Proportional response dynamics in the Fisher market

Li Zhang

Microsoft Research Silicon Valley Lab
1065 La Avenida
Mountain View, CA 94043
lzha@microsoft.com

Abstract

In this paper, we show that the proportional response dynamics, a utility based distributed dynamics, converges to the market equilibrium in the Fisher market with constant elasticity of substitution (CES) utility functions. By the proportional response dynamics, each buyer allocates his budget proportional to the utility he receives from each good in the previous time period. Unlike the tâtonnement process and its variants, the proportional response dynamics is a large step discrete dynamics, and the buyers do not solve any optimization problem at each step. In addition, the goods are always cleared and assigned to the buyers proportional to their bids at each step. Despite its simplicity, the dynamics converges fast for strictly concave CES utility functions, matching the best upperbound of computing the market equilibrium via solving a global convex optimization problem.

1. Introduction

The market equilibrium characterizes the efficient outcome in a competitive market and is a central notion in Economics. While much recent studies have been devoted to computing the market equilibrium, it is desirable, from both economic and computational perspective, to know how such equilibrium emerges when the agents dynamically respond to the market condition in a distributed fashion. In this paper, we show that for certain widely studied markets, namely, the Fisher market with constant elasticity of substitution (CES) utility functions, there is a utility based proportional response dynamics that converges to the market equilibrium, and it may converges fast, matching the bound by solving a global convex program via the ellipsoid or interior point method.

Preprint submitted to Elsevier *October 30, 2009*

We consider the Fisher market in which there are distinct sellers and buyers. Further, each seller has one unit of divisible good for sale (so we do not distinguish seller and good), and each buyer i has a budget b_i and a utility function with the form $u_i(x_1, \cdots, x_n) = \sum_j (w_{ij}x_j)^{\rho_i}$ for some $0 < \rho_i \leq 1$. Such utility functions are the standard constant elasticity of substitution (CES) utility functions. [1] It includes the well studied linear Fisher market by setting $\rho_i = 1$ for each i. We consider the market rule that after the buyers place bids to the goods, each good is allocated to a buyer proportional to the buyer's bid, or equivalently, the price of a good is the sum of the bids placed to that good. By the proportional response dynamics, the buyer submits bids in discrete time steps and adjusts his bids according to the *utility* he receives from each good in the previous time step. Formally, if we denote by $b_{ij}(t)$ the bid of buyer i to good j at time t, then $b_{ij}(t+1) = b_i \frac{u_{ij}(t)}{u_i(t)}$ where $u_{ij}(t) = (w_{ij}\frac{b_{ij}(t)}{p_j(t)})^{\rho_i}$ is the utility received by the buyer i from the good j, $p_j(t) = \sum_i b_{ij}(t)$ is the total bids submitted to the good j, and $u_i(t) = \sum_j u_{ij}(t)$ is the total utility of the buyer i.

From the above description, we can see that the proportional response dynamics is characteristically different from the standard tâtonnement market dynamics. In the tâtonnement process, the price of each good is gradually adjusted according to the excess of demand in the previous time step. The proportional response dynamics does not explicitly involve a price mechanism as it is based on the user's utility. Consequently, it requires much less information and no need to solve an optimization problem at each step. It is naturally distributed and guarantees the market clearance at each step. In addition, it is a large step discrete dynamics in the sense that the buyer does not gradually change his bid, and therefore there is no need to choose a sufficiently small step size as typically done in tâtonnement process. Yet, for CES utility functions, we show that the proportional response dynamics converges to the market equilibrium, and in the case when each $\rho_i < 1$, the proportional response dynamics converges much faster than the discretized tâtonnement process.

In the tâtonnement process, at each time step, the buyer computes the optimum bundle. Such strategy bears similarity to the best response dynamics in multi-player games. In contrast, in the proportional response dynamics, each buyer adapts his bid according to the utility received in the

[1] The standard form is actually $u_i(x_{i1}, \cdots, x_{in}) = (\sum_j (w_{ij}x_{ij})^{\rho_i})^{1/\rho_i}$, to make it homogeneous with degree one.

previous time step. This is similar to the family of dynamics based on the payoff reinforcement learning. Examples include the replicator dynamics in evolutionary games [13] and the multiplicative update algorithm in zero-sum games [12]. In these dynamics, the probability of playing each strategy is adjusted by a multiplicative factor determined by the payoff corresponding to that strategy. As we shall see later in the convergence proof, the proportional response can be reformulated as a multiplicative process. As one contribution of our work, we demonstrate that there exists utility based dynamics that converges to the market equilibrium, a general equilibrium.

Related work. The Fisher market is a special case of the general exchange market. According to [5], it was first defined by American economist Irving Fisher. The linear Fisher market is equivalent to the pari-mutuel method studied in [11]. In [11], Eisenberg and Gale established that the market equilibrium, which they call equilibrium probabilities, is the solution to a convex program, now commonly referred to as the Eisenberg-Gale program, and laid the foundation for many subsequent works. In Computer Science community, [9] first presented a polynomial time algorithm to approximate the market equilibrium in the linear market with bounded number of goods. [10] proposed a polynomial time combinatorial algorithm for computing the market equilibrium for the linear Fisher market. In [14, 19], polynomial time algorithms are presented for computing the market equilibrium by solving the Eisenberg-Gale program.

There has also been a long history in studying the dynamics for converging to the market equilibrium. One particularly well studied dynamics is the tâtonnement process in which the price changes gradually according to the excess of demand. Tâtonnement dynamics was first described by Walras in his monumental work published in 1874 and later formulated and extensively studied in Economics. It was shown in [1, 2, 3] that tâtonnement converges locally for economies satisfying weak gross substitutability (WGS). In [17], Norvig showed that a greedy bidding strategy, which can be regarded as a variant of tâtonnement process, converges to the market equilibrium in the set up considered by Eisenberg and Gale in [11]. More recently, in [7, 6], it is shown that the discretized tâtonnement process converges for WGS utility functions, and [8] showed that an asynchronized variant also converges. In [4], a dynamics is presented for a perturbed keyword auction mechanism and shown to converge to the market equilibrium. That dynamics can also be regarded as a variant of tâtonnement process.

The proportional response dynamics has been studied in [18] for a market that models the bandwidth allocation in the peer to peer file sharing system. One important property in that model is that each good, the upload bandwidth, brings the same utility to the interested users. This is not the case in a Fisher market. Consequently, we can no longer apply the techniques in [18]. In particular, the proportional response is not equivalent to a matrix scaling process, an important tool used in that paper. In this paper, we show that the Kullback-Leibler divergence between the allocation defined by the dynamics and the market equilibrium approaches 0. Our proof is facilitated by the connection between the Eisenberg-Gale program and the market equilibrium. This also renders the proof simpler and the technique more general than that in [18].

Admittedly, compared to the tâtonnement process, the proportional response dynamics applies to more specific markets. It remains an interesting direction to discover similar dynamics that converge to the market equilibrium in more general economies.

2. Preliminaries

Fisher market. A Fisher market is a bipartite market which distinguishes the role of buyer and seller. Each buyer i has a budget b_i, and each seller has a unit of divisible good for sale (and therefore we do not distinguish the sellers and the goods). Suppose that there are m buyers and n goods. Each buyer's utility is defined as a function of the amount of each good he receives. In this paper, we consider the family of markets where a buyer's utility function has the form:

$$u_i(x_{i1}, \cdots, x_{in}) = \sum_{j=1}^{n} (w_{ij} x_{ij})^{\rho_i} ,$$

where $0 < \rho_i \leq 1$, $w_{ij} \geq 0$, and x_{ij} represents the amount of good j allocated to the user i. Such utility functions have constant elasticity of substitution (CES) property and are standard in Economics. One special case is the linear Fisher market when setting $\rho_i = 1$ for all i. Without loss of generality, we assume that for each i, there exists j such that $w_{ij} > 0$, and for each j, there exists i, such that $w_{ij} > 0$. We denote by $\rho_M = \max_i \rho_i \leq 1$.

Market equilibrium. In a Fisher market, for any price vector $\mathbf{p} = (p_1, \cdots, p_n)$ where p_j is the price of the good j, each buyer i can maximize his utility under his budget constraint. The optimum bidding is the solution to the following optimization problem.

$$
\begin{aligned}
&\max u_i(x_{i1}, \cdots, x_{in})\,, \text{s.t.} \\
&\forall j \quad x_{ij} = b_{ij}/p_j\,, \\
&\sum_j b_{ij} \le b_i\,, \\
&\forall j \quad b_{ij} \ge 0\,.
\end{aligned} \tag{1}
$$

If it happens that there exists a solution $\mathbf{b} = \{b_{ij}\}$ to the optimization problem for each buyer i and such that $\forall j \quad \sum_i b_{ij} = p_j$, we call the price vector together with the corresponding bidding and allocation a *market equilibrium*.

It is known that

Lemma 1. *The Fisher market with CES utility functions always has an equilibrium. At the equilibrium, each good's price and each buyer's utility is unique.*

Approximate market equilibrium. The notion of approximate market equilibria is useful for measuring the closeness of an allocation to an equilibrium. Suppose that $\mathbf{p}^*$ is the market equilibrium price. Following [9, 15], the bidding vector $\mathbf{b} = \{b_{ij}\}_{i,j}$, with price vector $\mathbf{p} = \{p_j = \sum_i b_{ij}\}_j$ is an *ε-approximate market equilibrium* if

1. For each j, $(1-\varepsilon)p_j^* \le p_j \le (1+\varepsilon)p_j^*$.
2. For each i, $u_i \ge (1-\varepsilon)\tilde{u}_i$ where $\tilde{u}_i$ is the maximum utility of buyer i given the price vector $\mathbf{p}$.

We also define a stronger notion of the approximate market equilibrium. A bidding vector $\mathbf{b} = \{b_{ij}\}$ is called a *strong ε-approximate market equilibrium* if there exists a market equilibrium $\mathbf{b}^*$ such that for all i, j, $(1-\varepsilon)b_{ij}^* \le b_{ij} \le (1+\varepsilon)b_{ij}^*$. It is easily seen that in the Fisher market with concave utility functions, a strong ε-approximate market equilibrium is an $O(\varepsilon)$-approximate market equilibrium. The reverse might not be true.

Proportional response dynamics. As standard in the study of market dynamics, we consider the setup where at each time step the buyers face the same market parameters, i.e. the same set of goods, budget constraint, and utility function while the buyers make their bidding decisions according to the previous market actions.

Denote by $b_{ij}(t)$ the bid of buyer i on the good j at time t. The proportional response dynamics considered in this paper is defined as $b_{ij}(t+1) =$

$b_i \frac{u_{ij}(t)}{u_i(t)}$, where $p_j(t) = \sum_i b_{ij}(t)$, $u_{ij}(t) = (w_{ij} b_{ij}(t)/p_j(t))^{\rho_i}$, and $u_i(t) = \sum_j u_{ij}(t)$. Therefore, at each step, each buyer allocates the bid proportional to the utility he receives from each good in the previous time period. In addition, we require that $b_{ij}(0) > 0$ whenever $w_{ij} > 0$. We have that,

Lemma 2. *A market equilibrium is a fixed point of the proportional response dynamics.*

Proof. Consider a market equilibrium with the price vector $\mathbf{p}$ and the bidding vector $\mathbf{b}$. By the definition, the market equilibrium is the solution of the optimization problem (1).

Using Lagrangian multiplier, we have that for each i, there exists λ_i such that if $b_{ij} > 0$, then $\rho_i \left(\frac{w_{ij}}{p_j}\right)^{\rho_i} b_{ij}^{\rho_i - 1} = \lambda_i$. Thus, $u_{ij} = \left(\frac{w_{ij} b_{ij}}{p_j}\right)^{\rho_i} = b_{ij} \lambda_i / \rho_i$. That is, for any j, k with $b_{ij}, b_{ik} > 0$, $u_{ij}/u_{ik} = b_{ij}/b_{ik}$. Hence, $\mathbf{b}$ is a fixed point of the proportional response dynamics. $\square$

Main results. The main result of the paper is

Theorem 3. *The proportional response dynamics converges to a market equilibrium in the Fisher market with CES utility functions.*

As for the convergence rate, we distinguish two cases, when $\rho_M < 1$ and when $\rho_M = 1$. Without loss of generality, let $\sum_i b_i = 1$ and $\sum_j w_{ij} = 1$ for every i. Let $W_1 = \frac{1}{\min_{w_{ij}>0} w_{ij}}$, $W_2 = \frac{1}{\min_i b_i}$, $W = nW_1W_2$, and $L = \log W$. Throughout this paper, log denoted the logarithm with base 2. We assume that initially $b_{ij}(0) = \Omega\left(\frac{b_i}{n^{O(1)}}\right)$. This includes the case where each buyer splits his bid evenly among all the goods. About the convergence rate, we have that

Theorem 4. *When $\rho_M < 1$, it takes $O\left(\frac{L + \log(1/\varepsilon)}{(1-\rho_M)^2}\right)$ steps to reach a strong ε-approximate market equilibrium. When $\rho_M = 1$, it takes $O(W^3/\varepsilon^2)$ steps to reach an ε-approximate market equilibrium.*

Since each step takes $O(mn)$ arithmetic computation, when $\rho_M < 1$, the overall running time is bounded by $O(mn(L + \log(1/\varepsilon))/(1-\rho_M)^2)$. We remark that this bound is comparable to the bound of $O((mnL)^{O(1)} \log(1/\varepsilon))$ obtained by solving a convex program via the ellipsoid or interior point methods, and it is much faster than the discretized tâtonnement process [6, 8].

3. The convergence proof

Our proof relies on a characterization of market equilibrium by Eisenberg and Gale. In [11], it is shown that the market equilibrium in the linear Fisher market is the solution to a convex program. Their result easily extends to CES utility functions. Consider the Eisenberg-Gale program defined as

$$\begin{aligned} &\max \sum_i \frac{b_i}{\rho_i} \log u_i\,, \quad \text{s.t.} \\ &\forall i \quad u_i = \sum_j (w_{ij} x_{ij})^{\rho_i}\,, \\ &\forall j \quad \sum_i x_{ij} = 1\,, \\ &\forall i, j \quad x_{ij} \geq 0\,. \end{aligned} \tag{2}$$

The following statement is a straight forward extension of [11].

Lemma 5. *For the Fisher market with CES utility functions, an allocation $\boldsymbol{x} = \{x_{ij}\}$ is an equilibrium if and only if it is a solution to (2). Further, the value of each u_i is unique at a solution of (2).*

Now we proceed to prove the convergence of the proportional response dynamics.

Proof (Theorem 3). The convergence proof consists of two steps. Consider the sequence of bids $\mathbf{b}(t) = \{b_{ij}(t)\}$ for $t = 0, 1, \cdots$. We first show that any limiting point of this sequence is a market equilibrium. This is done by showing that the Kullback-Leibler(KL) divergence between $b_{ij}(t)$ and the market equilibrium converges. This is sufficient to guarantee the convergence when there is a unique market equilibrium, such as in the case where $\rho_M < 1$. When $\rho_M = 1$, an additional argument is needed to show that there can be at most one limiting point starting from any given initial condition.

1. Any limiting point of the dynamics is a market equilibrium.

For any two vectors $\mathbf{x} = \{x_1, \cdots, x_n\}$ and $\mathbf{y} = \{y_1, \cdots, y_n\}$ with non-negative entries, let $\mathrm{KL}(\mathbf{x}\|\mathbf{y})$ denote the KL-divergence

$$\mathrm{KL}(\mathbf{x}\|\mathbf{y}) = \sum_i x_i \log(x_i/y_i)\,.$$

It is well known that when $\sum_i x_i = \sum_i y_i$, then $\mathrm{KL}(\mathbf{x}\|\mathbf{y}) \geq 0$ and the equality holds only when $\mathbf{x} = \mathbf{y}$. Given any bidding vector $\mathbf{b}$, we denote by $\mathbf{b}_i = \{b_{i1}, \cdots, b_{in}\}$ the bidding by the i-th buyer. Take any market equilibrium $\mathbf{b}^* = \{b^*_{ij}\}$. We define the potential function

$$\Phi(t) = \sum_i \frac{1}{\rho_i} \Phi_i(t)\,,$$

where $\Phi_i(t) = \mathrm{KL}(\mathbf{b}^*_i \| \mathbf{b}_i(t))$.

To gain intuition on the definition of the potential function, we have the following dependence of $b_{ij}(t+1)$ on $b_{ij}(t)$.

$$b_{ij}(t+1) = b_i \frac{u_{ij}(t)}{u_i(t)} = \left(\frac{w_{ij}}{p_j(t)}\right)^{\rho_i} \frac{b_i}{u_i(t)} b_{ij}(t)^{\rho_i}\,.$$

Hence, the dynamics can be regarded as a multiplicative process which motivates the use of the KL divergence in the potential function.

We will now show that $\Phi(t)$ converges which in turn implies that any limiting point of the dynamics is a market equilibrium.

Let u^*_i represent the utility of the buyer i and p^*_j the price of good j at the equilibrium. When $b^*_{ij} > 0$, by Lemma 2, we have that $b^*_{ij}/b_i = u^*_{ij}/u^*_i$. Therefore

$$\begin{aligned}
& b^*_{ij} \log \frac{b^*_{ij}}{b_{ij}(t+1)} \\
= \;& b^*_{ij} \log \frac{b^*_{ij} u_i(t)}{b_i u_{ij}(t)} \quad (\text{by that } b_{ij}(t+1) = \tfrac{u_{ij}(t)}{u_i(t)} b_i) \\
= \;& b^*_{ij} \log \frac{u^*_{ij} u_i(t)}{u^*_i u_{ij}(t)} \quad (\text{by Lemma 2 } b^*_{ij}/b_i = u^*_{ij}/u^*_i) \\
= \;& b^*_{ij} \log \frac{u^*_{ij}}{u_{ij}(t)} - b^*_{ij} \log \frac{u^*_i}{u_i(t)}\,.
\end{aligned} \tag{3}$$

Since $u^*_{ij} = (w_{ij} b^*_{ij}/p^*_j)^{\rho_i}$ and $u_{ij}(t) = (w_{ij} b_{ij}(t)/p_j(t))^{\rho_i}$, we have that $\frac{u^*_{ij}}{u_{ij}(t)} = \left(\frac{b^*_{ij}/b_{ij}(t)}{p^*_j/p_j(t)}\right)^{\rho_i}$. Plugging it in (3), we have that

$$\begin{aligned}
& b^*_{ij} \log \frac{b^*_{ij}}{b_{ij}(t+1)} \\
= \;& \rho_i b^*_{ij} \log \frac{b^*_{ij}}{b_{ij}(t)} - \rho_i b^*_{ij} \log \frac{p^*_j}{p_j(t)} - b^*_{ij} \log \frac{u^*_i}{u_i(t)}\,.
\end{aligned} \tag{4}$$

Hence, we have that

$$
\begin{aligned}
& \Phi(t+1) \\
= \; & \sum_i \frac{1}{\rho_i} \Phi_i(t+1) \\
= \; & \sum_i \frac{1}{\rho_i} \sum_j b^*_{ij} \log \frac{b^*_{ij}}{b_{ij}(t+1)} \\
= \; & \sum_i \frac{1}{\rho_i} \sum_j \left(\rho_i b^*_{ij} \log \frac{b^*_{ij}}{b_{ij}(t)} - \rho_i b^*_{ij} \log \frac{p^*_j}{p_j(t)} - b^*_{ij} \log \frac{u^*_i}{u_i(t)} \right) \\
= \; & \sum_i \Phi_i(t) - \sum_j \sum_i b^*_{ij} \log \frac{p^*_j}{p_j(t)} - \sum_i \frac{1}{\rho_i} \sum_j b^*_{ij} \log \frac{u^*_i}{u_i(t)} \\
& (\text{by that } \textstyle\sum_i b^*_{ij} = p^*_j \text{ and } \textstyle\sum_j b^*_{ij} = b_i) \\
= \; & \sum_i \Phi_i(t) - \sum_j p^*_j \log \frac{p^*_j}{p_j(t)} - \sum_i \frac{b_i}{\rho_i} \log \frac{u^*_i}{u_i(t)} \\
= \; & \sum_i \Phi_i(t) - \mathrm{KL}(\mathbf{p}^* \| \mathbf{p}(t)) - \sum_i \frac{b_i}{\rho_i} \log \frac{u^*_i}{u_i(t)} \, .
\end{aligned}
$$

Write $\Psi(t) = \mathrm{KL}(\mathbf{p}^* \| \mathbf{p}(t)) + \sum_i \frac{b_i}{\rho_i} \log \frac{u^*_i}{u_i(t)}$. Then $\Phi(t+1) = \sum_i \Phi_i(t) - \Psi(t)$. We have that

Lemma 6. 1. *For each i, $\Phi_i(t) \geq 0$, and $\Phi_i(t) \leq \Phi(t)$.*

2. *$\Psi(t) \geq 0$, and $\Psi(t) = 0$ if and only if $u_i(t) = u^*_i$ for each i and $p_j(t) = p^*_j$ for each j.*
3. *$\Phi(t+1) \leq \rho_M \Phi(t) - \Psi(t) \leq \Phi(t) - \Psi(t)$.*

Proof. 1. Follows from that for each i, $\sum_j b^*_{ij} = \sum_j b_{ij}(t) = b_i$.

2. By that $\sum_j p^*_j = \sum_j p_j(t) = \sum_i b_i$, we have that $\mathrm{KL}(\mathbf{p}^* \| \mathbf{p}(t)) \geq 0$. By Lemma 5, $\sum_i \frac{b_i}{\rho_i} \log u_i(t) \leq \sum_i \frac{b_i}{\rho_i} \log u^*_i$. That is, $\sum_i \frac{b_i}{\rho_i} \log \frac{u^*_i}{u_i(t)} \geq 0$. Therefore $\Psi(t) \geq 0$. The second half of the statement follows from the equality condition in the above two inequalities and Lemma 5.

3. By that $\Phi_i(t) \geq 0$ and $\rho_M = \max_i \rho_i$, we have that

$$
\Phi(t+1) = \sum_i \Phi_i(t) - \Psi(t) \leq \sum_i \frac{\rho_M}{\rho_i} \Phi_i(t) - \Psi(t) = \rho_M \Phi(t) - \Psi(t) \, .
$$

Since $\rho_M \leq 1$ and $\Phi(t) \geq 0$, $\Phi(t+1) \leq \Phi(t) - \Psi(t)$. $\square$

By the above lemma, we now show that $\Psi(t) \to 0$ when $t \to \infty$. By repeatedly applying Lemma 6.3, we have that $\Phi(t+1) \le \Phi(0) - \sum_{\tau=0}^{t} \Psi(\tau)$. That is

$$\sum_{\tau=0}^{t} \Psi(\tau) \le \Phi(0) - \Phi(t+1) \le \Phi(0)\,. \tag{5}$$

Since $b_{ij}(0) > 0$ if $w_{ij} > 0$, $\Phi(0)$ is upper bounded. Together with the fact that $\Psi(t) \ge 0$, (5) implies that $\Psi(t) \to 0$ when $t \to \infty$. By Lemma 6.2, this in turn implies that $u_i(t) \to u_i^*$ for any i and $p_j(t) \to p_j^*$ for any j. By Lemma 5, any limiting point of the dynamics is a market equilibrium.

2. The dynamics always converges to a single market equilibrium.

When $\rho_M < 1$, the market equilibrium is unique. The proportional dynamics converges to that unique market equilibrium from any initial condition. However, when some $\rho_i = 1$, there may exist multiple market equilibria. We shall show that it is impossible that the sequence $b_{ij}(t)$ has two distinctive limiting points. Suppose that $\mathbf{b}' = \{b'_{ij}\}$ is a limiting point of the sequence $b(t_0), b(t_1), \cdots$. By 1, we know that $\mathbf{b}'$ is a market equilibrium. Since we can choose any market equilibrium in the definition of Φ, we now choose $\mathbf{b}^* = \mathbf{b}'$. Since $\mathbf{b}(t_k) \to \mathbf{b}'$, we have that $\Phi_i(t_k) \to 0$ and therefore $\Phi(t_k) \to 0$ when $k \to \infty$. By that $\Phi(t)$ is monotonically decreasing, and that $\Phi(t) \ge 0$, we have for any infinite strictly increasing sequence $s_0, s_1, \cdots$, $\Phi(s_k) \to 0$ when $k \to \infty$. Therefore, $\mathbf{b}(s_k) \to \mathbf{b}'$. That is, the dynamics always converges to a single market equilibrium. $\square$

4. The rate of convergence

We now bound the convergence rate of the proportional response dynamics. We consider two cases, when $\rho_M < 1$ and when $\rho_M = 1$. In the former case, we are able to show a fast convergence of the dynamics; and in the latter case, we show that the dynamics converges in pseudo-polynomial time.

Without loss of generality, we may scale b_i and w_{ij} such that $\sum_i b_i = 1$ and for each i, $\sum_j w_{ij} = 1$. Recall that $W_1 = \frac{1}{\min_{w_{ij}>0} w_{ij}}$ and $W_2 = \frac{1}{\min_i b_i}$, $W = nW_1W_2$, and $L = \log W$. We have that

Lemma 7. *At the equilibrium, $p_j^* \ge \frac{1}{W}$ for any j. When $\rho_M < 1$, $b_{ij}^* = \Omega\left(\left(\frac{1}{W^2}\right)^{1/(1-\rho_M)}\right)$ whenever $w_{ij} > 0$.*

PROOF. For any j, consider a buyer i with $w_{ij} > 0$. If for every other k, $w_{ik} = 0$, then $p_j^* \geq b_i \geq 1/W_2$. Otherwise, suppose that $w_{ij}, w_{ik} > 0$. As in the proof of Lemma 2,

$$\left(\frac{w_{ij}}{p_j^*}\right)^{\rho_i} {b_{ij}^*}^{\rho_i - 1} = \left(\frac{w_{ik}}{p_k^*}\right)^{\rho_i} {b_{ik}^*}^{\rho_i - 1} . \tag{6}$$

By that $b_{ij}^* \leq p_j^*$, $b_{ik}^* \leq p_k^*$, and $\rho_i \leq 1$, we have

$$\left(\frac{w_{ij}}{p_j^*}\right)^{\rho_i} {p_j^*}^{\rho_i - 1} \leq \left(\frac{w_{ik}}{b_{ik}^*}\right)^{\rho_i} {b_{ik}^*}^{\rho_i - 1} .$$

Rearranging the terms, we have that $p_j^* \geq \left(\frac{w_{ij}}{w_{ik}}\right)^{\rho_i} b_{ik}^* \geq \frac{1}{W_1} b_{ik}^*$. Since there exists k such that $b_{ik}^* \geq b_i/n$, $p_j^* \geq \frac{b_i}{nW_1} \geq \frac{1}{W}$.

When $\rho_M < 1$, by applying (6) again, we have that

$$\begin{aligned} b_{ij}^* &= b_{ik}^* \left(\frac{w_{ij} p_k^*}{w_{ik} p_j^*}\right)^{\rho_i/(1-\rho_i)} \geq b_{ik}^* \left(\frac{1}{W_1} \cdot \frac{1}{nW_1 W_2}\right)^{\rho_i/(1-\rho_i)} \\ &\geq b_{ik}^* \left(\frac{1}{nW_1^2 W_2}\right)^{\rho_M/(1-\rho_M)} = \Omega\left(\left(\frac{1}{W^2}\right)^{1/(1-\rho_M)}\right) . \end{aligned}$$

□

We will need the following technical lemma that bounds the difference between two vectors from their KL divergence.

Lemma 8. *For two positive sequences x_j and y_j for $j = 1, \cdots, n$ that satisfy $\sum_j x_j = \sum_j y_j$, let $\eta = \max_j \frac{|x_j - y_j|}{x_j}$. Then*

$$\mathrm{KL}(\boldsymbol{x} \| \boldsymbol{y}) \geq \frac{1}{16} \min(1, \eta) \eta \min_j x_j .$$

PROOF. We use the well known inequality

$$\sum_j x_j \log(x_j/y_j) \geq \frac{1}{2} \sum_j (\sqrt{x_j} - \sqrt{y_j})^2 .$$

Suppose that $k = \arg\max_j \frac{|x_j - y_j|}{x_j}$. Then

$$\begin{aligned} \mathrm{KL}(\mathbf{x} \| \mathbf{y}) &\geq \frac{1}{2} (\sqrt{x_k} - \sqrt{y_k})^2 = \frac{1}{2} \left(\frac{x_k - y_k}{\sqrt{x_k} + \sqrt{y_k}}\right)^2 \\ &= \frac{1}{2} \eta^2 \frac{x_k^2}{(\sqrt{x_k} + \sqrt{y_k})^2} . \end{aligned}$$

When $y_k \le 2x_k$, $x_k^2/(\sqrt{x_k}+\sqrt{y_k})^2 \ge x_k^2/(\sqrt{x_k}+\sqrt{2x_k})^2 \ge \frac{1}{8}x_k$.
When $y_k > 2x_k$, i.e. when $y = (\eta+1)x_k$ and $\eta > 1$,

$$\frac{x_k^2}{(\sqrt{x_k}+\sqrt{y_k})^2} = \frac{x_k^2}{(\sqrt{x_k}+\sqrt{(1+\eta)x_k})^2} = \frac{1}{(1+\sqrt{1+\eta})^2}x_k \ge \frac{1}{8\eta}x_k\,.$$

The last inequality follows from that $\eta > 1$. Hence, we have that

$$\mathrm{KL}(\mathbf{x}\|\mathbf{y}) \ge \frac{1}{16}\min(1,\eta)\eta x_k \ge \frac{1}{16}\min(1,\eta)\eta \min_j x_j\,.$$

$\square$

According to Lemma 8, in order to show $\max_j \frac{|x_j-y_j|}{x_j} \le \varepsilon < 1$, it suffices to show that $\mathrm{KL}(\mathbf{x}\|\mathbf{y}) = O(\varepsilon^2)\min_j x_j$. Now, we are ready to prove Theorem 4.

Proof (Theorem 4). (1) When $\rho_M < 1$, we show that the dynamics reaches a strong ε-approximate market equilibrium. By Lemma 6.2 and 3, $\Phi(t+1) \le \rho_M\Phi(t) - \Psi(t) \le \rho_M\Phi(t)$. Applying the inequality iteratively, we have that $\Phi(t) \le \rho_M^t\Phi(0)$, and $\Phi_i(t) \le \Phi(t) \le \rho_M^t\Phi(0)$.

As observed earlier, $\Phi_i(t) = \sum_j b_{ij}^*\log\frac{b_{ij}^*}{b_{ij}(t)}$ is the KL divergence between b_{ij}^* and $b_{ij}(t)$. By Lemma 8, $b_{ij}(t)$ is a strong ε-approximate equilibrium as long as $\Phi_i(t) \le \frac{1}{16}\varepsilon^2\min_{j:b_{ij}^*>0} b_{ij}^*$ for $\varepsilon < 1$. Of course, how fast the process converges also depends on the initial choice of $b_{ij}(0)$. By the assumption that $b_{ij}(0) = \frac{b_i}{n^{O(1)}}$, $\Phi_i(0) = \sum_j b_{ij}^*\log\frac{b_{ij}^*}{b_{ij}(0)} = \sum_j b_{ij}^*\log\frac{b_i}{b_i/n^{O(1)}} = O(b_i\log n)$. Hence $\Phi(0) = O(\log n)$ by that $\sum_i b_i = 1$.

By choosing $t = c\frac{L+\log(1/\varepsilon)}{(1-\rho_M)^2}$ for sufficiently large c, we have that $\Phi(t) \le \rho_M^t\Phi(0) = O((\frac{1}{W^2})^{1/(1-\rho_M)}\varepsilon^2)$. By Lemma 7 and 8, we have that $\frac{|b_{ij}^*-b_{ij}(t)|}{b_{ij}^*} \le \varepsilon$ for any i,j, that is $\mathbf{b}(t) = \{b_{ij}(t)\}$ is a strong ε-approximation to $\mathbf{b}^*$.

(2) When $\rho_M = 1$, the convergence could be slower. Indeed, it may never converge to a strong ε-approximate equilibrium as at the market equilibrium, some b_{ij}^* may be 0 even when $w_{ij} > 0$. We will show that it converges to the standard notion of ε-approximate equilibrium in $O(W^3/\varepsilon^2)$ steps. We will establish the bound in the worst case of $\rho_i = 1$ for all i.

By Lemma 6.3, we have that $\Phi(t+1) \le \Phi(t) - \Psi(t)$. We claim that if $\Psi(t) \le \delta$ for $\delta = \frac{1}{256WW_2^2}\varepsilon^2$, then $\mathbf{b}(t)$ is an ε-approximate equilibrium. Recall that $\Psi(t) = \mathrm{KL}(\mathbf{p}^*\|\mathbf{p}(t)) + \sum_i b_i\log\frac{u_i^*}{u_i(t)}$. In what follows, we omit (t)

for the simplicity of the notations. Since $\mathrm{KL}(\mathbf{p}^*\|\mathbf{p}) \geq 0$ and $\sum_i b_i \log \frac{u_i^*}{u_i} \geq 0$. By that $\Psi \leq \delta$, we have that

$$\begin{aligned} \mathrm{KL}(\mathbf{p}^*\|\mathbf{p}) &\leq \delta\,, && (7)\\ \sum_i b_i \log \frac{u_i^*}{u_i} &\leq \delta\,. && (8) \end{aligned}$$

By (7), Lemma 7 and 8, we have that for every j, $(1-\varepsilon')p_j^* \leq p_j \leq (1+\varepsilon')p_j^*$ where $\varepsilon' = \varepsilon/(4W_2)$. Let $\tilde{u}_i$ denote the maximum utility of the buyer i under the price vector $\mathbf{p}$. It remains to show that for each i, $u_i \geq (1-\varepsilon)\tilde{u}_i$. Since $\tilde{u}_i = b_i \max_j w_{ij}/p_j$ and $u_i^* = b_i \max_j w_{ij}/p_j^*$, we have that $u_i^* \geq (1-\varepsilon')\tilde{u}_i \geq (1-\varepsilon')u_i$. With $\varepsilon' < 1/2$, for any i,

$$\begin{aligned} b_i \log \frac{u_i^*}{u_i} &= \sum_j b_j \log \frac{u_j^*}{u_j} - \sum_{j\neq i} b_j \log \frac{u_j^*}{u_j}\\ &\leq \delta - \sum_{j \neq i} b_j \log(1-\varepsilon')\\ &\leq \delta - \log(1-\varepsilon') \leq \delta + 2\varepsilon'\,. \end{aligned}$$

Therefore, $\frac{u_i^*}{u_i} \leq 2^{(\delta+2\varepsilon')/b_i} \leq 1 + W_2(\delta+2\varepsilon')$ as δ, ε' are sufficiently small. Hence $u_i \geq (1-W_2(\delta+2\varepsilon'))u_i^* \geq (1-W_2(\delta+2\varepsilon'))(1-\varepsilon')\tilde{u}_i$. By the choice of δ, we have that $\delta, \varepsilon' \leq \varepsilon/(4W_2)$ and consequently $u_i \geq (1-\varepsilon)\tilde{u}_i$. Hence, $\mathbf{b}$ is an ε-approximate market equilibrium.

By (5), $\sum_{\tau=0}^{t} \Psi(\tau) \leq \Phi(0) = O(\log n)$. Hence, when $t \geq c \log n/\delta$ for some constant $c > 0$, there exists $\tau \leq t$ such that $\Psi(\tau) \leq \delta$. We thus have that the dynamics converges to an ε-approximate market equilibrium in $O(WW_2^2 \log n/\varepsilon^2) = O(W^3/\varepsilon^2)$ steps. □

While we do not know if the above convergence rate for the linear Fisher market is tight, the following example shows that it requires $\Omega(1/\varepsilon)$ steps to converge to an ε-approximate equilibrium. Consider the market with three buyers and two goods. The budgets of the buyers are $b_1 = 2$, $b_2 = 1$, and $b_3 = 1$, respectively. The weights are $w_{11} = w_{21} = w_{22} = w_{32} = 1$ and 0 otherwise. At the equilibrium, the prices of both goods are 2, and $b_{21}^* = 0$, $b_{22}^* = 1$. However, it is easy to see that when $b_{21} = 2\varepsilon$, it takes $\Theta(1/\varepsilon)$ steps to reduce it to ε.

5. Conclusion

One crucial property used in the convergence proof is the equivalence between the market equilibrium and the solution to the Eisenberg-Gale program. It would be interesting to know if the technique can extend to other

Eisenberg-Gale markets as defined in [16]. We note that the proportional response dynamics most naturally applies to separable utility functions, i.e. the utility of each buyer is the sum of the utility obtained from different goods. It would be interesting to know if similar dynamics can be defined for more general families of such utility functions. It is also interesting to know if the dynamics converges under certain asynchronous model.

References

[1] Arrow, K., Block, H. D., Hurwicz, L., 1959. On the stability of the competitive equilibrium, II. Econometrica 27 (4), 82–109.

[2] Arrow, K., Hurwicz, L., 1958. On the stability of the competitive equilibrium, I. Econometrica 26 (4), 522–552.

[3] Arrow, K. J., Hurwicz, L., 1960. Competitive stability and weak gross substitutability: the Euclidean distance approach. International Economic Review 1 (1), 38–49.

[4] Borgs, C., Chayes, J. T., Immorlica, N., Jain, K., Etesami, O., Mahdian, M., 2007. Dynamics of bid optimization in online advertisement auctions. In: WWW. pp. 531–540.

[5] Brainard, W., Scarf, H., 2000. How to compute equilibrium prices in 1891. Cowles Foundation Discussion Paper 1270.

[6] Codenotti, B., McCune, B., Varadarajan, K. R., 2005. Market equilibrium via the excess demand function. In: STOC. pp. 74–83.

[7] Codenotti, B., Pemmaraju, S. V., Varadarajan, K. R., 2005. On the polynomial time computation of equilibria for certain exchange economies. In: SODA. pp. 72–81.

[8] Cole, R., Fleischer, L., 2008. Fast-converging tatonnement algorithms for one-time and ongoing market problems. In: STOC. pp. 315–324.

[9] Deng, X., Papadimitriou, C., Safra, S., 2002. On the complexity of equilibria. In: STOC. pp. 67–71.
URL `citeseer.ist.psu.edu/deng02complexity.html`

[10] Devanur, N., Papadimitriou, C., Saberi, A., Vazirani, V., 2002. Market equilibrium via a primal-dual-type algorithm. In: FOCS. pp. 389–395.
URL `citeseer.ist.psu.edu/article/devanur02market.html`

[11] Eisenberg, E., Gale, D., 1959. Consensus of subjective probabilities: the Pari-Mutuel method. Annals of Mathematical Statistics 30, 165–168.

[12] Freund, Y., Schapire, R. E., 1999. Adaptive game playing using multplicative weights. Games and Economic Behavior 29, 79–103.

[13] Fudenberg, D., Levine, D. K., 1998. The Theory of Learning in Games. The MIT Press, Ch. 3.

[14] Jain, K., 2004. A polynomial time algorithm for computing the Arrow-Debreu market equilibrium for linear utilities. In: FOCS. pp. 286–294.

[15] Jain, K., Mahdian, M., Saberi, A., 2003. Approximating market equilibria. In: RANDOM-APPROX. pp. 98–108.

[16] Jain, K., Vazirani, V. V., 2007. Eisenberg-gale markets: algorithms and structural properties. In: STOC. pp. 364–373.

[17] Norvig, T., 1967. Consensus of subjective probabilities: A convergence theorem. Annals of Mathematical Statistics 38 (1), 221–225.

[18] Wu, F., Zhang, L., 2007. Proportional response dynamics leads to market equilibrium. In: STOC. pp. 354–363.

[19] Ye, Y., 2008. A path to the Arrow-Debreu competitive market equilibrium. Mathematical Programming 111 (1-2), 315–348.