On-Line Algorithms for Market Equilibria

Spyros Angelopoulos* Atish Das Sarma* Avner Magen Anastasios Viglas *
University of Waterloo IIT Bombay University of Toronto University of Sydney

Abstract

We consider a variation of the classical problem of finding prices which guarantee equilibrium in linear markets consisting of divisible goods and agents with money. Specifically, we consider *on-line* algorithms for this problem in which goods are considered on-line, and each good is assigned an irrevocable price. Since exact equilibria will not be found in such a setting, we appeal to the concept of approximate equilibrium defined in previous studies of the problem, to characterize the quality of our solutions. We consider both deterministic and randomized algorithms for finding approximate equilibria. We prove a tight bound on the performance of deterministic algorithms, and show that under certain natural conditions, randomized algorithms lead to market prices which are closer to equilibrium.

1 Introduction

The existence of equilibria in markets is a central problem in mathematical economics, and has attracted enormous interest since the pioneering work of Walras [10] and Fisher [3]. The problem is the following. Consider a market which consists of buyers, each with a certain amount of money, and divisible goods of a certain amount each. The desirability of goods to each of the buyers is expressed by the utility functions of the buyers. The goal is to assign prices to goods such that the buyers can buy their individually optimal bundle of goods in terms of the utility they get and there is no surplus of goods, i.e, the market clears. The theorem of Arrow and Debreu [1] established the existence of equilibrium prices in a very general setting. The proof is nonconstructive, yet recently polynomial-time algorithms for the case of linear utility functions have been presented. These algorithms assume complete knowledge of the entire market meaning that the utility functions of all goods are known beforehand. This might be the case in a static market, but it does not reflect the situation in a *dynamic* market, in which little, if anything, is known about goods that will appear in the market in the future.

This work focuses on the performance of on-line algorithms for computing or predicting equilibrium prices. The online algorithm assigns prices incrementally trying to approximate the actual equilibrium prices that correspond to the actual offline market problem. Prices are assigned without complete knowledge of the market and cannot change. In a more general setting we can formulate the restriction on re-assigning prices by associating a change in a previously assigned price of a good with a cost which measures how undesirable such a change would be. The online setting addressed in this work captures the special case where this cost is infinite and therefore any price assignment is irrevocable. In other words, the assigned price is advertised and must be honoured as advertised even after other products appear on the

*This work was done while the author was at the Deprtment of Computer Science, the University of Toronto.

market. Also, when a good appears, it comes with a survey that specifies how much various customers desire it, i.e. with a complete apecification of its utilities to all agents.

As we mentioned, in our setting the algorithm assigns prices trying to approximate the price equilibrium that we would get in an offline setting. It quickly becomes obvious that in the setting described above, the exact equilibrium are not generally computable. Hence we appeal to the concept of an *approximate* equilibrium as introduced by Deng, Papadimitriou and Safra [5] (precise definitions are given in Section 2). Prices set by online algorithms will not necessarily correspond to an exact market equilibrium, but we may still be looking for the best possible approximation, which is an approximate equilibrium. The two main parameters of an approximate equilibrium are market clearence and optimizing the pay-off for each agent (personal optimal bundle). We distinguish between two types of approximate market clearance: we can require that every good is cleared to a certain extent, or we can require that on average the goods clear. The former definition associates each good with a distinct seller and is the one adopted in the study of (off-line) equilibria, while the latter captures a setting in which there is only one seller that is satisfied when goods are cleared on average.

Our model is quite different from previous auction-type algorithms [2] for market clearence, based on online matching algorithms. Those models enforce the extra restriction that a price *and* an allocation or sale must be done online in an irrevocable way. Our model only requires that the price is set and advertised, while the allocation of goods, or the sales can be done at the end when all prices have been set. This gives a characterization of the best approximation of equilibrium prices in a different market setting, and our results indicate that it is much more difficult to get good equilibrium approximations.

Previous Work. The Arrow-Debreu theorem [1] states that in general divisible markets, equilibrium prices always exist. The theorem applies in the general setting in which each agent (buyer) has an initial endowment of goods, which she can then trade to acquire other goods, but it also applies to markets in which the agents have money with which they buy goods (and hence there is a clear distinction between sellers and buyers). Such markets are called *Fisher* markets [3, 6], and they are focus of this paper.

Computing market equilibria has been a long-standing problem in economics [9] but it was only recently that it was approached from the point of view of algorithmic solutions with strict guarantees. Deng, Papadimitriou and Safra [5] introduced the concept of an approximate equilibrium: essentially, one seeks an allocation of goods which approximately clears the market and for which every agent is approximately maximally happy. Deng *et al* considered the problem of approximating equilibria in endowment markets with linear utility functions (linear markets). They showed that, for indivisible goods, the problem is NP-hard to approximate within 1/3, and it is NP-hard even when the number of agents is two. In contrast, for the case where the number of agents is fixed, they provided a $(1+\epsilon)$-approximation. For divisible goods they showed an exact (polynomial-time) algorithm provided that the number of goods, or the number of agents, is bounded.

For linear Fisher markets with divisible goods, Devanur, Papadimitriou, Saberi and Vazirani [4] presented a polynomial-time exact algorithm based on the primal-dual schema. Jain, Mahdian and Saberi [8] used the algorithm of [4] to provide a FPTAS in the more general setting of endowment markets. This algorithm was recently improved to a strongly polynomial algorithm by Devanur and Vazirani [6]. More work on market equilibria was presented in [7] on the spending constraint model.

Our results. We provide upper and lower bounds on the deviation-from-equilibrium of an algorithm for the problem with respect to the two definitions of clearance. Here, the deviation-from-equilibrium is defined as the smallest k for which the prices are $(1-1/k)$-equilibrium prices (the precise definition of this measure and its rational is given in Section 2). For deterministic algorithms we show a (tight) bound of $\Theta(\min\{\sqrt{m}, n\})$ for both individual and global clearance (n and m are the number of agents and goods respectively). We then turn our attention to the application of randomization in the context of this problem; this is motivated by the observation that an algorithm with access to random bits could possibly avoid (on an average case) bad prices for the whole sequence of goods. We show a randomized algorithm that acheives a better deviation-from-equilibrium that its deterministic counterpart, when clearence is required "on the average". Specifcally, we show how to get deviation-from-equilibrium $O(\min\{m^{1/3}\sqrt{\log m}, n\})$ albeit for a somewhat restricted, but still fairly broad and natural family of inputs. Furthermore, we provide a lower bound of $\Omega(\min\{m^{1/3}, \sqrt{n}\})$. For the case of individual clearance we show that randomization does not actually help, as we show a lower bound of $\Omega(\min\{\sqrt{m}, n\})$ to the deviation-from-equilibrium under this criterion.

2 Problem Definition and Preliminaries

A market consists of a set A of n *agents* and a set G of m divisible *goods*. Each good is characterized of its *size* b_j. Agent i possesses a certain amount of *money* $e_i \in \mathbb{R}^+$, which she can use to buy goods in G. A *bundle* of goods for agent i is a vector $\mathrm{x}_i \in \mathbb{R}^m$ such that $\mathrm{x}_{ij} \leq b_j$. A *feasible allocation* (or simply allocation) $\mathbf{x}$ is a collection of n bundles $\mathrm{x}_1, \ldots, \mathrm{x}_n$ (one for each agent), such that, for every j, $\sum_{i=1}^n x_{ij} \leq b_j$. The *utility function* of i is a function $u_i : \mathbb{R}^m \to \mathbb{R}$; in particular, $u_i(\mathrm{x}_i)$ specifies the utility of agent i for bundle x_i (informally, it represents the happiness of i if she buys a quantity $x_{ij} \leq b_j$ of good j). Throughout this paper, we assume *linear* utility functions, that is $u_i(\mathrm{x}_i) = \sum_{j=1}^m u_{ij} x_{ij}$, for non-negative constants u_{ij}. We call u_{ij} the *utility of j for i*.

Suppose that good j is assigned a *price* $p_j \in \mathbb{R}^+$. Since agent i wants to maximize her utility, the *optimal bundle* for i is the bundle $\tilde{\mathrm{x}}_i$ which is the solution to the following maximization program:

$$\begin{aligned} &\text{maximize} \quad u_i(\mathrm{x}_i) \\ &\text{subject to} \quad \sum_{j=1}^m p_j \mathrm{x}_{ij} \leq e_i. \end{aligned} \tag{1}$$

Informally, the optimal bundle for agent i maximizes the utility i can make without taking into consideration the presence of other agents. Can we set prices so that we can find an allocation that consists of bundles that are close to optimal, for all users? Clearly, if prices are high enough then this requirement is met, as there are no conflicts between different buyers demanding more than the supply. However, we also wish to achieve *market clearance*, in the sense that there is no surplus or deficiency of goods. The concept of a *market equilibrium* aims to strike a balance between these two conflicting goals. More formally, an *ϵ-approximate equilibrium* ($0 \leq \epsilon \leq 1$), or ϵ-equilibrium for brevity, is a price vector $\mathbf{p} \in \mathbb{R}^m_+$ such that there exists an allocation $\mathrm{x} = \{\mathrm{x}_1, \ldots, \mathrm{x}_n\}$ with the following two properties:

1. For all i, $u_i(\mathrm{x}_i) \geq (1-\epsilon) u_i(\tilde{\mathrm{x}}_i)$. $\tilde{x}_i$ is the solution to the maximization program (1).

2. The market approximately clears. Let $c_j = \frac{1}{b_j}\sum_{i=1}^n x_{ij}$, namely the fraction of good j that was bought. In *individual clearance* we require that $\min_j c_j \geq 1-\epsilon$, in other words we consider the worst case clearance of the goods. Alternatively, we can look at a more relaxed definition and require that the *average* of the c_j-s is at least $1-\epsilon$. We call this type of clearance *global clearance.*

A 0-equilibrium is simply called an *equilibrium*. We call the price vector relative to which one can find an allocation that leads to an ϵ-equilibrium an *ϵ-equilibrium price*.

Without loss of generality, in linear markets and for both variants of clearance, we can normalize the size of each good to 1 unit by scaling the utilities appropriately. We notice that subject to this normalization global clearance simply states that the total number of units bought is at least $(1-\epsilon)m$.

The setting of the market equilibrium problem assumes that all information is available to an algorithm for the problem. In this paper we consider the *on-line* version of the market equilibrium problem. More specifically, we assume that goods arrive on-line. Every time a good j appears, the utilities u_{ij} are revealed, for all $i \in [n]$. The on-line algorithm must assign an *irrevocable* price to each good at the time of its appearance, that is the price of goods cannot be modified throughout the algorithm's execution[1].

We motivate our definition of a *deviation-from-equilibrium* of an on-line algorithm. When considering approximation algorithms for a maximization problem which achieves a value v while the optimal value is τ, we either say the approximation ratio is $k = \tau/v$ or, when v approaches τ, e.g. in the case of a PTAS, we rather consider that minimal ϵ for which $v \geq (1-\epsilon)\tau$. A similar situation arises here. We define the deviation-from-equilibrium of an on-line algorithm as the smallest k for which the prices set by the algorithm are $(1-1/k)$-equilibrium prices.

Additional definitions. A good j is called *uniform* if for every two agents i, i', $u_{ij} = u_{i'j}$. A set of goods is called *monotone* if for every two goods j, j' and every two agents i, i', we have $u_{ij} \leq u_{ij'} \Rightarrow u_{i'j} \leq u_{i'j'}$. The definition asserts that the ordering of monotone goods by utility is the same for every agent. The *aspect ratio* of the market is defined as $\max_{i,j,j'} \frac{u_{ij}}{u_{ij'}}$. The good j that maximizes the ratio u_{ij}/p_j is called the *best* good for agent i.

3 Deterministic Algorithms

Theorem 3.1 *There exists a deterministic on-line algorithm for the problem with deviation-from-equilibrium $O(\min\{\sqrt{m}, n\})$, for individual (and therefore also for global) clearance. Furthermore, this bound is tight.*

Proof:
Upper bound. We consider individual clearance, which clearly implies global clearance with the same guarantees. First, we provide some intuition behind the assignment of prices. If the price of each good is set as high as E (recall that E is the total money of all agents), then clearly there is an allocation in which every agent gets as much utility as from its optimal bundle; however at most one unit of good is allocated, and we are far from market clearance. On the other extreme, if the prices are very low, e.g., E/m, then all goods can be allocated, however there will be contention between agents for goods that are important to a large subset of agents. We reconcile the two extremes by assigning prices as follows: the price of the j-th good, for all $j \leq n^2$ is set to $\frac{E}{\sqrt{j}}$. For more details see the appendix.

Lower Bound. The adversary will present to n agents, each having a unit amount of money, a sequence of m uniform goods; that is, for every good j and agent i, $u_{ij} = u_j$. The intuition behind the adversarial input is that when the algorithm considers a good of very high utility

[1]Alternatively, one could assume a model in which agents appear on-line, each revealing how much utility she can make from each good. Such a model would make sense only if the price of every good is set before any allocation (even for the very first agent) takes place. However, if prices are set, the problem is trivial to solve for linear markets, by using linear programming.

for all agents, then it has to set a high price to it, otherwise there will be heavy contention between agents. By providing a sequence of "progressively better" goods, in terms of their utility, the adversary will force the algorithm to assign high prices to every good, which implies poor market clearance. For more details see the appendix. ∎

The upper bound of Theorem 3.1 is tight, provided that the aspect ratio of the market is sufficiently high. What if the aspect ratio is bounded? The following theorem builds upon the idea behind the lower-bound proof of Theorem 3.1. The intuition here is that goods in the sequence become progressively better, but only as much as possible given the bound on α.

Theorem 3.2 *The deviation-from-equilibrium of every deterministic algorithm for a market with aspect ratio* $\alpha > 1$ *is in* $\Omega(\min\{nb, \sqrt{mb}\})$, *where* $b = \frac{\beta^m(\beta-1)}{\beta^{m+1}-1}$, *and* $\beta = \alpha^{\frac{1}{m}}$.

Proof: We will consider global clearance; the lower bound then carries over to individual clearance. As in the proof of Theorem 3.1 the adversary will present to n agents, each having a unit amount of money, a sequence of m uniform goods (we will denote by u_j the utility of the j-th good in the sequence, with $u_1 = 1$). Every time the algorithm assigns a "low" price to good $j < m$, then good $j+1$ is is such that $u_{j+1} = \beta \cdot u_j$ (hence if all j first goods were assigned low prices, then $u_{j+1} = \beta^j$). Otherwise, the adversary presents $m-j$ goods, each of (low) utility β^{j-m}, and terminates the game. We will assume (without loss of generality) that the algorithm knows m in advance, and that the deviation-from-equilibrium of the algorithm, say k, for $k \geq 1$ is a function of m, n, $\mathbf{e}$ and α *only*.

Consider good j in the sequence, assuming that no good in $[j]$, with $l \leq m-1$ has received a low price. We claim that if the price p_j is "low" this provides a lower bound to the deviation-from-equilibrium k by the following relaton.

$$\beta^j \cdot \min\{1, \frac{1}{p_j}\} < k \cdot (\frac{\sum_{l=1}^{j} \beta^l + (m-j)\beta^{j-m}}{n}). \tag{2}$$

To see this, note that in such a case, the adversary will terminate the game by providing $m-j$ goods of utility β^{j-m}. Then, for every allocation of the m goods to agents, there exists one agent, say i, who will receive utility at most $(\sum_{l=1}^{j} \beta^l + (m-j)\beta^{j-m})/n$. On the other hand, the optimal bundle for i yields utility at least equal to the LHS of (2). Taking into account the fact that $\frac{\beta^j}{\sum_{l=1}^{j} \beta^l + (m-j)\beta^{j-m}} \geq \frac{\beta^m}{\sum_{l=1}^{m} \beta^l} = b$, we conclude that, for the algorithm to have deviation-from-equilibrium k, the price p_j must not be low, namely it must be such that

$$n \min\{1, \frac{1}{p_j}\} \cdot b \leq k, \tag{3}$$

If $p_j \leq 1$, (3) gives $k \geq nb$. Otherwise, (3) implies that $p_j \geq \frac{n}{k} b$, for all $j \in [m-1]$. Hence at most $(Ek)/(bn) + 1 = k/b + 1$ goods can be allocated to agents, and since we require that the algorithm has deviation-from-equilibrium k, it must be that $\frac{m}{k/b+1} \leq k$, thus $k \in \Omega(\sqrt{mb})$. Summarizing, $k \in \Omega(\min\{nb, \sqrt{mb}\})$. ∎

Theorem 3.2 demonstrates that when α is exponential on m, then the algorithm used in the proof of Theorem 3.1 is asymptotically optimal.

It is worth mentioning that for the special case where all goods are identical (and hence the only information not known to the on-line algorithm is the number of goods), a variant of this approach can be employed to show that no deterministic algorithm has deviation-from-equilibrium better than a certain constant bigger than 1. We omit the details.

4 Randomized Algorithms

A randomized on-line algorithm is an algorithm which assigns irrevocable prices to goods according to a certain probability distribution. We start by stating precisely the definition of deviation-from-equilibrium of such an algorithm. Consider an allocation of goods after the prices have been set. Let F_i denote the random variable which is the utility agent i can make from this allocation. Let also G_i denote the random variable that corresponds to the utility of the optimal bundle for agent i. Last, define H_j to be the random variable which denotes how much of good j was bought, and $H = \sum_j H_j$. Then, $K_G = \max\{\max_i G_i/F_i, m/H\}$ and $K_I = \max\{\max_i G_i/F_i, \max_j 1/H_j\}$ are the random variables that denote the global and individual clearance respectively. Let $k_G = \mathbb{E}[K_G]$ and $k_I = \mathbb{E}[K_I]$ be the corresponding expectations of these variables and they are the ones we consider. We start with a negative result.

Theorem 4.1 *The deviation-from-equilibrium of every randomized on-line algorithm for markets with global clearance $\Omega(\min\{m^{\frac{1}{3}}\sqrt{n}\})$. For markets with individual clearance it is $\Omega(\min\{\sqrt{m}, n\})$.*

Proof: Using Yao's principle, we present a distribution on inputs on which every deterministic algorithm has high expected deviation-from-equilibrium . Let I_j be the input $u_1 << u_2 << \ldots << u_j$ and $u_{j+1}, \ldots, u_m = 0$. Take input I_j for $j < m$ with probability $\frac{1}{2m}$ and I_m with probability $\frac{1}{2} + \frac{1}{2m}$. Let $p_1, p_2, \ldots, p_m$ be the algorithm answers to I_m. There is a subtle point to note here: since I_j and I_m are consistent for the first j goods, it must be the case that prices set by the algorithm for the input I_j are with agreement to those set for the input I_m, namely $p_1, \ldots, p_j$. Arguments similar to the ones in 3.2 show that considering agent j, at the event of input I_j we have the bound $K_G, K_I \geq n \cdot \min\{1, 1/p_j\}$, and at the event of inout I_m we have (from clearance constriants) that $K_G \geq m/l$ where l is the maximal number of goods totaling to at most n. So

$$k_G = \mathbb{E}[K_G] \geq \max\left\{\frac{n}{2m}\sum_i \min\{1, 1/p_j\}, \frac{m}{2l}\right\}.$$

We apply the first part of Lemma 4.2 to conclude $k_G \geq \frac{1}{2}\min\{\sqrt{n}, m^{\frac{1}{3}}\}$.

Considering K_I, we notice that if the individual clearance is at most K, a quantity of $1/K$ of each good must be purchased by the agents. This means that $\frac{1}{K} \cdot \sum_i p_i \leq E = n$, and therefore we get

$$k_I = \mathbb{E}[K_I] \geq \max\left\{\frac{n}{2m}\sum_i \min\{1, 1/p_j\}, \frac{\sum_i p_i}{2n}\right\}.$$

By the second part of Lemma 4.2 we get $k_I = \Omega(\min\{\sqrt{m}, n\})$. ∎

Lemma 4.2 *Let n, m be positive integer numbers, p_i be nonnegative reals, and $l = l(p_1, \ldots, p_m; n)$ is the maximal l such that the sum of the l smallest p_i does not exceed n. Then*

$$\max\left\{\frac{n}{m} \cdot \sum_i \min\{1, 1/p_i\}, \frac{m}{l}\right\} \geq \min\{\sqrt{n}, m^{\frac{1}{3}}\},$$

and also

$$\max\left\{\frac{n}{2m}\sum_i \min\{1, 1/p_j\}, \frac{\sum_i p_i}{2n}\right\} = \Omega\{m^{1/2}, n\}.$$

Proof: Let $S = \{j : p_j < 1\}$ and L be the set if indices of the l smallest p_i-s (cutting ties arbitrarily). Further, let $B = L \setminus S$, $s = |S|$ and $b = |b|$. Clearly $l = s + b$. Now,

$$\sum_{i \in B} \min\{1, 1/p_i\} = \sum_{i \in B} 1/p_i = b \cdot \frac{\sum_{i \in B} 1/p_i}{b} \geq b \cdot \frac{b}{\sum_{i \in B} p_i} \geq b^2/n.$$

The first inequality above is the Arithmetic-Harmonic-Mean inequality. So

$$\sum_i \min\{1, 1/p_i\} \geq \sum_{i \in S} \min\{1, 1/p_i\} + \sum_{i \in B} \min\{1, 1/p_i\} \geq s + b^2/n$$

and we get

$$\max\left\{\frac{n}{m} \cdot \sum_i \min\{1, 1/p_i\}, \frac{m}{l}\right\} \geq \max\left\{\frac{n}{m} \cdot (s + b^2/n), \frac{m}{s+b}\right\}.$$

It is now enough to show that for every choice of nonnegative b and s the inequality $\max\left\{\frac{n}{m} \cdot (s + b^2/n), \frac{m}{s+b}\right\} \geq \min\{m^{\frac{1}{3}}, \sqrt{n}\}$ holds. First assume $s = 0$. Here we need to optimize $\max\{b^2/m, m/b\}$ which is clearly at least $m^{\frac{1}{3}}$. We now turn to the case $s > 0$. We will use the simple inequality $\max\left\{\frac{n}{m} \cdot (s + b^2/n), \frac{m}{s+b}\right\} \geq \frac{1}{2}\left(\frac{n}{m} \cdot (s + b^2/n) + \frac{m}{s+b}\right)$, and will now lower bound the latter expression. Assume b is fixed and we need to find that value of s minimizing the expression. Easy calculus shows that $s = m/\sqrt{n} - b$ is that value. We can safely assume $m/\sqrt{n} - b > 0$ otherwise $s = 0$ would be the best choice which is a case we already covered. Substituting for s we get

$$\frac{n}{m} \cdot (s + b^2/n) + \frac{m}{s+b} \geq \frac{n}{m} \cdot (m/\sqrt{n} - b + b^2/n) + \sqrt{n}.$$

Now we optimize over b and get that $b = n/2$ is the minimizing value for the last expression. By the same argument we may assume here that $n/2 \leq m/\sqrt{n}$. We substitute b for $n/2$ and get

$$\frac{n}{m} \cdot (m/\sqrt{n} - b + b^2/n) + \sqrt{n} = 2\sqrt{n} - \frac{1}{4}n^2/m \geq \sqrt{n}.$$

For the second part of the lemma, Using the same definition $S = \{j : p_j < 1\}$ we get

$$\sum_i \min\{1, 1/p_i\} = \sum_{i \in S} \min\{1, 1/p_i\} + \sum_{i \notin S} \min\{1, 1/p_i\} \geq s + \sum_{i \notin S} 1/p_i \geq$$

$$\geq s + \frac{(m-s)^2}{\sum_{i \notin S} p_i} = s + \frac{(m-s)^2}{P},$$

where $s = |S|$ and $P = \sum_{i \notin S} p_i$. Now

$$k \geq \frac{1}{2} \max\left\{\frac{n}{m} \sum_i \min\{1, 1/p_j\}, \frac{\sum_i p_i}{n}\right\} \geq \frac{1}{4}\left(\frac{n}{m} \sum_i \min\{1, 1/p_j\} + \frac{\sum_i p_i}{n}\right) \geq$$

$$\frac{1}{4}\left(\frac{n}{m}(s + (m-s)^2/P) + \frac{P}{n}\right) = \Omega\{m^{1/2}, n\}.$$

The last quantity is easily verified by checking separately for $s < m/2$ and for $s \geq m/2$. ∎

Can we get a better upper bound by using randomization? The following result shows that for monotone goods the answer is positive. It should be noted that the adversarial input of both our deterministic and randomized lower bounds complies to the condition of monotonicity.

Theorem 4.3 *For markets with* monotone *goods and* global *clearance, there exists a randomized on-line algorithm with deviation-from-equilibrium* $O(\min\{m^{1/3}\sqrt{\log m}, n\})$.

Proof: We first show how to get a $O(\min\{m^{1/3}\log m, n\})$ first. Let $\tilde{j}$ denote $2^{\lceil \log j \rceil}$. The algorithm assigns prices to goods according to the following probability distribution. For every good $j \leq n^3$

$$p_j = \begin{cases} E/\tilde{j}^{2/3} & \text{with probability } 1/\tilde{j}^{1/3} \\ E/\tilde{j}^{1/3} & \text{otherwise} \end{cases}$$

We call a good j *cheap* when it receives price $E/\tilde{j}^{2/3}$ and *expensive* otherwise [2]. Next, we set the prices of all goods $j > n^3$, if any, to be arbitrarily low; we call such goods *free* goods.

First, we want to show that with this setting the market approximately clears. Suppose first that $m \leq n^3$. Consider the set of goods $B = \{j : \tilde{j} \geq m/4\}$ and let B_c be the cheap goods in B. Clearly $|B| \geq m/2$ and that $\mathbb{E}[|B_c|] \geq |B|/(2m)^{\frac{1}{3}} = \Omega(m^{2/3})$. Further, Chernoff bound guarantees that $\mathbb{P}[|B_c| \leq \frac{1}{2}\mathbb{E}[|B_c|] = \exp(-\Omega(m^{2/3}))$. Look at the following strategy. It is easy to see that agents can allocate $1/4$ of their money to buy $\Omega(m^{1/3})$ goods j for which $m/8 \leq \tilde{j} < m/4$, as the prices for these goods is at most $O(E/m^{1/3})$. Another $1/4$ is used to buy as many goods in B_c. Since prices there are $O(E/m^{2/3})$ we can buy as many as $\min\{|B_c|, m^{2/3}\}$ such goods. We get that the number of goods that can be bought g is $\Omega(m^{2/3})$ with probability $1 - \exp(-\Omega(m^{2/3}))$ and $\Omega(m^{1/3})$ otherwise which gives $\mathbb{E}[m/g] = \Omega(m^{1/3})$. In the case where $m > n^3$ the agents can still achieve an expected ratio of $n^{3\cdot\frac{1}{3}} = n$ on the first n^3 by exactly the same argument, and then buy all $m - n^3$ free goods at arbitrary low cost which can only improve the ratio. Hence, we only need to be concerned about finding an allocation that yields high utility to each agent. Since each agent spent at most half of her money for the allocations above and since constant factors do not affect our bounds, we may assume for simplicity that these allocations did not cost money at all.

Clearly, if $n < m^{1/3}$ we are done, since agents can always achieve at least $1/n$ of their optimal bundle. We now get to the interesting part of the theorem where $m \leq n^3$. Notice that we can think of the goods as organized in *bunches* by their $\tilde{j}$ value. Namely the bunches are of size $1, 2, 4, \ldots$ (with the exception of the last one which may be smaller). In the allocation we will describe, every agent will spend at most a fraction of $1/w$ of her money, for some $w > 1$ to a good that has price E/w. Whenever all agents comply to this restriction there is no deficiency of the goods, or informally, there is no contention between agents about any good, and we can consider the allocation of each agent individually.

Let j_i be any of the best goods for agent i; recall that these are the good maximizing the ratio utility per price. In the event where j_i is an expensive good, its price is at least $E/m^{1/3}$. In this case agent i will spend $e_i/m^{1/3}$ on buying j_i and so she receives at least $1/m^{1/3}$ of her optimal utility. Otherwise j_i is a cheap good. We now introduce some additional notation. Let u_i be the utility of j_i, b_i the bunch containing j_i and C_i are the goods in bunch b_i with utility at least u_i. Finally, let $t_i = \min\{|C_i|, T_i\}$, S_i be the size of bunch b_i and $T_i = S_i^{1/3}$. Agent i spends e_i/T_i on each of t_i goods from C_i. Therefore the quantity of goods in C_i which agent i gets is exactly $t_i \cdot \frac{e_i/T_i}{E/T_i} = e_i t_i/E$, and since the goods in C_i have at least utility u_i we get a total utility of at least $u_i e_i t_i/E$. On the other hand, an upper bound to the optimal utility for agent i is $\frac{u_i e_i}{E/T_i^2}$ Letting K_i be the ratio between the utility from optimal bundle and

[2] Note that by this definition there are cheap goods that cost more than expensive ones..

utility from allocation for i we get

$$K_i \leq \frac{T_i^2 e_i u_i}{e_i u_i t_i} = T_i^2/t_i. \tag{4}$$

At this point, we show that instead of considering the n different random variables K_i we can look at $\log m$ very related variables: For each bunch b let us choose an ordering that is consistent with the utilities for all agents. In other words if j and j' are goods in bunch b and $u_{ij} > u_{ij'}$ then good j appears before j' in the ordering. Such ordering is possible thanks to the monotonicity of goods. Let r_b be the minimal ordinal of a cheap good according to the ordering. As usual S_b is the size bunch b and $T_b = S_b^{1/3}$. Finally, let $\nu_b = \min\{r_b, T_b\}$. It is now easy to see that if $b_i = b$ then $\nu_b \leq t_i$. Defining $L_b = \max\{K_i : b_i = b\}$ we now get that

$$L_b \leq \max_{\{i|b_i=b\}} T_i^2/t_i \leq T_b^2/\nu_b.$$

The maximum ratio of optimal utility over utility through allocation over all agents is now

$$\max\{m^{1/3}, \max_i K_i\} = \max\{m^{1/3}, \max_b L_b\} \leq m^{1/3} + \sum_b L_b.$$

and so $\mathbb{E}[K] \leq m^{1/3} + \sum_b \mathbb{E}[T_b^2/\nu_b]$.

Lemma 4.4 *Let $S = \{1, 2, \ldots, s\}$ and pick elements in S independently with probability $q = 1/T$ where $T \leq s$. Let r be the minimal number picked or s if none exists. Also, let $\nu = \min\{r, T\}$. Then $\mathbb{E}[1/\nu] = \log T/T$.*

Proof: As long as $r \leq T$, it is distributed geometrically with parameter q. Now

$$\begin{aligned}\mathbb{E}[1/\nu] &= \mathbb{P}[r \leq T] \cdot \mathbb{E}[1/\nu | r \leq T] + \mathbb{P}[r > T] \cdot \mathbb{E}[1/\nu | r > T] \leq \sum_{r=1}^{T} \frac{q(1-q)^r}{r} + 1/T \leq \\ &\leq qH(T) + 1/T \leq O(\log T/T).\end{aligned}$$

∎

Applying the lemma we get that $\mathbb{E}[1/\nu_b] = O(\log T_b/T_b)$. Summing up we have

$$\mathbb{E}[K] \leq m^{1/3} + \sum_b \mathbb{E}[T_b^2/\nu_b] \leq m^{1/3} + \sum_b O(T_b \log T_b) = O(m^{1/3} \log m).$$

To get the desired $O(m^{1/3}\sqrt{\log m})$ bound we notice the imbalance between the two criteria for the equilibrium. We can multiply the prices of cheap goods by a factor of $\sqrt{\log \tilde{j}}$. Obviously the clearance does not deteriorates by more than a factor of $\sqrt{\log m}$. On the other hand, we get the improved upper bound for L_b for the second criterion, $L_b \leq T_b^2/(\sqrt{\log S_b}\nu_b)$. From this the improvement in the bound follows easily. ∎

5 Discussion and Future Work

In this work we formulate and address the problem of on-line equilibria in linear Fisher markets with divisible goods. A number of unresolved issues remain; these are not only of theoretical interest, but they relate to situations which are expected to occur in practice. Among them we first distinguish the following two: First, is it possible to show better (deterministic or randomized) upper bounds for markets with subexponential aspect ratio α? Ideally, we would like to provide an on-line algorithm, which has no knowledge of the aspect ratio of the entire sequence, and whose deviation-from-equilibrium is a function of α.

Second, is it possible to remove the monotonicity condition from Theorem 4.3? The result is meant to show that randomization is helpful when considering global clearance, but we would like to extend it to capture more general markets. We believe that an elaborate probabilistic argument will be needed to address this issue.

A different extension deals with markets with indivisible goods. We have some preliminary results for this type of markets. It is worth noting that here one has to provide certain restrictions on the sequence of the goods the adversary will provide, otherwise no on-line algorithm may achieve bounded deviation-from-equilibrium . Note also that since an (exact) equilibrium does not necessarily exist in this setting, we must relate to the best possible (approximate) equilibrium that an off-line algorithm can achieve.

As argued in the introduction, our work is motivated by dynamic markets, where a cost is associated with any change in the price of an existing good. Our on-line upper and lower bounds address the case in which the cost is infinite. But what about other cost functions? For instance, suppose that the cost to change a price is constant, and there is a strict bound on the total cost due to price changes every time a good arrives. What is the best approximate equilibrium we can guarantee in such a setting?

References

[1] K. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22:265–290,1954.

[2] Avrim Blum and Tuomas Sandholm and Martin Zinkevich, Online algorithms for market clearing, *SODA '02: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, 2002, pp 971–980

[3] W.C. Brainard and H.E. Scarf. How to compute equilibrium prices in 1891. *Cowles Foundation Discussion Paper 1270*, 2000.

[4] N. Devanur, C. Papadimitriou, A. Saberi and V. Vazirani. Market equilibrium via a primal-dual type algorithm. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pp. 389–395, 2002.

[5] X. Deng, C. Papadimitriou and S. Safra. On the complexity of equilibria. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pp. 67–71, 2002.

[6] N. Devanur and V. Vazirani. An improved approximation scheme for computing market equilibrium. In *Proceedings of the 23rd Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, 2003.

[7] Nikhil R. Devanur, Vijay V. Vazirani. Extensions of the spending constraint-model: existence and uniqueness of equilibria (extended abstract). *ACM Conference on Electronic Commerce* 2003: 202-203

[8] K. Jain and M. Mahdian and A. Saberi. Approximating market equilibria. In *Proceedings of the 6th Workshop on Approximation algorithms for Combinatorial Optimization Problems (APPROX)*, 2003.

[9] H.E. Scarf. *The computation of Economic Equilibria (with collaboration of T. Hansen).* Cowles Foundation Monograph No. 24. Yale University Press, 1973.

[10] L. Walras. *Éléments d'économie politique pure; ou, théorie de la richesse sociale (Elements of Pure Economics, or the theory of social wealth).* Lausanne, Paris, 1874. (1899, 4th ed.; 1926 rev. ed. , 1954 Engl. transl.).

6 Appendix

Proof: (of Theorem 3.1)

Proof details for the upper bound: For all $j > n^2$, the price of good j is set to be infinitesimally small; without loss of generality, we can assume it is zero. Suppose first that $m \leq n^2$, then the allocation of goods to agent i is done in two phases. First, i uses a fraction of $1/(4\sqrt{m}) \leq 1/4$ of her money toward her best good (that is she buys $\frac{1}{4}e_i/(\sqrt{m}p_j) \leq \frac{1}{4}$ of that good. Note that the utility of i is at least $1/(4\sqrt{m})$ times the utility of the optimal bundle of i. In the second phase agents buy, collectively, a fraction of $1/(4\sqrt{m})$ of each good and it is easy to see that at most $E/2$ money is needed in this phase.

Details of the lower bound proof: We consider global clearance; the lower bound then carries over to individual clearance. The game between the adversary and the algorithm begins with the adversary presenting a good with $u_1 = 1$, and proceeds as follows. Every time the algorithm assigns a price to good $j < m$ that is not "low" (the meaning of "low" price will become clear shortly), then the $j+1$ good is such that $u_{j+1} \gg u_j$ (more precisely, $u_{j+1} > n \cdot \sum_{l=1}^{j} u_l$). Otherwise, the adversary presents $m - j$ goods, each of very low utility (without loss of generality, we may assume the their utility is zero), and terminates the game. For convenience, we may assume that the algorithm knows m in advance–this does not affect the lower bound. We will also assume that the deviation-from-equilibrium of the algorithm, say k, is a function of m, n, and $\mathbf{e}$ *only*.

Consider good j in the sequence, where none of the goods $1, \ldots, j-1$ received low prices. For every allocation of goods $1, \ldots, j$ to agents, there is at least one agent, say i, who is allocated at most a fraction of $1/n$ of good j. Therefore, i can receive utility at most $\frac{2u_j}{n}$ from goods in $[j]$. On the other hand, the optimal bundle (among goods $1, \ldots, j$ only) for i allocates as much of good j to i as her money can buy. Suppose that the price assigned to good j by the algorithm is low, namely that

$$\min\{u_j, \frac{e_1}{p_j}u_j\} > k(2u_j/n).$$

In this event, the adversary will terminate the game by providing $m-j$ goods of zero utility, as explained earlier, and it is easy to verify that the deviation-from-equilibrium of the algorithm is worse than k. Therefore, p_j must be chosen such that

$$\min\{u_j, \frac{e_1}{p_j}u_j\} \leq k(2u_j/n). \tag{5}$$

If $p_j \leq e_i$, then (5) implies that $k \geq n/2$. Otherwise, (5) gives $p_j \geq n/(2k)$ (recall that $e_i = 1$ for all i).

When the game terminates, either $p_j \geq n/(2k)$ for every good $j < m$ or $k \geq n/2$. Remains to consider only the former case. Observe that, at such prices, at most $\frac{E}{n/(2k)} + 1 = 2k+1$ units of goods can be allocated to agents. For the algorithm to has a deviation-from-equilibrium of k, we require that $\frac{m}{2k+1} \leq k$, hence $k \geq \sqrt{m}/2 - 1$. Thus $k \geq \min\{n/3, \sqrt{m}/2 - 1\}$. ∎