_i^* v_i$ ($i \in I_k$) is linear on $[a_{k-1}, a_k]$, we can partition $[a_{k-1}, a_k]$ into consecutive intervals $\Theta_{i,k}$, each having $p^*(\Theta_{i,k}) = b_{i,k}$, $i \in I_k$.

An illustrative example. Consider n buyers with distinct linear valuations v_i on the item space $\Theta = [0, 1]$ such that $v_i(\Theta) = 1$ for all i . In this case, $p^* = \max_i \beta_i^* v_i$ is piecewise linear with *exactly* n pieces, since (i) it has at most n pieces by Lemma 6 and (ii) each buyer i “wins” at least one piece, i.e., $p^* = \beta_i^* v_i$ on a nonempty closed interval (otherwise, by (4) in Lemma 3, buyer i has $x_i^* = 0$ and gets $u_i^* = 0$, contradicting to Lemma 4). We construct a small instance with $n = 4$ and solve its convex program (1) using Algorithm 1 (SDA) to obtain an approximate solution $\tilde{\beta} \approx \beta^*$. Then, we construct an approximate equilibrium allocation, which is given by the 4 intervals as well as the (unique) winners of each interval, as shown in Figure 1. The $n = 4$ intervals are, from left to right, allocated to buyers 3,




Figure 1: A division of $[0, 1]$ given by $\tilde{\beta}$, where each buyer i gets the interval shown in the legend, on which $\tilde{p} = \tilde{\beta}_i v_i$

2, 4, 1, respectively.

Discussion and Conclusions

Motivated by applications in ad auctions and fair recommender systems, we considered infinite-dimensional Fisher markets with a continuum of items and the concept of a market equilibrium in this setting. We proposed infinite-dimensional Eisenberg-Gale-type convex programs whose optimal solutions are ME, and vice versa. We established existence of optimal solutions (and hence existence of ME) and optimality conditions of the convex programs that parallel various structural properties of ME. We also showed that a ME exhibits various efficiency and fairness guarantees. Utilizing a finite-dimensional reformulation of a convex program, we proposed efficient optimization algorithms for computing approximate equilibrium utility prices. Lastly, we discussed the construction of an equilibrium allocation from the optimal solution and gave an illustrative example.

References

- Aziz, H.; and Mackenzie, S. 2016. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, 416–427. IEEE.
- Aziz, H.; and Ye, C. 2014. Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In *International Conference on Web and Internet Economics*, 1–14. Springer.
- Balseiro, S. R.; Besbes, O.; and Weintraub, G. Y. 2015. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science* 61(4): 864–884.
- Ben-Tal, A.; and Nemirovski, A. 2019. Lectures on modern convex optimization. *Online version: http://www2.isye.gatech.edu/~nemirovs/Lect_ModConvOpt*.
- Bertsekas, D. 1999. *Nonlinear programming*. Athena Scientific, 2nd edition.
- Brams, S. J.; Feldman, M.; Lai, J. K.; Morgenstern, J.; and Procaccia, A. D. 2012. On maxsum fair cake divisions. In *Twenty-sixth AAAI Conference on Artificial Intelligence*.

- Brams, S. J.; and Taylor, A. D. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press.
- Bubeck, S. 2015. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning* 8(3-4): 231–357.
- Chen, Y.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2013. Truth, justice, and cake cutting. *Games and Economic Behavior* 77(1): 284–297.
- Cohler, Y. J.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. 2011. Optimal envy-free cake cutting .
- Cole, R.; Devanur, N. R.; Gkatzelis, V.; Jain, K.; Mai, T.; Vazirani, V. V.; and Yazdanbod, S. 2017. Convex program duality, fisher markets, and Nash social welfare. In *18th ACM Conference on Economics and Computation, EC 2017*. Association for Computing Machinery, Inc.
- Conitzer, V.; Kroer, C.; Panigrahi, D.; Schrijvers, O.; Sodomka, E.; Stier-Moses, N. E.; and Wilkens, C. 2019. Pacing Equilibrium in First-Price Auction Markets. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. ACM.
- Conitzer, V.; Kroer, C.; Sodomka, E.; and Stier-Moses, N. E. 2018. Multiplicative Pacing Equilibria in Auction Markets. In *International Conference on Web and Internet Economics*.
- Day, M. M. 1973. Normed linear spaces. In *Normed Linear Spaces*, 27–52. Springer.
- Eisenberg, E. 1961. Aggregation of utility functions. *Management Science* 7(4): 337–350.
- Eisenberg, E.; and Gale, D. 1959. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics* 30(1): 165–168.
- Gao, Y.; and Kroer, C. 2020. First-order methods for large-scale market equilibrium computation. *arXiv preprint arXiv:2006.06747* .
- Jain, K.; and Vazirani, V. V. 2007. Eisenberg-gale markets: Algorithms and structural properties. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, 364–373.
- Jain, K.; and Vazirani, V. V. 2010. Eisenberg–Gale markets: Algorithms and game-theoretic properties. *Games and Economic Behavior* 70(1): 84–106.
- Kroer, C.; and Peysakhovich, A. 2019. Scalable Fair Division for At Most One Preferences. *arXiv preprint arXiv:1909.10925* .
- Kroer, C.; Peysakhovich, A.; Sodomka, E.; and Stier-Moses, N. E. 2019. Computing large market equilibria using abstractions. *arXiv preprint arXiv:1901.06230* .
- Lacoste-Julien, S.; Schmidt, M.; and Bach, F. 2012. A simpler approach to obtaining an $O(1/t)$ convergence rate for the projected stochastic subgradient method. *arXiv preprint arXiv:1212.2002* .
- Legut, J. 2017. Optimal fair division for measures with piecewise linear density functions. *International Game Theory Review* 19(02): 1750009.
- Legut, J. 2020. How to obtain an equitable optimal fair division. *Annals of Operations Research* 284(1): 323–332.
- Luenberger, D. G. 1997. *Optimization by vector space methods*. John Wiley & Sons.
- Nesterov, Y. 2009. Primal-dual subgradient methods for convex problems. *Mathematical programming* 120(1): 221–259.
- Nisan, N.; and Ronen, A. 2001. Algorithmic mechanism design. *Games and Economic behavior* 35(1-2): 166–196.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic game theory*. Cambridge University Press.
- Nocedal, J.; and Wright, S. 2006. *Numerical optimization*. Springer Science & Business Media.
- Ponstein, J. 2004. *Approaches to the Theory of Optimization*, volume 77. Cambridge University Press.
- Procaccia, A. D. 2013. Cake cutting: not just child’s play. *Communications of the ACM* 56(7): 78–87.
- Procaccia, A. D. 2016. Cake cutting algorithms. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of computational social choice*, chapter 13, 311–332. Cambridge University Press.
- Weller, D. 1985. Fair division of a measurable space. *Journal of Mathematical Economics* 14(1): 5–17.
- Xiao, L. 2010. Dual Averaging Methods for Regularized Stochastic Learning and Online Optimization. *Journal of Machine Learning Research* 11: 2543–2596.