

Distributed Algorithms via Gradient Descent for Fisher Markets

Benjamin Birnbaum^{*}
University of Washington
birnbaum@cs.washington.edu

Nikhil R. Devanur
Microsoft Research, Redmond
nikdev@microsoft.com

Lin Xiao
Microsoft Research, Redmond
lin.xiao@microsoft.com

ABSTRACT

Designing distributed algorithms that converge quickly to an equilibrium is one of the foremost research goals in algorithmic game theory, and convex programs have played a crucial role in the design of algorithms for Fisher markets. In this paper we shed new light on both aspects for Fisher markets with linear and spending constraint utilities. We show fast convergence of the *Proportional Response* dynamics recently introduced by Wu and Zhang [WZ07]. The convergence is obtained from a new perspective: we show that the Proportional Response dynamics is equivalent to a gradient descent algorithm (with respect to a Bregman divergence instead of euclidean distance) on a convex program that captures the equilibria for linear utilities. We further show that the convex program easily extends to the case of spending constraint utilities, thus resolving an open question raised by [Vaz10b]. This also gives a way to extend the Proportional Response dynamics to spending constraint utilities. We also prove a technical result that is interesting in its own right: that the gradient descent algorithm based on a Bregman divergence converges with rate $O(1/t)$ under a condition that is weaker than having Lipschitz continuous gradient (which is the usual assumption in the optimization literature for obtaining the same rate).

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics; F.2.0 [Analysis of Algorithms and Problem Complexity]: General

General Terms

Algorithms, Economics, Theory

Keywords

Market equilibrium, Proportional Response, Gradient descent, Distributed algorithm

^{*}Part of this work was done while at Microsoft Research, Redmond.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'11, June 5–9, 2011, San Jose, California, USA.

Copyright 2011 ACM 978-1-4503-0261-6/11/06 ...\$10.00.

1. INTRODUCTION

The computational complexity of equilibria has been the focus of much research in Algorithmic Game Theory over the last decade, not to mention considerable amount of work in the economics literature. Although most papers consider algorithms in the centralized model of computation, what are more relevant are *distributed algorithms* that converge to an equilibrium. Such algorithms are especially important in the context of automated markets that have emerged due to the Internet, such as the market for online advertising. In such markets trading is often done by automated agents. Distributed algorithms can be thought of as protocols that these agents are programmed to follow. The design and analysis of distributed algorithms converging to equilibria in the context of games has also received considerable attention, most commonly convergence of best response dynamics [AGM⁺08, CS07, Ros73, AAE⁺08].

For markets, essentially the only distributed algorithms that have been studied are variants of the tatonnement process. Wu and Zhang [WZ07] proposed a new distributed algorithm, called the *Proportional Response* (PR) dynamics, as a protocol for trading bandwidth on a peer-to-peer network.¹ The PR dynamics involves a sequence of bids by the buyers and prices by the sellers that respond to each other. The seller prices are simply set to be the sum of all the bids they receive. The buyers set their bids proportional to the utility they would obtain with the bids and the prices in the previous round.

PR seems to converge really fast to equilibrium in practice, and yet the reason for this remained a mystery. Zhang [Zha09] analyzed PR dynamics for Fisher markets and showed the following: “if τ is large enough, then there exists a $t < \tau$ such that the bids and prices in round t are close to the equilibrium.” This does not quite show *convergence* of the dynamics to an equilibrium. Moreover the proof uses a potential function that gives little insight into why this happens.

In this paper we *demystify* the PR dynamics: we show that the PR dynamics are equivalent to a generalized gradient descent algorithm with Bregman divergences (instead of Euclidean distance) on a convex program introduced by Shmyrev [Shm09] that captures the equilibria of Fisher markets with linear utilities. This insight gives an intuitive understanding of why such dynamics work. PR dynamics fit so well with this perspective that we can confidently say this is the “right” way of thinking about them. Moreover, we get true convergence, of the form, “if t is large enough, then the bids and prices in round t are close to an equilibrium.”

Further we show that the convex program and the PR dynam-

¹It has since been shown to be an effective solution for BitTorrent [LLSB08].

ics generalize to the *spending constraint* utilities, thus resolving an open question of Vazirani [Vaz10b].

1.1 Fisher's Market Model with Linear and Spending Constraint Utilities

In the Fisher market model with linear utilities, there are n buyers and m perfectly divisible goods. There is a unit² supply of each good. Each buyer i has a budget B_i , and for each item j , a number u_{ij} , which represents the utility of that buyer for one unit of item j . Given an allocation vector $x \in \mathbb{R}_+^{n \times m}$, the utility of buyer i is $\sum_j u_{ij}x_{ij}$.

An allocation $x \in \mathbb{R}_+^{n \times m}$ and price vector $p \in \mathbb{R}_+^m$ are an *equilibrium* if two conditions are satisfied. The first, *buyer optimality*, is that given p , each player maximizes his utility subject to his budget constraint. In other words, for all i , the allocation x optimizes the following program (where p is a constant):

$$\begin{aligned} & \text{maximize} && \sum_j u_{ij}x_{ij} \\ & \text{subject to} && \sum_j p_j x_{ij} \leq B_i, \\ & && x_{ij} \geq 0, \forall j. \end{aligned}$$

The second equilibrium condition, *market clearance*, is that $\sum_i x_{ij} = 1$ for all j .

In the *spending constraints* model, the utility of a buyer also depends on the prices, as follows. The utility is still additive across goods, that is, the total utility for a bundle of goods is the sum of utilities for each good separately. For a given good, the utility function is divided into segments; each segment l has a constant rate of utility u_{ij}^l , and a budget B_{ij}^l . The utility of the buyer for x_{ij}^l amount of good j under segment l is $u_{ij}^l x_{ij}^l$, but is subject to the constraint that $p_j x_{ij}^l \leq B_{ij}^l$. Hence the utility maximizing program for a given buyer is now as follows.

$$\begin{aligned} & \text{maximize} && \sum_{j,l} u_{ij}^l x_{ij}^l \\ & \text{subject to} && \sum_{j,l} p_j x_{ij}^l \leq B_i, \\ & && 0 \leq x_{ij}^l \text{ and } p_j x_{ij}^l \leq B_{ij}^l \forall j, l. \end{aligned}$$

1.2 Proportional Response Dynamics

The *Proportional Response (PR) Dynamics* is a distributed algorithm for computing market equilibrium and was introduced by Wu and Zhang [WZ07] in the context of a bandwidth trading market model for peer-to-peer networks and by Zhang [Zha09] in the current context of the Fisher markets. The algorithm proceeds in discrete time steps, and for each time $t = 0, 1, 2, \dots$, buyer i submits a bid $b_{ij}(t)$ for each good j . Given the bids, a price vector $p(t)$ is computed by summing the bids for each item: $p_j(t) = \sum_i b_{ij}(t)$. Each buyer i is allocated an amount $x_{ij}(t) = b_{ij}(t)/p_j(t)$ of item j . The buyers then update their bids for the next time step so that their new bid for a good is proportional to the *utility obtained* from this good in the current round, that is $b_{ij}(t) \propto u_{ij}x_{ij}(t)$. In summary, the new bid vector is computed according to the following recursion:

$$b_{ij}(t+1) = B_i \frac{u_{ij}x_{ij}(t)}{\sum_{j'} u_{ij'}x_{ij'}(t)}. \quad (1)$$

The PR dynamics has many desirable properties that are also unique.

- The PR dynamics are simple, distributed and require no global communication.³ In every step, a buyer only has to know the

²This is without loss of generality, by the choice of unit, since the goods are divisible.

³Except for the fact that they need to be synchronized, in the sense that the dynamics happen in rounds and in every round, every bidder updates his bids.

prices of the goods he is interested in, and a good only has to know the bids placed on it.

- PR dynamics is also *stateless*: the bids at any time step depend only on the bids of the previous step, and there is no special starting point.
- Unlike all other dynamics, which depend on choosing the right step size, PR dynamics does not need a step size to choose and adapt.

In this paper, we analyze the PR dynamics from the perspective of general convex optimization algorithms. Along the way, we give a new convex program for the linear Fisher market and establish a new convergence condition for generalized gradient descent methods based on Bregman divergences. With this general framework, we are able to extend both the new convex program and the PR dynamics to markets with spending constraint utilities.

2. MAIN RESULTS

2.1 The Convex Program

A well-known convex program called the Eisenberg-Gale [EG59] convex program captures equilibrium allocations of a linear Fisher market. Here we consider an alternate convex program discovered recently by Shmyrev [Shm09] that also does the same. As we shall see this convex program is closely related to the Proportional Response (PR) dynamics. The convex program is as follows.

$$\begin{aligned} & \text{maximize} && \sum_{i,j} b_{ij} \log u_{ij} - \sum_j p_j \log p_j \\ & \text{subject to} && \forall j, \sum_i b_{ij} = p_j, \\ & && \forall i, \sum_j b_{ij} = B_i, \\ & && \forall i, j, 0 \leq b_{ij}. \end{aligned} \quad (2)$$

The variable p_j corresponds to the price of good j . The variable b_{ij} represents the amount of money spent on good j by buyer i . Hence given a solution to the above program, the allocation x_{ij} is given by b_{ij}/p_j . It is easy to see that for any feasible solution to the program, the corresponding prices and allocations are such that the market always clears, and the buyers always exhaust their budgets. The only other condition required for equilibrium is that each buyer be allocated his optimal bundle of goods.

THEOREM 1 ([SHM09]). *Let b^* and p^* be an optimum solution to Convex Program (2). Then p^* and allocation $x_{ij}^* = b_{ij}^*/p_j^*$ is an equilibrium for the Fisher market with linear utilities given by the u_{ij} 's and the B_i 's.*

This convex program can be generalized to markets with spending constraint utilities. This resolves an important open problem raised by Vazirani [Vaz10b].

$$\begin{aligned} & \text{maximize} && \sum_{i,j,l} b_{ij}^l \log u_{ij}^l - \sum_j p_j \log p_j \\ & \text{subject to} && \forall j, \sum_{i,l} b_{ij}^l = p_j, \\ & && \forall i, \sum_{j,l} b_{ij}^l = B_i, \\ & && \forall i, j, l, 0 \leq b_{ij}^l \leq B_{ij}^l. \end{aligned} \quad (3)$$

As before, the allocation is given by $x_{ij}^l = b_{ij}^l/p_j$, and the only condition that needs verifying is buyer optimality.

THEOREM 2. *An optimum solution to Convex Program (3) corresponds to an equilibrium for the Fisher market with spending constraint utilities given by the u_{ij}^l 's, B_{ij}^l 's and the B_i 's.*

2.2 Generalized Gradient Descent with Bregman Divergence

Both of the convex programs (2) and (3) are special cases of the constrained optimization problem

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && x \in C \end{aligned} \quad (4)$$

where C is a compact convex set, and $f(x)$ is a differentiable convex function with $\text{dom } f \supseteq C$ (i.e., f is finite on C). The conventional *projected gradient descent* method (e.g., [Ber99]) for solving (4) is

$$x(t+1) = \pi_C(x(t) - \mu_t \nabla f(x(t))) \quad (5)$$

where μ_t is an appropriate stepsize, and $\pi_C(\cdot)$ denotes Euclidean projection on the set C . Let

$$\ell_f(x; y) := f(y) + \langle \nabla f(y), x - y \rangle$$

be the linear approximation of f using the gradient at y . The method (5) is equivalent to

$$x(t+1) = \arg \min_{x \in C} \left\{ \ell_f(x; x(t)) + \frac{1}{2\mu_t} \|x - x(t)\|_2^2 \right\}, \quad (6)$$

where the *proximal* term $(1/2)\|x - x(t)\|_2^2$ is the squared Euclidean distance between two points. In the *generalized gradient descent* method the squared Euclidean distance in (6) is replaced by a generalized distance function $d(x, x(t))$, i.e.,

$$x(t+1) = \arg \min_{x \in C} \{ \ell_f(x; x(t)) + \gamma_t d(x, x(t)) \}, \quad (7)$$

where the parameter $\gamma_t > 0$, and its reciprocal $1/\gamma_t$ play the role of stepsize μ_t in the Euclidean case. In particular, we let $d(x, y)$ be the *Bregman divergence* [Bre67] of a differentiable convex function $h(x)$ (assuming $\text{dom } h \supseteq C$), defined as

$$d_h(x, y) = h(x) - \ell_h(x; y), \quad \forall x \in C, y \in \text{rint } C,$$

where $\text{rint } C$ denotes the *relative interior* of C [Roc70, Section 6]. As an example, $d_h(x, y) = (1/2)\|x - y\|_2^2$ if $h(x) = (1/2)\|x\|_2^2$. We often drop the subscript h whenever there is no confusion from context. We say h is the *kernel* of d , and d is generated by h . We assume that both f and h are differentiable and they satisfy

$$f(x) \leq \ell_f(x; y) + \gamma d_h(x, y), \quad \forall x \in C, y \in \text{rint } C, \quad (8)$$

for some constant $\gamma > 0$. We show convergence of the generalized gradient descent method under the assumption of (8), which is weaker than the typical assumption of Lipschitz continuity of the gradient of f .

THEOREM 3. *Suppose that the sequence $x(t)$ is defined by Equation (7), f and h are convex and differentiable functions and that they satisfy Condition (8). Then for all t ,*

$$f(x(t)) - f(x^*) \leq \frac{\gamma d(x^*, x(0))}{t}.$$

In particular we show the following two lemmas from which the above theorem follows immediately.

LEMMA 4. *The total deviation of f from the optimum throughout the run of the algorithm is bounded by a value that is independent of how long the algorithm is run.*

$$\sum_t (f(x(t)) - f(x^*)) \leq \gamma d(x^*, x(0)).$$

LEMMA 5. *For all t , $f(x(t+1)) \leq f(x(t))$.*

2.3 PR as Gradient Descent

One can easily see the similarities between the PR dynamics (1) and Convex Program (2) in the way the prices and allocations are defined. In particular, the market clears automatically in the PR dynamics and the only condition we need to worry about is buyer optimality. In fact as detailed below, we show that the PR dynamics is exactly the gradient descent algorithm given by (7) on Convex Program (2).

First of all, we can eliminate the p_j variables from the program (2); in fact, p_j can be thought of as a function of the b_{ij} 's, defined as $p_j(b) = \sum_i b_{ij}$. We stick to the notation p_j without the argument when there is no confusion⁴ from the context. Let the objective function⁵ of the convex program be denoted by $\varphi(b) = -\sum_{i,j} b_{ij} \log u_{ij} + \sum_j p_j(b) \log(p_j(b))$, and the feasible region by

$$S = \left\{ b \in \mathbb{R}^{n \times m} : \sum_j b_{ij} = B_i \forall i, b_{ij} \geq 0 \forall i, j \right\}.$$

Without loss of generality assume that $\sum_i B_i = 1$. Also define the Bregman divergence d with the kernel function being the unnormalized negative entropy, $h(b) = \sum_{i,j} (b_{ij} \log b_{ij} - b_{ij}) = \sum_{i,j} b_{ij} \log b_{ij} - \sum_i B_i$ for $b \in S$. A straightforward calculation shows that

$$d(a, b) = \sum_{i,j} a_{ij} \log \frac{a_{ij}}{b_{ij}} = D(a \| b),$$

where $D(\cdot \| \cdot)$ is the *KL-divergence* between two probability distributions.

THEOREM 6. *The proportional response dynamics (1) is equivalent to the gradient descent algorithm (7) with $d(a, b) = D(a \| b)$ as above, $f = \varphi$, $C = S$ and $\gamma = 1$.*

In order to use our results about convergence of the above algorithm (Theorem 3), we need to show that Condition (8) is satisfied, which we do in the following lemma.

LEMMA 7. *For all $a, b \in S$,*

$$\varphi(b) \leq \ell_\varphi(b; a) + D(b \| a).$$

Thus, PR dynamics satisfies the convergence properties as specified in Theorem 3 with $f = \varphi$ and $\gamma = 1$. If the initial bids are $b_{ij}(0) = B_i/m$ for all i and j , then $D(b^* \| b(0)) \leq \log(mn)$ (see Lemma 13). A more detailed analysis of the convergence is given in Section 5.

For markets with spending constraints, the algorithm (7) leads to an *iterative PR and capping* algorithm, which has the same convergence properties as the PR dynamics (see Section 4.1).

2.4 Related Work and Motivation

The algorithmic study of equilibrium concepts in general, and of market equilibria in particular, are motivated by the question of whether markets can feasibly operate at equilibrium. After almost a decade since this line of research started, the computational complexity of equilibrium concepts is fairly well understood. The hardness results [CT09, CDDT09, CD06, DGP06, CSVY06] point

⁴Similarly, we use $p_j(t)$ to denote $p_j(b(t))$ and p_j^* to denote $p_j(b^*)$.

⁵To maintain consistency with our notation in the previous subsection, we rephrase the convex program as minimizing a convex function instead of maximizing a concave function.

towards the unlikeliness that markets in general operate at equilibria, with recent results [CT09, CDDT09, VY10] showing that even for the simple class of separable piecewise linear functions, it is PPAD-hard to compute an equilibrium. The algorithmic results [DPS02, DPSV08, Jai07, DV04, CPV05, CMV05, JV10, DK08, Ye08] provide hope for certain special classes of markets. A good survey of algorithms for market equilibria can be found in [Vaz07] and [CV07], with the summary being that we can find equilibria for markets with weak gross substitutes property, for certain resource allocation markets [JV10], and when there are a bounded number of goods [DK08].

Algorithmic results in a centralized model of computation do not directly address the question of market *dynamics*: how might agents interacting in a market arrive at an equilibrium? Here, the quest is for simple and distributed algorithms that are guaranteed to converge fast. Such distributed algorithms are especially applicable when the agents involved are automated, and one has to prescribe a particular *protocol* for them to follow. This is true in many networking applications, and also in markets such as search advertising where often the bids are updated automatically. There is also a huge amount of work in the networking community on designing such distributed algorithms that achieve *proportional fairness* (e.g., [Kel97, KMT98]). It is known that proportional fairness is equivalent to market equilibrium in many settings (including linear utilities), via an Eisenberg-Gale type convex program [KV]. Our results could potentially be useful for these applications as well, and is a direction for future work.

Almost all dynamics considered in the literature are variants of the *tatonnement* process first defined by Leon Walras [Wal74]. The most general version of tatonnement just says that prices are increased when demand exceeds supply and are decreased when supply exceeds demand. Many versions of this [ABH59, Uza60, CMV05, FGK⁺08, GK06] have been analyzed over the years, but none of these until recently gave simple, fast and distributed algorithms. Some of the early work in economics [ABH59, Uza60] ignored the convergence rate entirely. [CMV05] showed polynomial time convergence of a version of tatonnement for markets with weak gross substitutes, but had to transform the market *a priori* which then translates to passing around global information in every round. The auction algorithm of Garg and Kapoor [GK06] requires that it be started from a particular state. The algorithm of [FGK⁺08] suffers from disadvantages such as having to take the average of all the prices and the average of all the allocations to get an equilibrium. The exception is the algorithm given by [CF08] which is distributed and asynchronous and converges fast for markets with weak gross substitutes property. The convergence time however depends on certain market parameters which tend to infinity for linear Fisher markets. One reason for the difficulty in designing tatonnement style processes for linear Fisher markets is that the demand of a buyer is not always uniquely determined. Thus at equilibrium one needs to specify the allocation in addition to the prices. The PR dynamics elegantly takes care of the problem of non-uniqueness of demand, since the bids uniquely determine the prices and the allocations.

PR was first proposed as a protocol for trading bandwidth on a peer-to-peer network by Wu and Zhang [WZ07] and has been shown to be an effective solution for BitTorrent [LLSB08]. Zhang [Zha09] shows the following about PR dynamics in Fisher markets: “if τ is large enough, then there exists a $t < \tau$ such that $b(t)$ is close to optimal.” We get true convergence, of the form, “if t is large enough, then $b(t)$ is close to optimal.” In fact, we rephrase Zhang’s proof and give a detailed comparison with ours in Section 5.2. We extend the results to spending constraint util-

ities as well. Clearly the extension was only possible because of our insights into the connection between PR dynamics and the convex program. Moreover, Zhang [Zha09] also considers Constant Elasticity of Substitution (CES) utilities, $U_i = \sum_j (u_{ij} x_{ij})^{\rho_i}$ with $0 < \rho_i \leq 1$. He gets a better bound on the convergence rate for these utilities: PR converges to an ϵ -approximate equilibrium in

$$O\left(\frac{L + \log(1/\epsilon)}{1 - (\max_i \rho_i)^2}\right)$$

rounds, where L is the bit-complexity of the input. Our techniques extend to CES utilities as well, and we strengthen Zhang’s results as before, by obtaining the stronger form of convergence. The details of this will be presented in the full version.

The convex program (2) was recently discovered by Shmyrev [Shm09], and was actually rediscovered by us. The convex program (2) is in fact related to the EG convex program as follows: take the dual of the EG program, perform a change of variables, and if you take the dual of this program again, you get the convex program (2). See [Dev09] for details on this connection.

The spending constraint model is an interesting special class of markets that is amenable to efficient algorithms. The spending constraint utilities for the Fisher was introduced by Vazirani [Vaz10b] and extended to the more general Arrow-Debreu model by Devanur and Vazirani [DV04]. Vazirani [Vaz10b] extended the primal-dual algorithm for linear utilities of Devanur et. al. [DPSV08] to the spending constraint utilities. The primal-dual interpretation of the algorithm of Devanur et. al. was based on the Eisenberg-Gale convex program, but no such program is known for spending constraints. Based on this, Vazirani [Vaz10b] conjectured that there must exist a convex program that captures the spending constraint model as well. We finally resolve this conjecture positively in this paper.

The generalized gradient descent method, also known as *mirror-descent* has been well studied in the optimization community, see, e.g., [NY83, BT03]. In the optimization literature however, the kernel function h is usually assumed to be strongly convex. Then with a diminishing stepsize $1/\gamma_t = O(1/\sqrt{t})$, the generalized gradient method (7) converges with rate $O(1/\sqrt{t})$, even if f is non-differentiable (e.g., [BT03, Nem05]). If f is further differentiable and ∇f is Lipschitz-continuous, then (7) converges with rate $O(1/t)$ using a constant γ that is no smaller than the Lipschitz constant. (e.g., [Ber99, Nes04]). In this case, more sophisticated variations of (7) can achieve a convergence rate $O(1/t^2)$; see, e.g., [Nes83, Nes04, Nes05, Tse08]. In this paper, we show that the method (7) can have convergence rate $O(1/t)$ with Condition (8). Condition (8) is implied by the assumptions that ∇f is Lipschitz continuous and h is strongly convex (e.g., [Nes04]) and is hence a weaker assumption. Without these stronger assumptions, Condition (8) often requires a close connection between the functions f and h , which is the case for our analysis of the PR dynamics in Section 4.

Organization: In Section 3 we prove Theorem 3. Theorem 6 is proved in Section 4. In Section 5, we define various measures of distance to equilibrium and show convergence of PR dynamics with respect to these measures. We also compare these in detail with Zhang’s results in Section 5.2. We conclude with directions for future research in Section 6.

3. GRADIENT DESCENT

In this section, we prove convergence properties of the generalized gradient method (7) based on the assumption (8) as stated in Theorem 3. We start by noting some key properties of Bregman divergences:

- $d(x, y) \geq 0$ for all $x, y \in \text{dom } h$. If h is strictly convex, then $d(x, y) = 0$ if and only if $x = y$.
- In general $d(x, y) \neq d(y, x)$, and it does not satisfy the triangle inequality.
- The following *three-point identity* follows directly from definition:

$$d(x, z) = d(x, y) + d(y, z) + \langle \nabla h(y) - \nabla h(z), x - y \rangle. \quad (9)$$

If h is not strictly convex, then the solution to the minimization problem in (7) may not be unique. However, a solution always exists because we assume C is compact.

Recall that in order to prove Theorem 3, it is sufficient to prove Lemmas 5 and 4. We first prove Lemma 5.

PROOF OF LEMMA 5. We have

$$\begin{aligned} f(x(t+1)) &\leq \ell_f(x(t+1); x(t)) + \gamma d(x(t+1), x(t)) \\ &\leq \ell_f(x(t); x(t)) + \gamma d(x(t), x(t)) \\ &= f(x(t)). \end{aligned}$$

The first inequality follows from (8). The second inequality is by definition of $x(t+1)$. The last equality is also by definition, since $d(x, x) = 0$. \square

The key step in the proof of Lemma 4 is the following lemma. Let $x^* := \arg \min_{x \in C} \{f(x)\}$ be an optimum solution to (4).

LEMMA 8. For all t ,

$$f(x(t+1)) - f(x^*) \leq \gamma d(x^*, x(t)) - \gamma d(x^*, x(t+1)).$$

Before we prove Lemma 8, we show how Lemma 4 follows from it.

PROOF OF LEMMA 4. The lemma follows by simply summing over all t the conclusion of Lemma 8 and using the fact that $d(x, y)$ is always greater than 0. \square

In order to prove Lemma 8, we need a result on optimization with Bregman divergence. Consider the following optimization problem:

$$\begin{aligned} &\text{minimize} && g(x) + d(x, y) \\ &\text{subject to} && x \in C \end{aligned} \quad (10)$$

where $g(x)$ is a convex function and C is a compact convex set. The following lemma can be found in, e.g., [CT93].

LEMMA 9. If x^+ is the optimal solution to the optimization problem (10), then

$$g(x) + d(x, y) \geq g(x^+) + d(x^+, y) + d(x, x^+). \quad (11)$$

We apply Lemma 9 with $g(x) = (1/\gamma)\ell_f(x; x(t))$ to obtain the following corollary.

COROLLARY 10. For all $z \in C$,

$$\begin{aligned} &\ell_f(x(t+1); x(t)) + \gamma d(x(t+1), x(t)) \\ &\leq \ell_f(z; x(t)) + \gamma d(z, x(t)) - \gamma d(z, x(t+1)). \end{aligned}$$

Lemma 8 follows as a consequence of this corollary.

PROOF OF LEMMA 8. We have

$$\begin{aligned} f(x(t+1)) &\leq \ell_f(x(t+1); x(t)) + \gamma d(x(t+1), x(t)) \\ &\leq \ell_f(x^*; x(t)) + \gamma d(x^*, x(t)) \\ &\quad - \gamma d(x^*, x(t+1)) \\ &\leq f(x^*) + \gamma d(x^*, x(t)) - \gamma d(x^*, x(t+1)). \end{aligned}$$

The first and third inequalities follow from (8), and the second inequality follows from Corollary 10. \square

4. PR AS GRADIENT DESCENT

In this section, we show the equivalence of the PR dynamics and the gradient descent algorithm on the new convex program, that is, we prove Theorem 6 and Lemma 7.

Recall the function $\varphi(b) = -\sum_{i,j} b_{ij} \log(u_{ij}/p_j)$ and feasible region S defined in the introduction. The components of the gradient of φ are given as

$$(\nabla \varphi(b))_{ij} = 1 - \log\left(\frac{u_{ij}}{p_j}\right).$$

It is clear that $\nabla \varphi$ is not Lipschitz continuous on S . Therefore we need the weaker assumption (8) to show $O(1/t)$ convergence.

The gradient descent algorithm (7) when applied to φ , S , with d being the KL-divergence and with $\gamma = 1$ is

$$b(t+1) := \arg \min_{a \in S} \{\ell_\varphi(b(t); a) + D(a \| b(t))\}. \quad (12)$$

It is well known⁶ that this update rule takes the following form:

$$\begin{aligned} b_{ij}(t) &= \frac{1}{Z'_i(t)} b_{ij}(t) \exp\left((\nabla \varphi(b(t)))_{ij}\right) \\ &= \frac{1}{Z_i(t)} b_{ij}(t) \left(\frac{u_{ij}}{p_j(t)}\right), \end{aligned}$$

where $Z_i(t)$ is chosen such that $\sum_{j=1}^n b_{ij}(t) = B_i$. This shows Theorem 6.

We now prove Lemma 7. The proof mainly relies on the following characterization of φ .

LEMMA 11. For all $a, b \in S$,

$$\varphi(b) = \ell_\varphi(b; a) + D(p(b) \| p(a)).$$

PROOF. We have

$$\begin{aligned} \varphi(b) - \ell_\varphi(b; a) &= \varphi(b) - \varphi(a) - \langle \nabla \varphi(a), b - a \rangle \\ &= -\sum_{i,j} b_{ij} \log \frac{u_{ij}}{p_j(b)} + \sum_{i,j} a_{ij} \log \frac{u_{ij}}{p_j(a)} \\ &\quad - \sum_{i,j} \left(1 - \log \frac{u_{ij}}{p_j(a)}\right) (b_{ij} - a_{ij}) \\ &= \sum_{i,j} b_{ij} \log \frac{p_j(b)}{p_j(a)} = D(p(b) \| p(a)), \end{aligned}$$

where in the third equality we used $\sum_{i,j} b_{ij} = \sum_{i,j} a_{ij} = 1$. \square

PROOF OF LEMMA 7. The lemma follows from Lemma 11 and the fact that $D(p(b) \| p(a)) \leq D(b \| a)$, which we can prove using the convexity of the function $q(x, y) = x \log(x/y)$:

$$\begin{aligned} D(p(b) \| p(a)) &= n \sum_j \left(\frac{1}{n} q(p_j(b), p_j(a)) \right) \\ &= n \sum_j q\left(\frac{1}{n} \sum_i b_{ij}, \frac{1}{n} \sum_i a_{ij}\right) \\ &\leq n \sum_j \frac{1}{n} \sum_i q(b_{ij}, a_{ij}) = D(b \| a). \end{aligned}$$

This finishes the proof. \square

⁶Theorem 12 shows a more general statement.

4.1 Extension to Spending Constraint Utilities

For Fisher markets with spending constraint utilities, let φ be the negative objective function of the convex program (3):

$$\varphi(b) = - \sum_{i,j,l} b_{ij}^l \log(u_{ij}^l) + \sum_j p_j \log(p_j).$$

The components of its gradient is

$$(\nabla \varphi(b))_{i,j,l} = 1 - \log\left(\frac{u_{ij}^l}{p_j}\right).$$

We use the same gradient descent method (12), but with the constraint set

$$S = \left\{ a : \sum_{j,l} a_{ij}^l \leq B_i, \forall i, \text{ and } 0 \leq a_{ij}^l \leq B_{ij}^l, \forall i, j, l \right\}.$$

THEOREM 12. *Each bidding vector $b_i(t+1)$ can be computed separately by using the following iterative PR and capping algorithm.*

Algorithm: Iterative PR and Capping

input: B_i and B_{ij}^l , u_{ij}^l , and $b_{ij}^l(t)$ for all j and l .

initialize: let A_i be set of all pairs (j, l) , and $C_i = \emptyset$.

repeat:

1. Let $A_i = A_i \setminus C_i$, $\bar{B}_i = B_i - \sum_{(j,l) \in C_i} B_{ij}^l$, and $Q_i = \sum_{(j,l) \in A_i} (u_{ij}^l / p_j(t)) b_{ij}^l(t)$.

2. Let

$$b_{ij}^l(t+1) = \begin{cases} \frac{(u_{ij}^l / p_j(t)) b_{ij}^l(t)}{Q_i} \bar{B}_i & \text{if } (j, l) \in A_i, \\ B_{ij}^l & \text{otherwise.} \end{cases}$$

3. Let $C_i = \{(j, l) \in A_i : b_{ij}^l(t+1) > B_{ij}^l\}$. If $C_i = \emptyset$, stop and return $b_i(t+1)$; otherwise continue.

PROOF. The minimization problem (12) can be solved separately for each agent i ,

$$b_i(t+1) = \arg \min_{\substack{\sum_{j,l} a_{ij}^l \leq B_i, \\ 0 \leq a_{ij}^l \leq B_{ij}^l, \forall j, l}} \left\{ \sum_{j,l} a_{ij}^l (\nabla \varphi(b(t)))_{i,j,l} + \sum_{j,l} a_{ij}^l \log\left(\frac{a_{ij}^l}{b_{ij}^l(t)}\right) \right\}.$$

To solve this problem, we introduce Lagrange multipliers λ_i and μ_{ij}^l . The Lagrangian is

$$\begin{aligned} L_i(a_i, \lambda_i, \mu_i) &= \sum_{j,l} a_{ij}^l (\nabla \varphi(b(t)))_{i,j,l} + \sum_{j,l} a_{ij}^l \log\left(\frac{a_{ij}^l}{b_{ij}^l(t)}\right) \\ &\quad + \lambda_i \left(\sum_{j,l} a_{ij}^l - B_i \right) + \sum_{j,l} \mu_{ij}^l (a_{ij}^l - B_{ij}^l) \\ &= \sum_{j,l} a_{ij}^l \left((\nabla \varphi(b(t)))_{i,j,l} + \lambda_i + \mu_{ij}^l \right) \\ &\quad + \sum_{j,l} a_{ij}^l \log\left(\frac{a_{ij}^l}{b_{ij}^l(t)}\right) - \lambda_i B_i - \sum_{j,l} \mu_{ij}^l B_{ij}^l. \end{aligned}$$

Taking derivatives with respect to a_{ij}^l and setting them to be zero, i.e., let $(\nabla L_i(a_i, \lambda_i, \mu_i))_{jl} = 0$, we have

$$(\nabla \varphi(b(t)))_{i,j,l} + \lambda_i + \mu_{ij}^l + \log\left(\frac{a_{ij}^l}{b_{ij}^l(t)}\right) + 1 = 0.$$

The solution is

$$a_{ij}^l = \frac{1}{\exp(2 + \lambda_i + \mu_{ij}^l)} \frac{u_{ij}^l b_{ij}^l(t)}{p_j(t)}.$$

Using the optimality conditions (KKT conditions),

$$\mu_{ij}^l (a_{ij}^l - B_{ij}^l) = 0, \quad \forall i, j, l,$$

we can find the optimal solution $b_i(t+1)$ using an iterative PR and capping algorithm. \square

5. CONVERGENCE PROPERTIES

The PR dynamics may never get to an exact equilibrium. Therefore we consider approximate equilibria, based on a measure of how close a given solution is to equilibrium. There are different ways to measure the “distance” from equilibria. We show that all of these can be related to $\varphi(b(t)) - \varphi(b^*)$, the difference in the objective function between the current point and the optimum. (The additive difference is a reasonable measure since we normalized the sum of all budgets to 1.) Thus, any convergence bound we show with respect to the objective function can easily be translated to one for the particular measure of approximation.

The PR dynamics is defined for any initial bid vector b^0 , and the convergence time is proportional to $D(b^* || b(0))$ (Theorem 3). For the sake of concreteness, we state our bounds assuming that $b_{ij}(0) = B^i/m$, that is each buyer initially divides his budget equally between all the items. With this, we have the following bound on $D(b^* || b(0))$.

LEMMA 13. *If $b_{ij}(0) = B_i/m$ for all i and j , then $D(b^* || b(0)) \leq \log(mn)$.*

PROOF. We have

$$\begin{aligned} D(b^* || b(0)) &= \sum_{i,j} b_{ij}^* \log \frac{b_{ij}^*}{b(0)_{ij}} = \sum_{i,j} b_{ij}^* \log \frac{mb_{ij}^*}{B_i} \\ &= \log m + \sum_{i,j} b_{ij}^* \log b_{ij}^* - \sum_{i,j} b_{ij}^* \log B_i \\ &= \log m - \sum_i B_i \log B_i + \sum_{i,j} b_{ij}^* \log b_{ij}^* \\ &\leq \log m - \sum_i B_i \log B_i \\ &\leq \log m + \log n. \end{aligned}$$

\square

We now have the following convergence results.

$$\varphi(b(t)) - \varphi(b^*) \leq \frac{\log(mn)}{t}. \quad (13)$$

$$\sum_t (\varphi(b(t)) - \varphi(b^*)) \leq \log(mn). \quad (14)$$

$$\varphi(b(t+1)) \leq \varphi(b(t)). \quad (15)$$

In other words, if $t \geq \frac{\log(mn)}{\epsilon}$ then $\varphi(b(t)) - \varphi(b^*) \leq \epsilon$.

5.1 Measures of Approximation

As mentioned before, we now show how to translate the convergence results mentioned above to other ways of measuring the distance from an equilibrium. Recall that we assumed without loss of generality, that the input has been scaled so that there is exactly one unit of each item and $\sum_i B_i = 1$. We may further assume that $\sum_j u_{ij} = 1$ for all i . Let $u_{\min} = \min_{i,j: u_{ij} > 0} u_{ij}$ and $B_{\min} = \min_i B_i$.

- The KL-divergence is a well known (asymmetric) measure of dissimilarity between probability distributions, defined as $D(x \parallel y) := \sum_i x_i \log(x_i/y_i)$. Since we normalized the total budget to be 1, the sum of the prices is also 1. In other words, one can think of the price vectors as probability distributions and use KL-divergence to measure the distance between them. Also it is well known that $\frac{1}{2} \|x - y\|_1^2 \leq D(x \parallel y)$. We show that

$$D(p(t) \parallel p^*) \leq \varphi(b(t)) - \varphi(b^*) . \quad (16)$$

- Let ψ be the objective function of the Eisenberg-Gale convex program for $x_{ij} = b_{ij}/p_j$. We can measure $\psi(b(t)) - \psi(b^*) = \sum_i B_i \log(u_i^*/u_i(t))$, where u_i^* is the utility of buyer i at equilibrium. We show that

$$\psi(b(t)) - \psi(b^*) \leq \varphi(b(t)) - \varphi(b^*) . \quad (17)$$

- We can measure a relative notion of distance from the equilibrium price vector: $\eta := \max_j \left| \frac{p_j(t) - p_j^*}{p_j^*} \right|$. We show that

$$\eta^2 \leq O\left(\frac{n}{u_{\min}}\right) (\varphi(b(t)) - \varphi(b^*)) . \quad (18)$$

- We can measure the maximum sub-optimality of the allocations. Let $u_i(t)$ be the utility of buyer i given his current allocation and \tilde{u}_i^t be the maximum utility buyer i could have obtained if the prices were set to $p(t)$. We show that there exists a ζ such that $u_i(t) \geq (1 - \zeta)\tilde{u}_i(t)$ for all i , and

$$\zeta^2 \leq O\left(\frac{n}{u_{\min} B_{\min}^2}\right) (\varphi(b(t)) - \varphi(b^*)) . \quad (19)$$

- We can also measure an aggregate notion of suboptimality. To this end, let the vector δ be such that $u_i(t) = (1 - \delta_i)\tilde{u}_i(t)$ for all i . Then we measure $\xi := \sum_i \delta_i B_i$. We show that

$$\xi^2 \leq O\left(\frac{n}{u_{\min}}\right) (\varphi(b(t)) - \varphi(b^*)) . \quad (20)$$

5.2 Comparison with Zhang's results

We now rephrase the main result in [Zha09] for ease of comparison. They follow the same framework as outlined earlier, first show convergence of a potential function and then relate that potential function to other measures of approximation. Whereas we show convergence of $\varphi(b(t)) - \varphi(b^*)$, Zhang does the same for the quantity $\psi(b(t+1)) - \psi(b^*) + D(p^* \parallel p(t))$, where $\psi(b)$ is the value of the objective function of the Eisenberg-Gale convex program with $x_{ij} = b_{ij}/p_j$:

$$\psi(b) = - \sum_i B_i \log\left(\sum_j \frac{u_{ij} b_{ij}}{p_j}\right).$$

Zhang's main result can be summarized by the following lemmas. Let $p_{\min}^* = \min_j p_j^*$.

LEMMA 14 ([ZHA09]).

$$\sum_t \left(\psi(b(t+1)) - \psi(b^*) + D(p^* \parallel p(t)) \right) \leq D(b^* \parallel b(0)) .$$

PROOF IN OUR FRAMEWORK. First by using Corollary 10, we have

$$\begin{aligned} \varphi(b(t+1)) &\leq \ell_\varphi(b(t+1); b(t)) + D(b(t+1) \parallel b(t)) \\ &\leq \ell_\varphi(b^*; b(t)) + D(b^* \parallel b(t)) - D(b^* \parallel b(t+1)) . \end{aligned}$$

Next using Lemma 11, i.e., $\varphi(b^*) = \ell_\varphi(b^*; b(t)) + D(p^* \parallel p(t))$, we get

$$\begin{aligned} \varphi(b(t+1)) - \varphi(b^*) + D(p^* \parallel p(t)) \\ \leq D(b^* \parallel b(t)) - D(b^* \parallel b(t+1)) . \end{aligned}$$

Summing the above inequality over all t and noticing $D(b^* \parallel b(t+1)) \geq 0$ lead to

$$\sum_t \left(\varphi(b(t+1)) - \varphi(b^*) + D(p^* \parallel p(t)) \right) \leq D(b^* \parallel b(0)) .$$

Further using Lemma 19 finishes the proof. \square

LEMMA 15 ([ZHA09]). Let $\eta := \max_j \left| \frac{p_j(t) - p_j^*}{p_j^*} \right|$ and assume $\eta \leq 1$. Then $\eta^2 \leq 16 \frac{D(p(t) \parallel p^*)}{p_{\min}^*}$.

LEMMA 16 ([ZHA09]). For all i , let $u_i(t)$ be buyer i 's utility at time t , and let $\tilde{u}_i(t)$ be the maximum utility he can achieve given the prices $p(t)$. Then there exists a ζ such that $u_i(t) \geq (1 - \zeta)\tilde{u}_i(t)$ for all i and

$$\zeta^2 \leq \frac{256}{p_{\min}^* B_{\min}^2} \left(\psi(b(t)) - \psi(b^*) + D(p(t) \parallel p^*) \right) .$$

LEMMA 17 ([ZHA09]). $p_{\min}^* \geq u_{\min} B_{\min}/m$.

However, it is not known if $\psi(b(t+1)) - \psi(b^*) + D(p^* \parallel p(t))$ is monotonically decreasing. So you simply obtain that for all τ , there exists a $t \leq \tau$ such that

$$\psi(b(t)) - \psi(b^*) + D(p^* \parallel p(t)) \leq \frac{D(b^* \parallel b(0))}{\tau} . \quad (21)$$

instead of obtaining that for all t

$$\psi(b(t)) - \psi(b^*) + D(p^* \parallel p(t)) \leq \frac{D(b^* \parallel b(0))}{t} .$$

In fact, we can derive all of Zhang's results in the stronger form by bounding each of $\psi(b(t)) - \psi(b^*)$ and $D(p(t) \parallel p^*)$ separately by $\varphi(b(t)) - \varphi(b^*)$ (and the fact that φ is monotonically decreasing).

LEMMA 18. For all t , $D(p(t) \parallel p^*) \leq \varphi(b(t)) - \varphi(b^*)$.

PROOF. By Lemma 11,

$$\begin{aligned} \varphi(b(t)) &= \ell_\varphi(b(t); b^*) + D(p(t) \parallel p^*) \\ &= \varphi(b^*) + \langle \nabla \varphi(b^*), b(t) - b^* \rangle + D(p(t) \parallel p^*) \end{aligned}$$

By optimality, $\langle \nabla \varphi(b^*), b(t) - b^* \rangle \geq 0$, so the lemma follows. \square

LEMMA 19. For all $b \in S$, $\psi(b) - \psi(b^*) \leq \varphi(b) - \varphi(b^*)$.

PROOF. It suffices to show

$$\psi(b) \leq \varphi(b) - \sum_i B_i \log B_i$$

and

$$\psi(b^*) = \varphi(b^*) - \sum_i B_i \log B_i .$$

For the inequality above, we use convexity of $-\log$:

$$\begin{aligned}
\psi(b) &= -\sum_i B_i \log \left(\sum_j u_{ij} \frac{b_{ij}}{p_j} \right) \\
&= -\sum_i B_i \log \left(\sum_j \frac{u_{ij}}{p_j} \frac{b_{ij}}{B_i} \right) - \sum_i B_i \log B_i \\
&\leq -\sum_i B_i \sum_j \frac{b_{ij}}{B_i} \log \frac{u_{ij}}{p_j} - \sum_i B_i \log B_i \\
&= -\sum_{i,j} b_{ij} \log \frac{u_{ij}}{p_j} - \sum_i B_i \log B_i.
\end{aligned}$$

For the equality at b^* , we use the optimality conditions:

$$\begin{aligned}
\psi(b^*) &= -\sum_i B_i \log \left(\sum_j \frac{u_{ij}}{p_j^*} \frac{b_{ij}^*}{B_i} \right) - \sum_i B_i \log B_i \\
&= -\sum_i B_i \log \left(\max_j \frac{u_{ij}}{p_j^*} \right) - \sum_i B_i \log B_i \\
&= -\sum_{i,j} b_{ij}^* \log \frac{u_{ij}}{p_j^*} - \sum_i B_i \log B_i.
\end{aligned}$$

□

We also have the following bound on the total deviation of the *symmetric* KL-divergence between p^* and $p(t)$.

LEMMA 20.

$$\begin{aligned}
&\sum_t \left(D(p^* \| p(t)) + D(p(t) \| p^*) \right) \\
&\leq D(b^* \| b(0)) + (\varphi(b(0)) - \varphi(b^*)).
\end{aligned}$$

PROOF. First using Corollary 10, we have

$$\begin{aligned}
\varphi(b(t+1)) &\leq \ell_\varphi(b(t+1); b(t)) + D(b(t+1) \| b(t)) \\
&\leq \ell_\varphi(b^*; b(t)) + D(b^* \| b(t)) - D(b^* \| b(t+1)).
\end{aligned}$$

Plugging in $\ell_\varphi(b^*; b(t)) = \varphi(b(t)) + \langle \nabla \varphi(b(t)), b^* - b(t) \rangle$ gives

$$\begin{aligned}
\langle \nabla \varphi(b(t)), b(t) - b^* \rangle &\leq D(b^* \| b(t)) - D(b^* \| b(t+1)) \\
&\quad + \varphi(b(t)) - \varphi(b(t+1)).
\end{aligned}$$

The optimality of b^* implies $\langle \nabla \varphi(b^*), b(t) - b^* \rangle \geq 0$. Subtracting $\langle \nabla \varphi(b^*), b(t) - b^* \rangle$ from the left hand gives

$$\begin{aligned}
&\langle \nabla \varphi(b(t)) - \nabla \varphi(b^*), b(t) - b^* \rangle \\
&\leq D(b^* \| b(t)) - D(b^* \| b(t+1)) + \varphi(b(t)) - \varphi(b(t+1)).
\end{aligned}$$

Using $\nabla \varphi(b) = 1 - \log(u_{ij}/p_j)$, we have

$$\begin{aligned}
&\langle \nabla \varphi(b(t)) - \nabla \varphi(b^*), b(t) - b^* \rangle \\
&= \sum_{i,j} (b_{ij}(t) - b_{ij}^*) \log \frac{p_j(t)}{p_j^*} \\
&= \sum_j (p_j(t) - p_j^*) \log \frac{p_j(t)}{p_j^*} \\
&= D(p(t) \| p^*) + D(p^* \| p(t)).
\end{aligned}$$

Therefore,

$$\begin{aligned}
&D(p(t) \| p^*) + D(p^* \| p(t)) \\
&\leq D(b^* \| b(t)) - D(b^* \| b(t+1)) + \varphi(b(t)) - \varphi(b(t+1)).
\end{aligned}$$

Summing over all t gives the desired result. □

5.3 Improved bounds

The bounds mentioned in Section 5.1 (in particular, (18) and (19)) are stronger than the corresponding ones in [Zha09]. These are obtained by using Lemmas 15 and 16 with a better bound for p_{\min}^* , which is as in the following lemma.

LEMMA 21. $p_{\min}^* \geq u_{\min}/n$.

PROOF. Let b^* be an equilibrium bid vector, and let $\beta_i^* = \max_j u_{ij}/p_j^*$. Suppose for contradiction that there exists an item \tilde{j} such that $p_{\tilde{j}}^* < u_{\min}/n$. Then for all bidders i ,

$$\beta_i^* \geq \frac{u_{i\tilde{j}}}{p_{\tilde{j}}^*} \geq \frac{u_{\min}}{p_{\tilde{j}}^*} > n.$$

By the buyer optimality condition, $b_{ij}^* > 0$ implies $u_{ij}/p_j^* = \beta_i^*$. Combining this with the above yields the condition

$$b_{ij}^* > 0 \implies p_j^* < \frac{u_{ij}}{n}.$$

We now obtain the contradiction

$$\begin{aligned}
1 &= \sum_{i,j : b_{ij}^* > 0} b_{ij}^* \leq \sum_{i,j : b_{ij}^* > 0} p_j^* \\
&< \sum_{i,j : b_{ij}^* > 0} \frac{u_{ij}}{n} \leq \frac{1}{n} \sum_{i,j} u_{i,j} = 1.
\end{aligned}$$

Hence, we conclude that there can be no such item \tilde{j} , and therefore $p_{\min}^* \geq u_{\min}/n$. □

Finally, we also give a bound on the aggregate notion of suboptimality of allocations (20) that is new.

LEMMA 22. Suppose $\varphi(b(t)) - \varphi(b^*) \leq u_{\min}^2/(8n^2)$. Let the vector δ be such that

$$u_i(t) = (1 - \delta_i) \tilde{u}_i(t),$$

where $u_i(t) = \sum_j u_{ij} b_{ij}(t)/p_j(t)$ is the utility of buyer i at time t and $\tilde{u}_i(t) = B_i \max_j (u_{ij}/p_j(t)) = B_i \beta_i(t)$ is the maximum utility that buyer i could obtain given the prices $p(t)$. Then

$$\left(\sum_i B_i \delta_i \right)^2 = O\left(\frac{n}{u_{\min}} \right) (\varphi(b(t)) - \varphi(b^*)).$$

PROOF. To ease notation, we drop the argument t . For all i, j , let $\alpha_i = \log \beta_i = \max_j \log(u_{ij}/p_j)$, $\varepsilon_{ij} = \alpha_i - \log(u_{ij}/p_j)$, and $\varepsilon_i = \sum_j b_{ij} \varepsilon_{ij}/B_i$. Observe

$$\begin{aligned}
u_i &= \sum_j b_{ij} \frac{u_{ij}}{p_j} = \sum_j b_{ij} \beta_i \exp(-\varepsilon_{ij}) \\
&\geq \sum_j b_{ij} \beta_i (1 - \varepsilon_{ij}) \\
&= B_i \beta_i - \beta_i \sum_j b_{ij} \varepsilon_{ij} \\
&= B_i \beta_i - B_i \beta_i \varepsilon_i = (1 - \varepsilon_i) \tilde{u}_i',
\end{aligned}$$

and hence $\delta_i \leq \varepsilon_i$ for all i . Therefore,

$$\sum_i B_i \delta_i \leq \sum_i B_i \varepsilon_i = \sum_{i,j} b_{ij} \varepsilon_{ij},$$

and it suffices to bound $\sum_{i,j} b_{ij} \varepsilon_{ij}$. In particular, we show that $\sum_{i,j} b_{ij} \varepsilon_{ij} \leq O(\eta) + \varepsilon$. The proof then follows immediately

from (18). To do this, we start by rewriting the objective function $\varphi(b)$ as

$$\begin{aligned}\varphi(b) &= -\sum_{i,j} b_{ij} \left(\log \frac{u_{ij}}{p_j} - \alpha_i + \alpha_i \right) \\ &= \sum_{i,j} b_{ij} \varepsilon_{ij} - \sum_{i,j} b_{ij} \alpha_i \\ &= \sum_{i,j} b_{ij} \varepsilon_{ij} - \sum_i B_i \alpha_i.\end{aligned}$$

Let $\alpha_i^* = \max_j \log(u_{ij}/p_j^*)$. The optimality condition states that $\varphi(b^*) = -\sum_i B_i \alpha_i^*$. Let $\varepsilon = \varphi(b) - \varphi(b^*)$. Then

$$\sum_{i,j} b_{ij} \varepsilon_{ij} = \sum_i B_i (\alpha_i - \alpha_i^*) + \varepsilon.$$

Since $\sum_i B_i = 1$, we have $\sum_i B_i (\alpha_i - \alpha_i^*) \leq \max_i (\alpha_i - \alpha_i^*)$. Furthermore, for all i ,

$$\begin{aligned}\alpha_i - \alpha_i^* &= \max_j \log \frac{u_{ij}}{p_j} - \max_j \log \frac{u_{ij}}{p_j^*} \\ &= \max_j \left(\log \frac{u_{ij}}{p_j^*} + \log \frac{p_j^*}{p_j} \right) - \max_j \log \frac{u_{ij}}{p_j^*} \\ &\leq \max_j \log \frac{u_{ij}}{p_j^*} + \max_j \log \frac{p_j^*}{p_j} - \max_j \log \frac{u_{ij}}{p_j^*} \\ &= \max_j \log \frac{p_j^*}{p_j} \leq O(\eta).\end{aligned}$$

□

6. CONCLUSION AND FUTURE WORK

Our work opens up many interesting directions for future research.

- Eisenberg-Gale-type convex programs also capture equilibria for other classes of markets, especially for many resource allocation markets [JV10] and markets with Leontief utilities. It is not clear if there are analogs of convex program (2) for these markets. Convex program (2) seems to capture a different class of markets since it generalizes to the spending constraint utilities, for which the EG program does not generalize to. Such programs might lead to interesting algorithms and dynamics for these other markets as well. Also such programs have been used to design distributed algorithms to achieve proportional fairness. It is conceivable that our new convex program and/or the PR dynamics could get better results for the same.
- Can we get a better convergence bound for PR dynamics? While it seems difficult to get a convergence rate better than $O(1/\epsilon)$ w.r.t φ , it is possible one could get faster rates w.r.t other notions of approximate equilibria that we discuss.
- Can we design other similar dynamics that converge faster? We note that under the stronger condition of Lipschitz continuity of gradients, there are gradient algorithms [Nes83, Nes05, Tse08] that converge in time $O(1/\sqrt{\epsilon})$ (as opposed to the one we use which converges in time $O(1/\epsilon)$). However, these algorithms do not seem to work with the weaker assumption that we have. It would be very interesting (from a more general convex programming point of view) to get an $O(1/\sqrt{\epsilon})$ time algorithm with our weaker assumption.

- There is a natural asynchronous version of the PR dynamics: in each iteration, the bids of a single buyer and the corresponding prices are updated, with some non-saturation condition saying that every buyer updates his bids frequently enough. There are also randomized versions where a buyer wakes up with a Poisson clock and updates his bids and the corresponding prices. We believe our techniques (perhaps combined with some of the techniques from the book by Bertsekas and Tsitsiklis [BT97]) could be extended to handle these cases as well.

7. REFERENCES

- [AAE⁺08] Baruch Awerbuch, Yossi Azar, Amir Epstein, Vahab S. Mirrokni, and Alexander Skopalik. Fast convergence to nearly optimal solutions in potential games. In *ACM Conference on Electronic Commerce*, pages 264–273, 2008.
- [ABH59] K. Arrow, H. Block, and L. Hurwicz. On the stability of the competitive equilibrium: II. *Econometrica*, 27(1):82–109, 1959.
- [AGM⁺08] Heiner Ackermann, Paul W. Goldberg, Vahab S. Mirrokni, Heiko Röglin, and Berthold Vöcking. Uncoordinated two-sided matching markets. In *ACM Conference on Electronic Commerce*, pages 256–263, 2008.
- [Ber99] D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 2nd edition, 1999.
- [Bre67] L. Bregman. The relaxation method of finding the common points of convex sets and its application to the solution of problems in convex programming. *U.S.S.R. Comput. Math. and Math. Phys.*, 7:200–217, 1967.
- [BT97] Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1997.
- [BT03] A. Beck and M. Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31:167–175, 2003.
- [CD06] Xi Chen and Xiaotie Deng. Settling the complexity of two-player Nash equilibrium. In *FOCS*, pages 261–272, 2006.
- [CDDT09] Xi Chen, Decheng Dai, Ye Du, and Shang-Hua Teng. Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities. In *FOCS*, 2009.
- [CF08] Richard Cole and Lisa Fleischer. Fast-converging tatonnement algorithms for one-time and ongoing market problems. In *STOC*, pages 315–324, 2008.
- [CMV05] Bruno Codenotti, Benton McCune, and Kasturi R. Varadarajan. Market equilibrium via the excess demand function. In *STOC*, pages 74–83, 2005.
- [CPV05] Bruno Codenotti, Sriram V. Pemmaraju, and Kasturi R. Varadarajan. On the polynomial time computation of equilibria for certain exchange economies. In *SODA*, pages 72–81, 2005.
- [CS07] Steve Chien and Alistair Sinclair. Convergence to approximate Nash equilibria in congestion games. In *SODA*, pages 169–178, 2007.
- [CSVY06] Bruno Codenotti, Amin Saberi, Kasturi R. Varadarajan, and Yinyu Ye. Leontief economies encode nonzero sum two-player games. In *SODA*, pages 659–667, 2006.
- [CT93] G. Chen and M. Teboulle. Convergence analysis of a proximal-like minimization algorithm using Bregman functions. *SIAM Journal on Optimization*, 3(3):538–543, August 1993.
- [CT09] Xi Chen and Shang-Hua Teng. Spending is not easier than trading: On the computational equivalence of Fisher and Arrow-Debreu equilibria. *CoRR*, abs/0907.4130, 2009.
- [CV07] Bruno Codenotti and Kasturi Varadarajan. *Algorithmic Game Theory*, chapter Computation of Market Equilibria by Convex Programming. Cambridge University Press, 2007.
- [Dev09] Nikhil R. Devanur. Fisher markets and convex programs. Working paper available from

- <http://research.microsoft.com/en-us/um/people/vazirani/>, 2009.
- [DGP06] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. In *STOC*, pages 71–78, 2006.
- [DK08] Nikhil R. Devanur and Ravi Kannan. Market equilibria in polynomial time for fixed number of goods or agents. In *FOCS*, pages 45–53, 2008.
- [DPS02] Xiaotie Deng, Christos H. Papadimitriou, and Shmuel Safra. On the complexity of equilibria. In *STOC*, pages 67–71, 2002.
- [DPSV08] Nikhil R. Devanur, Christos H. Papadimitriou, Amin Saberi, and Vijay V. Vazirani. Market equilibrium via a primal–dual algorithm for a convex program. *J. ACM*, 55(5), 2008.
- [DV04] Nikhil R. Devanur and Vijay V. Vazirani. The spending constraint model for market equilibrium: algorithmic, existence and uniqueness results. In *STOC*, pages 519–528, 2004.
- [EG59] Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *Ann. Math. Statist.*, 30(1):165–168, 1959.
- [FGK⁺08] Lisa Fleischer, Rahul Garg, Sanjiv Kapoor, Rohit Khandekar, and Amin Saberi. A fast and simple algorithm for computing market equilibria. In *WINE*, pages 19–30, 2008.
- [GK06] Rahul Garg and Sanjiv Kapoor. Auction algorithms for market equilibrium. *Math. Oper. Res.*, 31(4):714–729, 2006.
- [Jai07] Kamal Jain. A polynomial time algorithm for computing an Arrow-Debreu market equilibrium for linear utilities. *SIAM J. Comput.*, 37(1):303–318, 2007.
- [JV10] Kamal Jain and Vijay V. Vazirani. Eisenberg-Gale markets: algorithms and game-theoretic properties. *Games and Economic Behavior*, 70(1), 2010.
- [Kel97] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.
- [KMT98] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [KV] F. P. Kelly and V. V. Vazirani. Rate control as a market equilibrium. Manuscript, 2002.
- [LLSB08] Dave Levin, Katrina LaCurts, Neil Spring, and Bobby Bhattacharjee. Bittorrent is an auction: analyzing and improving bittorrent’s incentives. In *SIGCOMM ’08: Proceedings of the ACM SIGCOMM 2008 conference on Data communication*, pages 243–254, New York, NY, USA, 2008. ACM.
- [Nem05] A. Nemirovski. Lectures on modern convex optimization. Lecture notes, Georgia Institute of Technology, 2005.
- [Nes83] Yu. Nesterov. A method for unconstrained convex minimization problem with the rate of convergence $o(1/k^2)$. *Soviet Math. Docl. (Doklady AN SSSR)*, 269:543–547, 1983.
- [Nes04] Yu. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer, Boston, 2004.
- [Nes05] Yu. Nesterov. Smooth minimization of nonsmooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- [NY83] A. Nemirovsky and D. Yudin. *Problem Complexity and Method Efficiency in Optimization*. J. Wiley & Sons, New York, 1983.
- [Roc70] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [Ros73] R.W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [Shm09] V. Shmyrev. An algorithm for finding equilibrium in the linear exchange model with fixed budgets. *Journal of Applied and Industrial Mathematics*, 3:505–518, 2009.
- [Tse08] P. Tseng. On accelerated proximal gradient methods for convex-concave optimization. Submitted to *SIAM Journal on Optimization*, 2008.
- [Uza60] H. Uzawa. Walras’ tatonnement in the theory of exchange. *Review of Economic Studies*, 27:182–94, 1960.
- [Vaz07] Vijay V. Vazirani. *Algorithmic Game Theory*, chapter Combinatorial Algorithms for Market Equilibria. Cambridge University Press, 2007.
- [Vaz10a] Vijay V. Vazirani. Rational convex programs, their feasibility, and the arrow-debreu nash bargaining game. Manuscript available from the author’s homepage, 2010.
- [Vaz10b] Vijay V. Vazirani. Spending constraint utilities, with applications to the adwords market. *Math. of Operations Research*, 35(2), 2010. A manuscript has been publicly available since 2003.
- [VY10] Vijay V. Vazirani and Mihalis Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. In *ICS*, 2010.
- [Wal74] 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.).
- [WZ07] Fang Wu and Li Zhang. Proportional response dynamics leads to market equilibrium. In *STOC ’07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 354–363, New York, NY, USA, 2007. ACM.
- [Ye08] Y. Ye. A path to the Arrow-Debreu competitive market equilibrium. *Mathematical Programming*, 111:315–348, 2008.
- [Zha09] Li Zhang. Proportional response dynamics in the Fisher market. In *ICALP ’09: Proceedings of the 36th International Colloquium on Automata, Languages and Programming*, pages 583–594, Berlin, Heidelberg, 2009. Springer-Verlag.