

Fisher Markets with Social Influence

Jiayi Zhao¹, Denizalp Goktas², Amy Greenwald²

¹ Department of Computer Science, Pomona College

² Department of Computer Science, Brown University

jzae2019@mymail.pomona.edu, denizalp_goktas@brown.edu, amy_greenwald@brown.edu

Abstract

A Fisher market is an economic model of buyer and seller interactions in which each buyer's utility depends only on the bundle of goods she obtains. Many people's interests, however, are affected by their social interactions with others. In this paper, we introduce a generalization of Fisher markets, namely influence Fisher markets, which captures the impact of social influence on buyers' utilities. We show that competitive equilibria in influence Fisher markets correspond to generalized Nash equilibria in an associated pseudo-game, which implies the existence of competitive equilibria in all influence Fisher markets with continuous and concave utility functions. We then construct a monotone pseudo-game, whose variational equilibria and their duals together characterize competitive equilibria in influence Fisher markets with continuous, jointly concave, and homogeneous utility functions. This observation implies that competitive equilibria in these markets can be computed in polynomial time under standard smoothness assumptions on the utility functions. The dual of this second pseudo-game enables us to interpret the competitive equilibria of influence CCH Fisher markets as the solutions to a system of simultaneous Stackelberg games. Finally, we derive a novel first-order method that solves this Stackelberg system in polynomial time, prove that it is equivalent to computing competitive equilibrium prices via *tâtonnement*, and run experiments that confirm our theoretical results.

Introduction

The branch of mathematical economics that attempts to explain the behavioral relationship among supply, demand, and prices via equilibria dates back to the work of French economist (Walras 1896), and today is known as general equilibrium theory (Mas-Colell, Whinston, and Green 1995). One of the seminal achievements in this area is the proof of existence of competitive equilibrium prices in Arrow-Debreu markets (1954). In such a market, traders seek to "purchase" goods from others, by exchanging a part of their endowment of goods for various other goods. A competitive equilibrium comprises an allocation of goods to traders together with good prices such that traders maximize their preferences over goods while ensuring that their spending does not exceed the value of their endowment, and the market clears: i.e., no more goods are allocated than the

market supply and Walras' law holds, meaning the value of demand equals the value of supply.

In much of mainstream consumer theory (Mas-Colell, Whinston, and Green 1995), and in Arrow-Debreu markets, each trader's preference depends only on its own consumption. Such models fail to capture the influence of social interactions on traders' interests. For example, the more friends one has who own an iPhone, the more one might prefer an iPhone. Likewise, if a celebrity, e.g., Beyonce, wears a particular brand of bag, e.g., Telfar, then one's preference for that brand of bag might increase. In an age of densifying social networks, it is becoming more and more essential that our economic models capture the effects of social interactions on individuals' preferences.

To try to better understand the implications of social networks on market equilibria, Chen and Teng (2011) recently proposed an extension of the Arrow-Debreu market model in which each trader's preference is influenced by the goods her neighbors obtains: **the Arrow-Debreu market with social influence**. Formally, Chen and Teng's model augments an Arrow-Debreu market with a social network connecting the traders, and then embeds this network's structure in each trader's utility function, thus inducing a preference relation over allocations of goods that depends both on the trader's and its neighbors' allocations. The authors then study a modest generalization of competitive equilibrium in which traders maximize their utility, assuming the allocations of the other traders in the market, including their neighbors, are fixed.

Chen and Teng analyze their model under two specific types of utilities: linear and threshold influence functions. They prove existence of competitive equilibrium in their setting, when the graph underlying the economy is strongly connected and the utility functions' parameters guarantee non-satiation of the preferences they represent. Under additional assumptions on the topology of the network, they also provide polynomial-time methods for computing competitive equilibria.

As the computation of competitive equilibrium in Arrow-Debreu markets is believed to be intractable, i.e., it is PPAD-complete (Chen and Deng 2006; Chen and Teng 2009), it seems unlikely that we can obtain positive computational results in such a broad setting. In the last two decades, however, **Fisher markets** have emerged as an interesting special

case of Arrow-Debreu markets in which competitive equilibria can be efficiently computed. The Fisher market is a one-sided Arrow-Debreu market comprising one seller and multiple buyers, the latter of whom are endowed with an artificial currency called their budget, rather than an endowment of goods.

During the last two decades, a wide array of polynomial-time computability results have been established for Fisher markets (Devanur et al. 2002; Jain, Vazirani, and Ye 2005; Gao and Kroer 2020; Goktas and Greenwald 2021). One of the most interesting findings is the observation that the primal and dual solutions, respectively, to the **Eisenberg-Gale convex program** (Eisenberg and Gale 1959), constitute competitive equilibrium allocations and prices in Fisher markets, and are computable in polynomial time assuming buyers with continuous, concave, and homogeneous utility functions representing locally non-satiated preferences (Devanur et al. 2002; Devanur et al. 2008; Jain, Vazirani, and Ye 2005). Moreover, Cheung, Cole, and Devanur show that solving the dual of the Eisenberg-Gale program via (sub)gradient descent amounts to solving the market via *tâtonnement*, an economic price-adjustment process dating back to Walras (1896), in which a fictional auctioneer increases (resp. decreases) the prices of goods that are overdemanded (resp. underdemanded) (Cheung, Cole, and Devanur 2013). Furthermore, Goktas and Greenwald (2021) show that the dual of the Eisenberg-Gale program corresponds to a zero-sum Stackelberg game, in which *tâtonnement* surfaces as a no-regret learning dynamic for the auctioneer (Goktas, Zhao, and Greenwald 2022).

With the aim of obtaining stronger results on the existence and computation of competitive equilibrium in markets with social influence, we introduce a special case of the Arrow-Debreu market with social influence and a generalization of Fisher markets (Brainard, Scarf et al. 2000), which we call **Fisher markets with social influence**, or **influence Fisher markets** for short. An influence Fisher market, as the name suggests, is a Fisher market in which buyers' utility functions depend not only on their own allocation, but also on their neighbors'. In this paper, we provide existence and polynomial-time computability results for competitive equilibrium in influence Fisher markets. We first extend Arrow and Debreu's competitive equilibrium existence argument using their theory of pseudo-games to prove that a competitive equilibrium exists in all influence Fisher markets with continuous utility functions that are concave in each buyer's allocation. Contrary to Chen and Teng, our existence result makes no assumptions about the topology of the network.

Next, for all influence Fisher markets with continuous and homogeneous utilities, we construct a similar pseudo-game with jointly convex constraints whose variational equilibria correspond to competitive equilibrium allocations. This pseudo-game is monotone assuming the buyers' utility functions are jointly-concave; thus, we can solve for its variational equilibria as a variational inequality problem (Facchinei, Fischer, and Piccialli 2007). This approach yields a polynomial-time algorithm that computes competitive equilibrium allocations in influence Fisher markets. Moreover, as the pseudo-game comprises n different opti-

mization problems, one per buyer, there are correspondingly n duals. Surprisingly, the solutions to all of these duals yield the same competitive equilibrium prices!

Finally, following Goktas and Greenwald (2021), who reformulate the dual of the Eisenberg-Gale program as a zero-sum Stackelberg game, we likewise reformulate the n duals of our pseudo-game as a system of n simultaneous zero-sum Stackelberg games. In Goktas and Greenwald's dual, the leader is a fictitious auctioneer who sets prices, while the followers are a set of buyers who effectively play as a team; in our n duals, each leader is again a fictitious auctioneer, but each follower is an individual buyer who best responds to the auctioneer's prices, *given the other buyers' allocations*. Thus, the buyers in this system play a Nash equilibrium. Also following Goktas and Greenwald, we show that running subgradient descent on each leader's value function, i.e., the leader's utility function assuming the follower best-responds, amounts to solving the market via *tâtonnement* in polynomial-time, as in (standard) Fisher markets. The main difference between our algorithm and theirs is that ours requires a Nash-equilibrium oracle, so that, given prices *and the other buyers' allocations*, buyers can play best responses to one another.

Related Work Gao and Kroer (2020) studied an alternative family of first-order methods for solving Fisher markets (only; not min-max Stackelberg games more generally), assuming linear, quasilinear, and Leontief utilities; such methods can be more efficient when markets are large.

Following Arrow and Debreu's introduction of GNE, Rosen (1965) initiated the study of the mathematical and computational properties of GNE in pseudo-games with jointly convex constraints, proposing a projected gradient method to compute GNE. Thirty years later, Uryas'ev and Rubinstein (1994) developed the first relaxation methods for finding GNEs, which were improved upon in subsequent works (Krawczyk and Uryasev 2000; Heusinger and Kanzow 2009). Two other types of algorithms were also introduced to the literature: Newton-style methods (Facchinei, Fischer, and Piccialli 2009; Dreves 2017; von Heusinger, Kanzow, and Fukushima 2012; Izmailov and Solodov 2014; Fischer et al. 2016; Dreves et al. 2013) and interior-point potential methods (Dreves et al. 2013). Many of these approaches are based on minimizing the exploitability of the pseudo-game, but others use variational inequality (Facchinei, Fischer, and Piccialli 2007; Nabetani, Tseng, and Fukushima 2011) and Lemke methods (Schiro, Pang, and Shanbhag 2013). Recently, this literature has established convergence guarantees for exploitability minimization (Goktas and Greenwald 2022) and relaxation (Jordan, Lin, and Zampetakis 2023) methods.

Preliminaries

In this section, we define our main modeling tool, pseudo-games, and then we introduce our object of study, Fisher markets with social influence, as a particular pseudo-game.

Notation

We use calligraphic uppercase letters to denote sets and set correspondences (e.g., \mathcal{X}); bold lowercase letters to denote vectors (e.g., \mathbf{p} , $\boldsymbol{\pi}$); bold uppercase letters to denote matrices and vector-valued random variables (e.g., \mathbf{X} , $\boldsymbol{\Gamma}$); lowercase letters to denote scalar quantities (e.g., x , γ); and uppercase letters to denote scalar-valued random variables (e.g., X , Γ). We denote the i th row vector of a matrix (e.g., \mathbf{X}) by the corresponding bold lowercase letter with subscript i (e.g., \mathbf{x}_i). Similarly, we denote the j th entry of a vector (e.g., \mathbf{p} or \mathbf{x}_i) by the corresponding lowercase letter with subscript j (e.g., p_j or x_{ij}). Lowercase letters also denote functions: e.g., f if the function is scalar valued, and \mathbf{f} if the function is vector valued. We denote the vector of ones of size n by $\mathbf{1}_n$, the set of integers $\{1, \dots, n\}$ by $[n]$, the set of natural numbers by \mathbb{N} , the set of real numbers by \mathbb{R} , and the positive and strictly positive elements of a set by a $+$ and $++$ subscript, respectively, e.g., \mathbb{R}_+ and \mathbb{R}_{++} . Finally, we denote the orthogonal projection operator onto a set C by Π_C , i.e., $\Pi_C(\mathbf{x}) = \arg \min_{\mathbf{y} \in C} \|\mathbf{x} - \mathbf{y}\|^2$.

Pseudo-games

A (concave) **pseudo-game** (Arrow and Debreu 1954) $\mathcal{G} \triangleq (n, \mathcal{A}, \mathcal{X}, \mathbf{g}, F)$ comprises $n \in \mathbb{N}_+$ players, each $i \in [n]$ of whom chooses an action $\mathbf{a}_i \in \mathcal{A}_i \subset \mathbb{R}^m$, with the players' joint action space $\mathcal{A} = \times_{i \in [n]} \mathcal{A}_i$. Each player i aims to maximize their continuous utility $f_i : \mathcal{A} \rightarrow \mathbb{R}$, which is concave in \mathbf{a}_i , by choosing a feasible action from a set of actions $\mathcal{X}_i(\mathbf{a}_{-i}) \subseteq \mathcal{A}_i$ determined by the actions $\mathbf{a}_{-i} \in \mathcal{A}_{-i} \subset \mathbb{R}^{(n-1)m}$ of the other players, where $\mathcal{X}_i : \mathcal{A}_{-i} \rightrightarrows \mathcal{A}_i$ is a non-empty, continuous, compact- and convex-valued action correspondence. We represent each such correspondence as a set $\mathcal{X}_i(\mathbf{a}_{-i}) = \{\mathbf{a}_i \in \mathcal{A}_i \mid g_{ik}(\mathbf{a}_i, \mathbf{a}_{-i}) \geq 0, \text{ for all } k \in [d]\}$, where g_{ik} is a continuous and concave function in \mathbf{a}_i that defines the constraints. Additionally, overloading notation, we define the joint action correspondence \mathcal{X} is simply a set of **jointly feasible action profiles**, namely $\{\mathbf{a} \in \mathcal{A} \mid \mathbf{g}(\mathbf{a}) \geq \mathbf{0}\}$, where $\mathbf{g} = (g_1, \dots, g_n)$. If $g_{ik} : \mathcal{A} \rightarrow \mathbb{R}$ is also concave in $\mathbf{a} \in \mathcal{A}$, then we say that the pseudo-game has **jointly convex constraints**, in which case, \mathcal{X} is simply a convex set. A pseudo-game is called **monotone**¹ if for all $\mathbf{x}, \mathbf{y} \in \mathcal{A}$, $\sum_{i \in [n]} (\nabla_{\mathbf{a}_i} f_i(\mathbf{x}) - \nabla_{\mathbf{a}_i} f_i(\mathbf{y}))^T (\mathbf{x}_i - \mathbf{y}_i) \leq 0$. Finally, a (concave) **game** (Nash 1950) is a pseudo-game where, for all players $i \in [n]$, \mathcal{X}_i is a constant correspondence, i.e., for all players $i \in [n]$, $\mathcal{X}_i(\mathbf{a}_{-i}) = \mathcal{X}_i(\mathbf{b}_{-i})$, for all $\mathbf{a}, \mathbf{b} \in \mathcal{A}$.

Given a pseudo-game \mathcal{G} , an ε -**generalized Nash equilibrium (GNE)** is an action profile $\mathbf{a}^* \in \mathcal{X}(\mathbf{a}^*)$ s.t. for all $i \in [n]$ and $\mathbf{a}_i \in \mathcal{X}_i(\mathbf{a}_{-i}^*)$, $f_i(\mathbf{a}^*) \geq f_i(\mathbf{a}_i, \mathbf{a}_{-i}^*) - \varepsilon$. An ε -**variational equilibrium (VE)** (or ε -**normalized GNE**) of a pseudo-game with joint constraints is an action profile $\mathbf{a}^* \in \mathcal{X}$ s.t. for all $i \in [n]$ and $\mathbf{a} \in \mathcal{X}$, $f_i(\mathbf{a}^*) \geq f_i(\mathbf{a}_i, \mathbf{a}_{-i}^*) - \varepsilon$. A GNE (VE) is an ε -GNE (VE) with $\varepsilon = 0$. While GNE

¹We call a pseudo-game monotone if $-(\nabla_{\mathbf{a}_1} u_1, \dots, \nabla_{\mathbf{a}_n} u_n)$ is a monotone operator. Such pseudo-games are also sometimes called dissipative, since $(\nabla_{\mathbf{a}_1} u_1, \dots, \nabla_{\mathbf{a}_n} u_n)$ is called a dissipative operator if $-(\nabla_{\mathbf{a}_1} u_1, \dots, \nabla_{\mathbf{a}_n} u_n)$ is a monotone operator.

are guaranteed to exist in all pseudo-games under standard assumptions (see Theorem 16), VE are only guaranteed to exist in pseudo-games with jointly convex constraints (see Theorem 17) (Arrow and Debreu 1954).

Fisher Markets with Social Influence

In this paper, we study a model of Fisher markets with social influence, in which a buyer's utility may be influenced by the goods allocated to her neighbors. A **Fisher market with social influence**, or an **influence Fisher market** for short, comprises $n \in \mathbb{N}_+$ buyers and $m \in \mathbb{N}_+$ divisible goods. Without loss of generality, we assume that exactly one unit of each good $j \in [m]$ is available.

The buyers are connected through a directed social influence graph $G = (V, E)$, where $V = [n]$ is the set of buyers, and for any $i, i' \in [n]$, there is an edge from i' to i iff the utility of i is influenced by the allocation $\mathbf{x}_{i'}$ of i' . We let $N_G(i) = \{i' \mid (i', i) \in E\}$ be the (incoming, and hence influential) neighbors of a buyer i , and we define $k_i = |N_G(i)|$, for all $i \in [n]$.

Each buyer $i \in [n]$ has a budget $b_i \in \mathbb{R}_+$ and a utility function $u_i : \mathbb{R}_+^{(k_i+1) \times m} \rightarrow \mathbb{R}$ that depends on not only her own allocation, but also her neighbors'. An instance of an influence Fisher market is thus given by a tuple (n, m, G, U, \mathbf{b}) , where G is the social network, $U = \{u_1, \dots, u_n\}$ is a set of utility functions, one per buyer, and $\mathbf{b} \in \mathbb{R}_+^n$ is a vector of buyer budgets. When n and m are clear from context, we denote influence Fisher markets simply by (G, U, \mathbf{b}) .

Given an influence Fisher market (G, U, \mathbf{b}) , an **allocation** $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T \in \mathbb{R}_+^{n \times m}$ is a map from goods to buyers, represented as a matrix, s.t. $x_{ij} \geq 0$ denotes the amount of good $j \in [m]$ allocated to buyer $i \in [n]$. Likewise, we denote by $\mathbf{x}_{N_G(i)} = (\mathbf{x}_{i'})_{i' \in N_G(i)} \in \mathbb{R}_+^{k_i \times m}$ the matrix representing the bundles of goods obtained by buyer i 's neighbors. A utility function is **locally non-satiated** if for all $\mathbf{x}_i \in \mathbb{R}_+^m$, $\mathbf{x}_{N_G(i)} \in \mathbb{R}_+^{k_i \times m}$, and $\varepsilon > 0$, there exists an $\mathbf{x}'_i \in \mathbb{R}_+^m$ with $\|\mathbf{x}'_i - \mathbf{x}_i\| \leq \varepsilon$ such that $u_i(\mathbf{x}'_i, \mathbf{x}_{N_G(i)}) > u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)})$. Related, a utility function satisfies **no saturation** if $\forall \mathbf{x}_i \in \mathbb{R}_+^m$ and $\mathbf{x}_{N_G(i)} \in \mathbb{R}_+^{k_i \times m}$, there exists an $\mathbf{x}'_i \in \mathbb{R}_+^m$ such that $u_i(\mathbf{x}'_i, \mathbf{x}_{N_G(i)}) > u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)})$. Note that if u_i is quasi-concave in \mathbf{x}_i and satisfies no saturation, then it is locally non-satiated (Arrow and Debreu 1954).

Feasibility asserts that no more of any good j is allocated than its available supply, i.e., $\forall j \in [m], \sum_{i \in [n]} x_{ij}^* \leq 1$. **Walras' law** states that, if a good j is not fully allocated, then its price must be zero; equivalently, if a good's price is positive, then it must be fully allocated. Mathematically, $\sum_{j \in [m]} p_j^* (\sum_{i \in [n]} x_{ij}^* - 1) = 0$.

A tuple $(\mathbf{X}^*, \mathbf{p}^*)$, which consists of an allocation \mathbf{X}^* and prices $\mathbf{p}^* = (p_1^* \dots, p_m^*)^T \in \mathbb{R}_+^m$, is a **competitive equilibrium (CE)** in an influence Fisher market (G, U, \mathbf{b}) if (1) fixing other buyers' allocations, buyers maximize their utilities constrained by their budget, i.e., $\forall i \in [n], \mathbf{x}_i^* \in \arg \max_{\mathbf{x}_i \in \mathbb{R}_+^m: \mathbf{x}_i \cdot \mathbf{p}^* \leq b_i} u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}^*)$, and (2) feasibility and Walras' law hold.

A Fisher market is a special influence Fisher market (G, U, \mathbf{b}) where $G = (V, E)$ satisfies $E = \emptyset$. In other words, each buyer i is isolated, so her utility $u_i : \mathbb{R}_+^m \rightarrow \mathbb{R}_+$ depends only on her own allocation. As G is simply a graph with n vertices and no edges, we can denote a Fisher market by the tuple (U, \mathbf{b}) .

When U is a set of specific utility functions, we refer to the influence Fisher market (G, U, \mathbf{b}) by the name of the utility function: e.g., if U is a set of linear utility functions then (G, U, \mathbf{b}) is a linear influence Fisher market.

Existence of Competitive Equilibrium via Pseudo-Games

In this section, we investigate the properties of competitive equilibrium in Fisher markets with social influence. Our main tool is the pseudo-game (or abstract economy) model introduced by Arrow and Debreu, as both a generalization of the standard normal-form game in game theory and of the Arrow-Debreu market in microeconomics (Arrow and Debreu 1954). We provide a proof of existence of competitive equilibrium in influence Fisher markets, using methods similar to those employed by Arrow and Debreu in their seminal proof of the existence of competitive equilibria in Arrow-Debreu economies.

Following Arrow and Debreu, we define a pseudo-game with an auctioneer who sets prices. Our pseudo-game then both generalizes and specializes theirs. While in theirs, each trader's utility depends only on their own allocations, ours captures social influence through augmented utility functions. While in theirs, buyers are constrained by their endowment, in ours, buyers are constrained by budgets.

Specifically, we construct an auctioneer-buyer pseudo-game comprising a single auctioneer and n individual buyers in which the auctioneer sets the good prices, while the buyers choose their allocations. Given an allocation $\mathbf{X} \in \mathbb{R}^{n \times m}$, let $\mathbf{z} = \left(\sum_{i \in [n]} \mathbf{x}_i \right) - \mathbf{1}_m$ be the vector of **excess demands**, i.e., the total amount by which the demand for each good exceeds its supply. In our auctioneer-buyer pseudo-game, each buyer i chooses allocations \mathbf{x}_i that maximize her utility subject to her budget constraint, given prices \mathbf{p} determined by the auctioneer, the auctioneer chooses prices that maximize her total profit, i.e., $\mathbf{p} \cdot \mathbf{z}$, fixing the allocation \mathbf{X} , subject to Walras' law. More specifically, we assume a numeraire (i.e., a good whose price we normalize to 1), and we view the buyers' budgets as quantities of this numeraire, in which case Walras' law can be restated as the sum of the prices being equal to the sum of the budgets: i.e., $\sum_{j \in [m]} p_j = \sum_{i \in [n]} b_i$.

In what follows, we show that the set of GNE in this auctioneer-buyer pseudo-game corresponds to the set of CE in an influence Fisher market; existence of a CE thus follows from existence of GNE.

Importantly, in this pseudo-game we require the auctioneer to choose prices whose sum is equal to the sum of the budgets, because the budgets correspond to a numeraire good, i.e., their price is set to 1 (this is without loss of generality since in any Arrow-Debreu market only the propor-

tions of the prices to one another matters), hence the prices of every other good have to be scaled appropriately, i.e., by the sum of the budget, so as to preserve the internal consistency of the Arrow-Debreu price system that emerges at the competitive equilibrium prices (since Arrow-Debreu assume wlog that prices are in the unit simplex).

Assumption 1. An influence Fisher market (G, U, \mathbf{b}) satisfies, for all buyers $i \in [n]$, u_i is 1. continuous in $(\mathbf{x}_i, \mathbf{x}_{N_G(i)})$, 2. concave in \mathbf{x}_i , and 3. satisfies no saturation.

Definition 2 (Auctioneer-Buyer Pseudo-game). Let (G, U, \mathbf{b}) be an influence Fisher market. The corresponding auctioneer-buyer pseudo-game $\mathcal{G} = (n + 1, \mathcal{A}, \mathcal{X}, \mathbf{g}, F)$ is defined by

- an auctioneer and n buyers.
- Each buyer chooses an allocation $\mathbf{x}_i \in \mathcal{A}_i = \mathbb{R}_+^m$, while the auctioneer chooses prices $\mathbf{p} \in \mathcal{A}_{n+1} = \mathbb{R}_+^m$.
- For all buyers $i \in [n]$, the feasible action set given the actions of other players is $\mathcal{X}_i(\mathbf{x}_{-i}, \mathbf{p}) = \{\mathbf{x}_i \in \mathcal{A}_i \mid \mathbf{g}(\mathbf{x}_i, \mathbf{x}_{-i}, \mathbf{p}) = b_i - \mathbf{x}_i \cdot \mathbf{p} \geq 0\}$.
- For the auctioneer, the feasible action set is the fixed set $\mathcal{X}_{n+1} = \{\mathbf{p} \in \mathcal{A}_{n+1} \mid \mathbf{p} \cdot \mathbf{1}_m = \mathbf{b} \cdot \mathbf{1}_n\}$.
- For all players $i \in [n + 1]$, i maximizes her utility $f_i : \prod_{i \in [n+1]} \mathcal{A}_i \rightarrow \mathbb{R}$, defined by $f_i(\mathbf{X}, \mathbf{p}) = u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)})$, for the buyers $i \in [n]$, and $f_{n+1}(\mathbf{X}, \mathbf{p}) = \mathbf{p} \cdot \mathbf{z}$ for the auctioneer.

Next, we prove the existence of CE in influence Fisher markets that satisfy Assumption 1.²

Theorem 3. The set of competitive equilibria of any influence Fisher market (G, U, \mathbf{b}) that satisfies Assumption 1 is equal to the set of generalized Nash equilibria of the associated auctioneer-buyer pseudo-game $\mathcal{G} = (n + 1, \mathcal{A}, \mathcal{X}, \mathbf{g}, F)$.

Existence of a CE in an influence Fisher market now follows immediately from existence of GNE in pseudo-games:

Corollary 4. There exists a CE $(\mathbf{X}^*, \mathbf{p}^*)$ in all influence Fisher markets (G, U, \mathbf{b}) satisfying Assumption 1.

Computation of Competitive Equilibrium via Pseudo-Games

Although we have established the existence of competitive equilibrium in all influence Fisher markets with continuous and concave utility functions, the proof itself provides little insight into equilibrium computation, as computing a GNE is PPAD-hard in general (Daskalakis, Goldberg, and Papadimitriou 2009). In order to gain further computational insights, we focus on a subset of influence Fisher markets in which each buyer's utility function is also homogeneous in its own allocation.

A utility function u_i is **homogeneous** in $\mathbf{x}_i \in \mathbb{R}_+^m$ if $u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)})$ is homogeneous for all $\mathbf{x}_{N_G(i)} \in \mathbb{R}_+^{k_i \times m}$, i.e., $u_i(\lambda \mathbf{x}_i, \mathbf{x}_{N_G(i)}) = \lambda u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)})$, for all $\lambda \geq 0$. As above, we also assume continuity and concavity. We call

²Proofs of all theorems appear in the appendix.

utility functions that satisfy all three of these assumptions CCH utility functions, and Fisher markets inhabited by buyers with such utility functions CCH Fisher markets.

We can compute competitive equilibria in CCH Fisher markets (without social influence) via the Eisenberg-Gale convex program and its dual (Eisenberg and Gale 1959). To generalize this convex program to CCH influence Fisher markets, we propose another pseudo-game, which we call the buyer (only) pseudo-game, that is jointly convex, and whose variational equilibria correspond to CE allocations. Moreover, we observe that the “dual” of this pseudo-game simultaneously characterizes CE prices. In other words, while the auctioneer-buyer pseudo-game explicitly models an auctioneer who updates prices in response to the buyers’ behavior, in the buyer (only) pseudo-game, the auctioneer is “fictitious,” as it is implicit in the dual.

If (U, \mathbf{b}) is a CCH Fisher market, then an optimal solution \mathbf{X}^* to the Eisenberg-Gale program (Eq. 1) constitutes a CE allocation, and an optimal solution to the Lagrangian that represents the feasibility constraints (Eq. 1a) are the corresponding equilibrium prices (Devanur et al. 2002; Jain, Vazirani, and Ye 2005).

Primal:

$$\begin{aligned} \max_{\mathbf{X} \in \mathbb{R}_+^{n \times m}} \quad & \sum_{i \in [n]} b_i \log(u_i(\mathbf{x}_i)) \quad (1) \\ \text{subject to} \quad & \forall j \in [m], \sum_{i \in [n]} x_{ij} \leq 1 \quad (1a) \end{aligned}$$

In Fisher markets (without social influence), each buyer’s utility maximization problem is independent of the others’, as each depends only on the buyer’s own allocation. The Eisenberg-Gale program takes advantage of this independence. It takes an aggregate perspective, maximizing the *sum* of the buyers’ utilities subject to their feasibility constraints, and nonetheless computes an optimal allocation that maximizes each buyer’s *individual* utility.

In influence Fisher markets, however, where this independence assumption does not hold, we can no longer compute CE from this aggregate perspective. In our solution—CE as the VE of a jointly-convex pseudo-game—each buyer maximizes her own utility, subject to a shared feasibility constraint. Note that the only players in this buyer pseudo-game are the n buyers; there is no auctioneer updating prices based on the buyers’ behavior.

Definition 5 (Buyer Pseudo-game). *Let (G, U, \mathbf{b}) be an influence Fisher market. The corresponding jointly-convex buyer pseudo-game $\mathcal{G} = (n, \mathcal{A}, \mathcal{X}, \mathbf{g}, F)$ is defined by*

- For all buyers $i \in [n]$, $\mathcal{A}_i = \mathbb{R}_+^m$.
- For all buyers $i \in [n]$, the feasible action set given the actions of other players is $\mathcal{X}_i(\mathbf{x}_{-i}) = \{\mathbf{x}_i \in \mathcal{A}_i \mid \mathbf{g}(\mathbf{x}_i, \mathbf{x}_{-i}) = \mathbf{1} - \sum_{i \in [n]} \mathbf{x}_i \geq \mathbf{0}\}$.
- For all buyers $i \in [n]$, i maximizes her utility $f_i : \times_{i \in [n]} \mathcal{A}_i \rightarrow \mathbb{R}$ defined by $f_i(\mathbf{X}) = b_i \log(u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}))$.

Assumption 6. *For each buyer $i \in [n]$, u_i is 1. continuous in $(\mathbf{x}_i, \mathbf{x}_{N_G(i)})$; and 2. concave and homogeneous³ in \mathbf{x}_i .*

Theorem 7. *Let (G, U, \mathbf{b}) be an influence Fisher market satisfying Assumption 6. Then, \mathbf{X}^* is a CE allocation of (G, U, \mathbf{b}) if and only if it is a variational equilibrium (VE) of the corresponding buyer pseudo-game $\mathcal{G} = (n, \mathcal{A}, \mathcal{X}, \mathbf{g}, F)$. Moreover, if \mathbf{X}^* is a VE of the buyer pseudo-game, then the corresponding KKT conditions are satisfied with optimal Lagrange multipliers $\lambda_1^* = \dots = \lambda_n^* = \mathbf{p}^*$, which correspond to CE prices.*

The construction of competitive equilibrium via the auctioneer-buyer pseudo-game (Theorem 3) is more general than the construction of competitive equilibrium via the buyer pseudo-game (Theorem 7); however, the existence of the auctioneer precludes monotonicity, and hence polynomial-time computability. To obtain efficient algorithms, we assume the buyers’ utilities are concave not only in their own allocations but in one another’s allocations as well, which implies monotonicity. We also require twice-continuous differentiability.

Assumption 8. *For each buyer $i \in [n]$, 1. The conditions in Assumption 6, and 2. u_i is jointly concave: i.e., concave in $(\mathbf{x}_i, \mathbf{x}_{N_G(i)})$, and 3. and twice-continuously differentiable in $(\mathbf{x}_i, \mathbf{x}_{N_G(i)})$.*

Under Assumption 8, an influence Fisher market can be expressed as a monotone variational inequality. There exist methods that converge in last iterates⁴ to a solution of any monotone variational inequality at a rate of $O(1/T)$ (e.g., Gorbunov, Loizou, and Gidel (2022)). Our next theorem follows from these two assertions:

Theorem 9. *There exist methods that converge in last iterates to the CE allocations of influence Fisher markets at a rate of $O(1/T)$ under Assumption 8. In such markets, approximate competitive equilibrium allocations can be computed in polynomial time.*

Computation of Competitive Equilibrium via Stackelberg Games

Recently, Cole and Tao (2019) presented a generalization of the Eisenberg-Gale dual for arbitrary CCH utility functions, which accurately characterizes competitive equilibrium prices, but fails to match the optimal objective value of the Eisenberg-Gale primal. Building on their results, Goktas, Viqueira, and Greenwald (2021) derived the exact Eisenberg-Gale dual, for which strong duality holds.

³We note that homogeneity implies no saturation, since for all $\mathbf{x} \in \mathbb{R}_+^m$ and $\mathbf{x}_{N_G(i)} \in \mathbb{R}_+^{k_i \times m}$, there exists an allocation $(1 + \varepsilon)\mathbf{x}$ for some $\varepsilon > 0$ s.t. $u_i((1 + \varepsilon)\mathbf{x}, \mathbf{x}_{N_G(i)}) = (1 + \varepsilon)u_i(\mathbf{x}, \mathbf{x}_{N_G(i)}) > u_i(\mathbf{x}, \mathbf{x}_{N_G(i)})$.

⁴Solodov and Svaiter (1999) and Ryu, Yuan, and Yin (2019) also provide methods that guarantee average-iterate convergence with this same rate in monotone variational inequalities.

Dual:

$$\begin{aligned} \min_{\mathbf{p} \in \mathbb{R}_+^m} \quad & \sum_{j \in [m]} p_j + \sum_{i \in [n]} b_i \log(u_i(\mathbf{x}_i^*)) - b_i \quad (2) \\ \text{s.t.} \quad & \forall i \in [n], \mathbf{x}_i^* \in \arg \max_{\mathbf{x}_i \in \mathbb{R}_+^m: \mathbf{x}_i \cdot \mathbf{p} \leq b_i} u_i(\mathbf{x}_i) \quad (2a) \end{aligned}$$

We begin this section by deriving the ‘‘duals’’ of our buyer pseudo-game.

In the buyer pseudo-game \mathcal{G} , each buyer is solving an optimization problem (Equation (3)) in which they maximize their utility function by choosing an optimal action in their feasible action set, given the other buyers’ VE actions. Based on this observation, we can derive the ‘‘dual’’ of our buyer pseudo-game; but as our pseudo-game comprises n different optimization problems, one for each buyer $i \in [n]$, instead of just one dual, we have n duals. Moreover, because any VE of a jointly-convex pseudo-game satisfies the corresponding KKT conditions with optimal Lagrange multipliers $\lambda_1^* = \dots = \lambda_n^*$ (Theorem 19 (Facchinei, Fischer, and Piccialli 2009)), all n duals yield the same CE prices! In other words, just as the dual of Eisenberg-Gale program characterizes the CE prices of a Fisher market, the n duals of our pseudo-game characterize the CE prices of an influence Fisher market (satisfying Assumption 6).

Theorem 10. *Let (G, U, \mathbf{b}) be an influence Fisher market satisfying Assumption 6, and let \mathcal{G} be the corresponding buyer pseudo-game $\mathcal{G} = (n, \mathcal{A}, \mathcal{X}, \mathbf{g}, F)$. For each buyer $i \in [n]$, fixing its neighbors’ allocations $\mathbf{x}_{N_G(i)}^*$, the dual of i ’s optimization problem,*

$$\max_{\mathbf{x}_i \in \mathbb{R}_+^m: \mathbf{x}_i + \sum_{k \neq i} \mathbf{x}_k^* \leq \mathbf{1}} b_i \log(u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}^*)) \quad (3)$$

is given by

$$\min_{\mathbf{p} \in \mathbb{R}_+^m} \sum_{j \in [m]} p_j \left(1 - \sum_{k \neq i} x_{kj}^* \right) + b_i \log(u_i(\mathbf{x}_i^*, \mathbf{x}_{N_G(i)}^*)) - b_i \quad (4)$$

$$\text{s.t.} \quad \mathbf{x}_i^* \in \arg \max_{\mathbf{x}_i \in \mathbb{R}_+^m: \mathbf{x}_i \cdot \mathbf{p} \leq b_i} u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}^*) \quad (4a)$$

Goktas and Greenwald (2021) further show that the dual of the Eisenberg-Gale program can be re-expressed as the solution to the following zero-sum convex-concave Stackelberg game characterizes the CE of any CCH Fisher market:

$$\min_{\mathbf{p} \in \mathbb{R}_+^m} \max_{\mathbf{X} \in \mathbb{R}_+^{m \times m}: \mathbf{X} \cdot \mathbf{p} \leq \mathbf{b}} \sum_{j \in [m]} p_j + \sum_{i \in [n]} b_i \log(u_i(\mathbf{x}_i)) \quad (5)$$

The leader in this game is a fictitious auctioneer (i.e., price setter), while the follower represents a set of buyers who effectively play as a team. The objective function is the sum of the auctioneer’s welfare (i.e., the sum of the prices) and the buyers’ Nash social welfare. Goktas and Greenwald also derive a first-order method that solves this game, which, via the aforementioned interpretation, can be understood as computing a competitive equilibrium of a Fisher market via *tatonnement*.

We argue that competitive equilibria in *influence* Fisher markets can likewise be characterized via Stackelberg equilibria. This more general setting requires not just one, but a system of n zero-sum convex-concave Stackelberg games (Goktas and Greenwald 2021), one per buyer. In each game, the leader once again is a fictitious auctioneer (i.e., price setter), but the follower is just an individual buyer, not the set of all buyers. Moreover, in each buyer’s Stackelberg game, the objective function is the sum of the auctioneer’s revenue (i.e., the sum of the good prices, each one discounted by the supply available to buyer i beyond what has been claimed by the others) and the individual buyer’s utility. These n Stackelberg games are played simultaneously with the fictitious auctioneer optimizing prices assuming all the buyers simultaneously best respond (i.e., play a Nash equilibrium), and the buyers best respond to the auctioneer’s prices, given the other buyers’ allocations.

Definition 11 (Buyer i ’s Stackelberg Game). *Let (G, U, \mathbf{b}) be an influence Fisher market. The corresponding Stackelberg game for buyer i is defined by*

$$\min_{\mathbf{p} \in \mathbb{R}_+^m} \max_{\mathbf{x}_i \in \mathbb{R}_+^m: \mathbf{x}_i \cdot \mathbf{p} \leq b_i} \sum_{j \in [m]} p_j \left(1 - \sum_{k \neq i} x_{kj}^* \right) + b_i \log(u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}^*)) \quad (6)$$

The following corollary follows from Theorems 7 and 10.

Corollary 12. *$(\mathbf{X}^*, \mathbf{p}^*)$ is a competitive equilibrium of an influence CCH Fisher market (G, U, \mathbf{b}) satisfying Assumption 6 iff $(\mathbf{x}_i^*, \mathbf{p}^*)$ is a Stackelberg equilibrium in buyer i ’s Stackelberg game, for all buyers $i \in [n]$.*

Proof. For all $i \in [n]$, $(\mathbf{x}_i^*, \mathbf{p}^*)$ is a Stackelberg equilibrium in buyer i ’s Stackelberg game iff $(\mathbf{x}_i^*, \mathbf{p}^*)$ solves

$$\begin{aligned} \min_{\mathbf{p} \in \mathbb{R}_+^m} \quad & \sum_{j \in [m]} p_j \left(1 - \sum_{k \neq i} x_{kj}^* \right) + b_i \log(u_i(\mathbf{x}_i^*, \mathbf{x}_{N_G(i)}^*)) - b_i \\ \text{s.t.} \quad & \mathbf{x}_i^* \in \arg \max_{\mathbf{x}_i \in \mathbb{R}_+^m: \mathbf{x}_i \cdot \mathbf{p} \leq b_i} u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}^*) \quad (7) \end{aligned}$$

By Theorem 10, \mathbf{p}^* is a solution to this bi-level optimization problem (Equation (4)) iff \mathbf{x}_i^* is a solution to buyer i ’s optimization problem (Equation (3)) in the buyer pseudo-game corresponding to (G, U, \mathbf{b}) . Finally, by Theorem 7, $(\mathbf{X}^*, \mathbf{p}^*)$ is a competitive equilibrium of (G, U, \mathbf{b}) . \square

Our Stackelberg game formulation of CE in influence Fisher markets enables us to compute CE by solving a system of buyer Stackelberg games: i.e., solving for a Stackelberg equilibrium in each of the buyer Stackelberg games together with a Nash equilibrium among the buyers in the system. Towards that end, for convenience, we define the **objective function** for buyer i ’s Stackelberg game:

$$f_i(\mathbf{x}_i, \mathbf{p}) := \sum_{j \in [m]} p_j \left(1 - \sum_{k \neq i} x_{kj}^* \right) + b_i \log(u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}^*)) \quad (8)$$

and the i th (fictional) auctioneer's **value function** in buyer i 's Stackelberg game:

$$V_i(\mathbf{p}) := \max_{\mathbf{x}_i \in \mathbb{R}_+^m : \mathbf{x}_i \cdot \mathbf{p} \leq b_i} \sum_{j \in [m]} p_j \left(1 - \sum_{k \neq i} x_{kj}^* \right) + b_i \log(u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}^*)) \quad (9)$$

Moreover, while each buyer is playing a Stackelberg game with its fictitious auctioneer, all buyers are also playing an n -buyer simultaneous game with one another, in which each buyer maximizes its objective function $f_i(\mathbf{x}_i, \mathbf{p})$ (Equation (8)), given the prices \mathbf{p} set by the auctioneer and the other buyers' allocations. We can characterize a Nash equilibrium of this n -buyer game as follows:

$$\mathbf{x}_i^* \in \arg \max_{\mathbf{x}_i \in \mathbb{R}_+^m : \mathbf{x}_i \cdot \mathbf{p} \leq b_i} \sum_{j \in [m]} p_j \left(1 - \sum_{k \neq i} x_{kj}^* \right) + b_i \log(u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}^*)) \quad (10)$$

$$= \arg \max_{\mathbf{x}_i \in \mathbb{R}_+^m : \mathbf{x}_i \cdot \mathbf{p} \leq b_i} u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}^*) \quad (10a)$$

As the first summand in Equation 10 and b_i are constants (i.e., they do not depend on \mathbf{x}_i), and \log is a monotonic function, buyer i simply seeks to maximize its utility subject to its budget constraint.

Using a subdifferential envelope theorem (Goktas and Greenwald 2021), we now derive the subgradient of each auctioneer's value function V_i (Equation (9)).

Theorem 13. *Given an influence CCH Fisher market (G, U, \mathbf{b}) , the subdifferential of the i th auctioneer's value function in buyer i 's Stackelberg game (Equation (9)) at given prices \mathbf{p} is equal to the negative excess demand at \mathbf{p} : i.e., $\partial_{\mathbf{p}} V_i(\mathbf{p}) = \mathbf{1} - \sum_{i \in [n]} \mathbf{x}_i^*$.*

Interestingly, this subgradient turns out to equal the negative excess demand in the market at the given prices. As excess demand is an aggregate quantity, it is independent of buyer i . Indeed, the subgradients of *all* the fictional auctioneers are the same; so there is effectively just one auctioneer.

Based on this observation, we now present our **Nash Equilibrium (NE)-oracle gradient descent** algorithm (Algorithm 1), which follows the subgradient of the auctioneer's value function, assuming access to a NE-oracle. Given prices $\mathbf{p} \in \mathbb{R}_+^m$, this oracle returns a Nash equilibrium \mathbf{X}^* of the n -buyer concave game specified by Equation (10). The algorithm then runs subgradient descent on the auctioneer's value function. Overall, this approach corresponds to solving for a CE allocation and prices via *tâtonnement*, assuming the NE oracle is exact. As NE-oracles are rarely exact, Algorithm 1 assumes a NE-oracle that finds a Nash equilibrium up to some approximation error δ .

Finally, under standard assumptions (i.e., Assumption 8), the auctioneer's value function (Equation (9)) is convex and ℓ_V -Lipschitz continuous in \mathbf{p} with $\ell_V = \max_{\mathbf{p} \in \mathbb{R}_+^m} \|\nabla_{\mathbf{p}} V(\mathbf{p})\|$.⁵ These properties imply that our

⁵Although $\nabla_{\mathbf{p}} V(\mathbf{p})$ is not necessarily bounded at $\mathbf{p} = 0$, we can remedy this fact by shifting \mathbf{p} by a small constant $\varepsilon > 0$, albeit losing some accuracy.

Algorithm 1: NE-Oracle *Tâtonnement* For Influence Fisher Markets

Inputs: $G, U, \mathbf{b}, \mathbf{p}^{(0)}, \eta, \delta$

Outputs: $\mathbf{X}^*, \mathbf{p}^*$

```

1: for  $t = 1, \dots, T$  do
2:   Find  $\mathbf{X}' \in \mathbb{R}_+^{n \times m}$  with  $\mathbf{X}' \cdot \mathbf{p}^{(t-1)} \leq \mathbf{b}$  such that:
3:   for all  $i \in [n]$ ,  $u_i(\mathbf{x}'_i, \mathbf{x}'_{N_G(i)}) \geq u_i(\mathbf{x}_i, \mathbf{x}'_{N_G(i)}) - \delta$ ,
4:
5:   for any  $\mathbf{x}_i \in \mathbb{R}_+^m$  satisfying  $\mathbf{x}_i \cdot \mathbf{p}^{(t-1)} \leq b_i$ 
6:   Set  $\mathbf{X}^{(t)} = \mathbf{X}'$ 
7:   Set  $\mathbf{p}^{(t)} = \Pi_{\mathbb{R}_+^m} \left( \mathbf{p}^{(t-1)} - \eta(1 - \sum_{i \in [n]} \mathbf{x}_i^{(t)}) \right)$ 
8: end for
9: return  $\mathbf{X}^{(T)}, \mathbf{p}^{(T)}$ 

```

NE-oracle gradient descent algorithm converges to competitive equilibrium at a rate of $O(1/\sqrt{T})$. We include a more detailed statement of the following theorem in the appendix.

Theorem 14. *Algorithm 1 (i.e., tâtonnement) converges to a competitive equilibrium in any influence CCH Fisher market (G, U, \mathbf{b}) satisfying Assumption 8 at a rate of $O(1/\sqrt{T})$.*

Remark 15. *We can implement an approximate NE oracle by computing the buyers' equilibrium allocations via extragradient ascent (Gorbunov, Loizou, and Gidel 2022), which is guaranteed to converge to a Nash equilibrium at a rate of $O(1/T)$, as the n -buyer concave game defined by Equation (10) is monotone. This observation gives rise to Algorithm 2 (see Appendix), which computes a competitive equilibrium in influence CCH Fisher markets in polynomial time.*

Experiments

We ran a series of experiments⁶ to see how the empirical convergence rates of Algorithm 1 compare to its theoretical guarantees under various utility structures. We considered three standard utility functions: linear, in which buyers practice utilitarian social welfare in their neighborhoods; Cobb-Douglas, in which practice Nash social welfare in their neighborhoods; and Leontief, in which practice egalitarian social welfare in their neighborhoods. Each utility structure endows the objective function (Equation (8)) and the value functions (Equation (9)) with different smoothness properties, which in turn varies the convergence properties of our algorithms.

Let $\theta_i \in \mathbb{R}^m$ be a vector of parameters that describes the utility function $f_i : \mathbb{R}_+^m \rightarrow \mathbb{R}_+$ of buyer $i \in [n]$. We consider the following (standard) utility functions: for all $i \in [n]$,

1. Linear: $u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}) = \sum_{k \in i \cup N_G(i)} f_k(\mathbf{x}_k)$, where $f_i(\mathbf{x}_i) = \sum_{j \in [m]} \theta_{ij} x_{ij}$
2. Cobb-Douglas: $u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}) = \prod_{k \in \{i\} \cup N_G(i)} f_k(\mathbf{x}_k)$, where $f_i(\mathbf{x}_i) = \prod_{j \in [m]} x_{ij}^{\theta_{ij}}$
3. Leontief: $u_i(\mathbf{x}_i, \mathbf{x}_{N_G(i)}) = \min_{k \in \{i\} \cup N_G(i)} f_k(\mathbf{x}_k)$, where $f_i(\mathbf{x}_i) = \min_{j \in [m]} \left\{ \frac{x_{ij}}{\theta_{ij}} \right\}$

⁶We include a detailed description of our experimental setup in the Appendix.




Figure 1: In blue, we depict a trajectory of average value of the objective function across experiments (Equation (8)), for Algorithm 1 with EG as the NE-oracle, in randomly initialized linear, Cobb-Douglas, and Leontief Fisher markets with social influence. In red, we plot an arbitrary $O(1/\sqrt{T})$ function.

Assuming any of these three utility functions, we can solve for a Nash equilibrium among buyers by formulating a monotone variational inequality problem, and solving it via the extragradient method (EG) in $O(1/T)$ iterations (Gorbunov, Loizou, and Gidel 2022). Then, by using EG as the NE-oracle, we can efficiently compute an optimal $\mathbf{X}^*(\mathbf{p})$ for any given \mathbf{p} , which yields Algorithm 2 (see Appendix), a specific implementation of Algorithm 1.

Figure 1 depicts the empirical convergence of Algorithm 1 with EG as the NE-oracle. We observe that convergence is fastest in influence Fisher markets with Cobb-Douglas utilities, followed by linear, and then Leontief. For influence Fisher markets with Cobb-Douglas utilities, both the value and the objective function are differentiable; in fact, they are both twice continuously differentiable, making them both Lipschitz-smooth. These factors combined seem to lead to a faster convergence rate than $O(1/\sqrt{T})$. On the other hand, for influence Fisher markets with linear utilities, we seem to obtain a tight convergence rate of $O(1/\sqrt{T})$, which seems plausible, as the value function is not differentiable assuming linear utilities, and hence we are unlikely to achieve a better convergence rate. Finally, influence Fisher markets with Leontief utilities, in which the objective function is not differentiable, are the hardest markets for our algorithm to solve. Nonetheless, we still observe a decent convergence rate, one that appears only slightly slower than $O(1/\sqrt{T})$.

Conclusion

In this paper, we studied a special case of Arrow-Debreu markets with social influence, which we call Fisher markets with social influence, or influence Fisher markets for short. First, we extended known results on the existence of competitive equilibrium in markets with social influence to a larger more natural class of markets. Our proof proceeds by reducing an influence Fisher market to an auctioneer-buyer pseudo-game such that every generalized Nash equilibrium in the pseudo-game is a competitive equilibrium of the influence Fisher market. The existence of generalized Nash

equilibrium in pseudo-games thus implies the existence of competitive equilibrium in influence Fisher markets.

We then introduced a monotone jointly convex buyer-only pseudo-game as a generalization of the Eisenberg-Gale program, whose variational equilibria correspond to the competitive equilibria in influence Fisher markets. In this pseudo-game, the duals of the individual buyers' utility-maximization problems constrained by the supply constraint comprise a system of n simultaneously-played zero-sum Stackelberg games, which simultaneously characterize the competitive equilibrium prices of the influence Fisher market. We then show that running gradient descent on the leaders'/auctioneers' value functions in these games is equivalent to solving the market via a variant of *tâtonnement*, where in addition to the auctioneers iteratively adjusting prices, the buyers iteratively learn a Nash equilibrium in response to these prices.

Our results pave the way for future work developing methods to compute competitive equilibria in more general types of influence markets beyond those considered in this paper (Chen and Teng 2011), and other market models with graphical structure, such as graphical economies (Kakade, Kearns, and Ortiz 2004).

Acknowledgments

This work was supported by NSF Grant CMMI-1761546.

References

- Arrow, K. J.; and Debreu, G. 1954. Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, 265–290.
- Brainard, W. C.; Scarf, H. E.; et al. 2000. *How to compute equilibrium prices in 1891*. Citeseer.
- Chen, X.; and Deng, X. 2006. Settling the complexity of two-player Nash equilibrium. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, 261–272. IEEE.

- Chen, X.; and Teng, S.-H. 2009. Spending is not easier than trading: on the computational equivalence of Fisher and Arrow-Debreu equilibria. In *International Symposium on Algorithms and Computation*, 647–656. Springer.
- Chen, X.; and Teng, S.-H. 2011. A Complexity View of Markets with Social Influence. *ArXiv*, abs/1009.0309.
- Cheung, Y. K.; Cole, R.; and Devanur, N. 2013. Tatonnement beyond Gross Substitutes? Gradient Descent to the Rescue. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '13, 191–200. New York, NY, USA: Association for Computing Machinery. ISBN 9781450320290.
- Cole, R.; and Tao, Y. 2019. Balancing the Robustness and Convergence of Tatonnement. *arXiv*:1908.00844.
- Daskalakis, C.; Goldberg, P. W.; and Papadimitriou, C. H. 2009. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1): 195–259.
- Devanur, N. R.; Papadimitriou, C. H.; Saberi, A.; and Vazirani, V. V. 2002. Market equilibrium via a primal-dual-type algorithm. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, 389–395.
- Devanur, N. R.; Papadimitriou, C. H.; Saberi, A.; and Vazirani, V. V. 2008. Market equilibrium via a primal-dual algorithm for a convex program. *Journal of the ACM (JACM)*, 55(5): 1–18.
- Dreves, A. 2017. Computing all solutions of linear generalized Nash equilibrium problems. *Mathematical Methods of Operations Research*, 85(2): 207–221.
- Dreves, A.; Heusinger, A.; Kanzow, C.; and Fukushima, M. 2013. A Globalized Newton Method for the Computation of Normalized Nash Equilibria. *J. of Global Optimization*, 56(2): 327–340.
- Eisenberg, E.; and Gale, D. 1959. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1): 165–168.
- Facchinei, F.; Fischer, A.; and Piccialli, V. 2007. On generalized Nash games and variational inequalities. *Operations Research Letters*, 35(2): 159–164.
- Facchinei, F.; Fischer, A.; and Piccialli, V. 2009. Generalized Nash equilibrium problems and Newton methods. *Mathematical Programming*, 117(1): 163–194.
- Fischer, A.; Herrich, M.; Izmailov, A. F.; and Solodov, M. V. 2016. A globally convergent LP-Newton method. *SIAM Journal on Optimization*, 26(4): 2012–2033.
- Gao, Y.; and Kroer, C. 2020. First-Order Methods for Large-Scale Market Equilibrium Computation. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Goktas, D.; and Greenwald, A. 2021. Convex-Concave Min-Max Stackelberg Games. *Advances in Neural Information Processing Systems*, 34.
- Goktas, D.; and Greenwald, A. 2022. Exploitability Minimization in Games and Beyond. *arXiv preprint arXiv:2210.10207*.
- Goktas, D.; Viqueira, E. A.; and Greenwald, A. 2021. A Consumer-Theoretic Characterization of Fisher Market Equilibria. In *Web and Internet Economics: 17th International Conference, WINE 2021, Potsdam, Germany, December 14–17, 2021, Proceedings*, 334–351. Berlin, Heidelberg: Springer-Verlag. ISBN 978-3-030-94675-3.
- Goktas, D.; Zhao, J.; and Greenwald, A. 2022. Robust no-regret learning in min-max Stackelberg games. *arXiv preprint arXiv:2203.14126*.
- Gorbunov, E.; Loizou, N.; and Gidel, G. 2022. Extragradient method: O(1/K) last-iterate convergence for monotone variational inequalities and connections with cocoercivity. In *International Conference on Artificial Intelligence and Statistics*, 366–402. PMLR.
- Heusinger, A.; and Kanzow, C. 2009. Relaxation Methods for Generalized Nash Equilibrium Problems with Inexact Line Search. *Journal of Optimization Theory and Applications*, 143(1): 159–183.
- Izmailov, A. F.; and Solodov, M. V. 2014. On error bounds and Newton-type methods for generalized Nash equilibrium problems. *Computational Optimization and Applications*, 59(1): 201–218.
- Jain, K.; Vazirani, V. V.; and Ye, Y. 2005. Market equilibria for homothetic, quasi-concave utilities and economies of scale in production. In *SODA*, volume 5, 63–71.
- Jordan, M. I.; Lin, T.; and Zampetakis, M. 2023. First-Order Algorithms for Nonlinear Generalized Nash Equilibrium Problems. *Journal of Machine Learning Research*, 24(38): 1–46.
- Kakade, S. M.; Kearns, M.; and Ortiz, L. E. 2004. Graphical Economics. In *Annual Conference Computational Learning Theory*.
- Krawczyk, J. B.; and Uryasev, S. 2000. Relaxation algorithms to find Nash equilibria with economic applications. *Environmental Modeling & Assessment*, 5(1): 63–73. This revised version was published online in July 2006 with corrections to the Cover Date.
- Mas-Colell, A.; Whinston, M. D.; and Green, J. R. 1995. *Microeconomic Theory*. Number 9780195102680 in OUP Catalogue. Oxford University Press. ISBN AR-RAY(0x4cf9c5c0).
- Nabetani, K.; Tseng, P.; and Fukushima, M. 2011. Parametrized Variational Inequality Approaches to Generalized Nash Equilibrium Problems with Shared Constraints. *Comput. Optim. Appl.*, 48(3): 423–452.
- Nash, J. F. 1950. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences*, 36(1): 48–49.
- Rosen, J. B. 1965. Existence and Uniqueness of Equilibrium Points for Concave N -Person Games. *Econometrica*, 33(3): 520–534.
- Ryu, E. K.; Yuan, K.; and Yin, W. 2019. ODE Analysis of Stochastic Gradient Methods with Optimism and Anchoring for Minimax Problems and GANs. *ArXiv*, abs/1905.10899.
- Schiro, D. A.; Pang, J.-S.; and Shanbhag, U. V. 2013. On the solution of affine generalized Nash equilibrium problems

with shared constraints by Lemke's method. *Mathematical Programming*, 142(1-2): 1–46. This work was based on research partially supported by the National Science Foundation under grant CMMI-0969600 and the Department of Energy under grant DOE DE-SC0003879.

Solodov, M.; and Svaiter, B. 1999. A Hybrid Approximate Extragradient-Proximal Point Algorithm Using The Enlargement Of A Maximal Monotone Operator. *Set-Valued Analysis*, 7: 323–345.

Uryas'ev, S.; and Rubinstein, R. 1994. On relaxation algorithms in computation of noncooperative equilibria. *IEEE Transactions on Automatic Control*, 39(6): 1263–1267.

von Heusinger, A.; Kanzow, C.; and Fukushima, M. 2012. Newton's method for computing a normalized equilibrium in the generalized Nash game through fixed point formulation. *Mathematical Programming*, 132(1): 99–123.

Walras, L. 1896. *Elements de l'economie politique pure, ou, Theorie de la richesse sociale*. F. Rouge.