

Convergence of Incentive-Driven Dynamics in Fisher Markets

Krishnamurthy Dvijotham* Yuval Rabani† Leonard J. Schulman‡

December 10, 2020

Abstract

We study out-of-equilibrium price dynamics in Fisher markets. We develop a general framework in which sellers have (a) a set of *atomic price update* rules (APU), which are simple responses to a price vector; (b) a *belief-formation procedure* that simulates actions of other sellers (themselves using the APU) to some finite horizon in the future. Sellers use an APU to respond to a price vector they generate with the belief formation procedure. The framework allows sellers to have inconsistent and time-varying beliefs about each other. Under mild and natural assumptions on the APU, we show that despite the inconsistent and time-varying nature of beliefs, the market converges to a unique equilibrium at a linear rate (distance to equilibrium decreases exponentially in time). If the APU are driven by weak-gross substitutes demands, the equilibrium point is the same as predicted by those demands.

Keywords: Fisher markets; disequilibrium; market dynamics; bounded rationality; best response; level k model.

1 Introduction

Motivation This paper deals with the question of why, and whether, a model of interacting *strategic* agents converges to equilibrium. We study this question in Fisher markets.

Over the years and in particular recently, several game and market dynamics have been studied, but they fall short of answering our question:

First, in game theory, dynamics are studied in the context of repeated games. Extensive form solution concepts such as subgame perfect or sequential equilibria assume that the agents unravel the entire evolution of the game and choose in advance their entire play optimally. This is likely to be computationally infeasible (e.g. [BCI⁺10], but see in contrast [HPS14]). The strategies are unrealistically prescient of the distant future, hindering on-the-fly adaptation to unexpected changes, and contradicting experience. Just as importantly, since the entire play is determined a-priori and an equilibrium is played throughout, such concepts do not capture out-of-equilibrium behavior (that may lead to equilibrium over time), so they are in fact a static notion.

Second, an alternative approach are game theoretic *regret-matching* dynamics [HMC12]. This is a genuinely dynamic, out-of-equilibrium framework, that can be shown to converge to an attractive solution concept (correlated equilibrium). However, in general the convergence is in the ergodic sense; the next-step strategies of the agents do not necessarily converge to equilibrium play of the game. The strategic justification for this is the ergodic guarantee of the dynamic and not the short-term future utility. In fact, these dynamics require agents to have perfect recollection of the past, and to solve the stage game optimally against past behavior of the other agents.

*Center for Mathematics of Information, California Institute of Technology, Engineering and Applied Science MC305-16, Pasadena CA 91125, USA, Engineering and Applied Science MC305-16, Pasadena CA 91125, USA, dvij@cs.washington.edu.

†The Rachel and Selim Benin School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 9190401, Israel, yraabani@cs.huji.ac.il.

‡California Institute of Technology, Engineering and Applied Science MC305-16, Pasadena CA 91125, USA, schulman@caltech.edu.

Third, in general equilibrium market models (which is the focus of our interest in this paper), Walrasian *tâtonnement*, as well as proportional response, are myopic regret minimizing strategies which can be shown in some cases to converge (non-ergodically) to a market equilibrium. However, the reactions of the agents have to be damped or calculated carefully for a desirable outcome to materialize; such reactions lack strategic justification.

We aim to circumvent some of the limitations of these approaches with a bounded-rationality-based approach. Various formulations of bounded rationality have provided a rich basis for progress in game theory and, over the last two decades, in its algorithmic aspects. The most basic approach in this vein is *best-response* dynamics (also known as Nash dynamics). Agents play myopically an optimal move at each round, assuming that the other agents will not deviate from their existing strategy. A strategy which is somewhat more sophisticated than best-response is limited-depth exploration of an extensive form game tree. This is an approach to complex games that was developed in the early days of AI (the exploration depth is sometimes called the *ply* of a search). Essentially the same concept is known in control theory as *receding-horizon control*. This is in contrast with the full-rationality approach underlying solution concepts such as the aforementioned subgame perfect or sequential equilibria. The idea that people compete in games by pursuing limited look-ahead situational analysis goes under the rubric of the *level k model*, initiated by [SW94, SW95] and [Nag95]; related ideas are also known as *cognitive hierarchy*, *higher-order rationality*, and *bounded depth of reasoning*. The idea has been subjected to many experimental tests—see [HCW98, CGCB01, Cra03, CHC04, CGC06, CI07b, CI07a]—and has emerged with considerable support. For recent theoretical work on the model see [Str14, Kne15, dCSS14, Gor18]; for a survey see [CCGI13].

From the same motivations (of bounded-rationality agents and changing conditions), it is important to study the dynamics and stability of *markets* composed of agents each of whom performs some limited look-ahead and plays optimally against that forecast. Limited look-ahead means that each agent j has a mental model of each other agent j' , where j' looks ahead some constant number of steps, and based on that look-ahead choose selfishly-optimal action for j' ; j now chooses an action that is selfishly-optimal conditional on those other imagined actions. Such a model does not make sense for computationally-unbounded agents who possess perfect common knowledge of the properties of the market. But such agents do not exist. Instead, the model is consistent with experience that markets are composed of many agents who, despite having limited ability to predict the actions of others, do their best to make such a prediction and then respond optimally to their own prediction. This is a very different approach to agent choice from the “solution concept” notion on which game theory rests: Nash equilibria, correlated equilibria, the core, and so forth. This approach even side-steps thinking of market equilibrium as a solution concept for general equilibrium theory.

In particular, one difference between the rational-agent approach and this approach to agent choice, is that in contrast with full rationality, in the limited look-ahead case the beliefs of the agents are not necessarily consistent with each other and with reality. From a purely mathematical perspective these inconsistencies might appear to be a fatal flaw. We hold differently, that this is part of the challenge of modeling out-of-equilibrium strategic play. The market is out of equilibrium *because* players do not have perfect models of each other, or because they are uncertain about exogenous factors a few steps into the future. We further hold that the predictive power in experiments of the *level k model* and its variations, is ample reason to study their dynamics. That is what we do here.

Our results This paper is devoted to studying the dynamics and stability of markets where the agents model their peers as using limited look-ahead. We focus on one of the best-understood cases of general equilibrium theory, namely Fisher markets that consist of sellers of goods and buyers endowed with budgets. Only the sellers are assumed to be strategic, and every seller controls the price of a single good.

We develop a fairly general look-ahead framework according to which boundedly-rational sellers can update their prices. The framework (Section 2) assumes a set of “atomic” price updates (APU)—this is a collection of price update rules which represent possible strategies that the sellers use to update prices assuming prices of other goods in the market remain fixed (a specific example includes best-response). Given the APU, we describe a general belief formation procedure, i.e., a procedure with which sellers simulate the state of the market. The simulation proceeds for a certain number of steps, and the players update their

prices using a rule from the APU given the final prices resulting from the simulation. The simulation also includes sellers’ mental models of the strategies and beliefs of other sellers—the overall belief formation procedure is described in terms of a recursion tree. Sellers’ beliefs about each other can be time-varying and inconsistent (i.e seller A’s belief about seller B may not be the same as seller C’s belief about seller B). This includes as a special case, but is considerably more general than, *level k* choices. We refer to dynamics of this sort as *belief-based price update* (abbreviated BBPU) dynamics.

BBPU dynamics, and even the special case of best-response (with no look-ahead), can be quite volatile, as compared with usual tâtonnement processes, because of the absence of any damping factor. Despite the volatility and the potential inconsistency of beliefs, we show that regardless of the specifics of the beliefs formed by the agents, the dynamic converges rapidly to market equilibrium provided the APU satisfies certain properties. More precisely, we analyze two versions of our process. In the synchronous case, all sellers update prices simultaneously. In this case, the distance to equilibrium decays exponentially in the number of steps (a.k.a. *linear convergence*), that is, the distance to equilibrium decreases as θ^t where t is the number of discrete time-steps (instants at which sellers update prices) and $\theta \in (0, 1)$ is a constant that depends on the APU but is independent of the belief formation process. In the asynchronous case, at each time step only a subset of one or more sellers update prices. In this case, the distance to equilibrium decays exponentially in the number of epochs, where an epoch consists of time intervals in which all the sellers update at least once.

Convergence of our dynamics in the limit is implied by work of Milgrom and Roberts [MR91], but any bounded rate of convergence, and definitely linear convergence, was not previously demonstrated even for the simplest version of our process, namely best-response. Our proof of convergence goes through showing that in a judiciously chosen metric, the Thompson metric, the BBPU dynamics form a contraction map.

In Section 3, we consider the concrete case when the APU is given by the best-response update for a given demand function. We shown that under certain assumptions on the demand elasticities, the best-response update satisfies the axioms required for the convergence of the BBPU dynamics. We bound the convergence rate in terms of the elasticity parameters of the demand function.

For the case of CES utilities (constant elasticity of substitution) in the weak gross substitutes (WGS) regime, our analysis of best-response updates appears to obtain a faster convergence rate than obtained in previous work analyzing tâtonnement [CFR10] (for further discussion see Sec. 3.1). In particular, our results show that a price vector at a certain distance from equilibrium can be obtained in a number of updates that is logarithmic in the initial distance to equilibrium (as opposed to the linear bound from Theorem 2 in [CFR10]). We believe that this is because best-response updates are undamped, unlike tâtonnement. However, best-response updates are computationally slightly harder to implement (we need to solve a 1-d nonlinear equation to compute the update, as opposed to the closed-form tâtonnement update).

Related work in markets General equilibrium theory (see [McK02, Muk02]) is largely responsible for the governing paradigm that a state of *equilibrium* which the participants in economic exchange do not wish to deviate from individually is, under mild conditions, attainable, and that this is normally roughly the state of the economy [AD54, McK54] (see also [Hil98]). This is a paradigm that can be observed “in the field” and also reproduced in controlled experiments, and it lends credence and concreteness to the famed *invisible hand* metaphor. However, there is little agreement as to *why* markets tend to reach a state of equilibrium. This is a question about the *stability* or *out-of-equilibrium* behavior of markets. It is important because in reality economic conditions are not static. They vary continually and suffer serious shocks occasionally. So justifying an equilibrium outcome requires a dynamic that moves an economy at disequilibrium back to a new equilibrium, and does so sufficiently quickly that the periods of disequilibrium due to fluctuations are relatively negligible (see [Dix90]). The classical mechanism proposed to explain general equilibrium is Walrasian tâtonnement [Wal74], a process that reacts to excess demand by raising the price and to excess supply by reducing the price. Variants of tâtonnement are known to converge to equilibrium, at least in some classes of markets including those we consider here (Sec. 3.1), see [Sam41, ABH59, CCD13]. However, the classical view of tâtonnement posits the existence of an imaginary “auctioneer” who controls the process by announcing prices. Recent work on the convergence of discrete-time tâtonnement in Fisher markets

attempts to present it as an in-market process in the context of the so-called *ongoing markets* [CF08, CCR12]. However, even this attempt requires a somewhat careful choice of the magnitude of the price adjustment which is not motivated by any agent considerations (aside from a common inexplicable passion to equilibrate the economy). Hence the difficulty is in formulating a theory of out-of-equilibrium behavior that makes sense in terms of the incentives of the participants.

A well-studied special case of our framework is that of Fisher markets where utilities of the buyers have constant elasticities of substitution (CES). It is well-known that market equilibria in Fisher markets with CES utilities can be expressed as solutions to a convex program first proposed by Eisenberg and Gale (see [JV10]). We observe that best-response dynamics (i.e., the simplest example of our setting) can, in fact, be explained as a specific implementation of coordinate descent (in the dual E-G program). The convergence of coordinate descent in E-G was established in [Tse01], without bounds on the rate. Recently, [ST13] established a sublinear convergence rate (the distance to the optimum decays linearly with the number of iterations), if the objective function satisfies some conditions. We note that the objective function of the dual E-G program satisfies these conditions. Our general result shows a linear convergence rate (the distance to equilibrium decays exponentially in the number of iterations), and this holds in particular in the case of best-response. To the best of our knowledge, this is not implied by previous results.

Two recent papers consider market dynamics under strategic behavior. Both bound the fraction of optimal welfare that is guaranteed. In [BLNPL14], strategic buyers play a Nash (or Bayesian) equilibrium in a market in which the sellers' prices are determined by Walrasian tâtonnement; note that here the tâtonnement is part of the *mechanism* defining the game, rather than the agents' strategies. In [BPLS15], sellers engage in best-response dynamics. In this setting the market does not actually have an equilibrium, but a fraction of the optimal welfare can be extracted by the dynamic. In both papers the market model is quite different from ours.

Related work in games In the game theory setting, various kinds of limited-rationality dynamics have been studied in recent years. Our model is related to those studied by Milgrom and Roberts [MR91]. (See also [MR90].) They consider repeated play of a non-cooperative stage game, and define a sequence of strategy sets parametrised by the look-ahead depth k . They define a class of *sophisticated learning dynamics* where players adapt to past behavior of the other players, while also taking into account the union over all finite k of the depth k strategy sets of the other players. In general, and also in the specific applications that they consider, their methods provide only a proof of convergence in the limit, with no bounds on the rate of convergence. One of the three applications that they analyze is continuous time proportional tâtonnement using *lagged price signals* in a somewhat more general market model than ours. Each seller has a fixed distribution over the lag time and the tâtonnement update uses past demand information according to this distribution. Thus, seller behavior and price updates are noticeably different than the ones we consider here. The dynamics that we consider can be defined within the framework of [MR91]. This implies convergence in the limit, but does not provide a bound on the convergence rate.

Relatedly, best-response dynamics in games have been studied extensively, mostly concerning bounds on the quality of the play and conditions that imply or prevent convergence to a Nash equilibrium [MV04, Rou15, FFM12, EFSW13]. The paper [NSVZ11] investigates conditions under which best-response is a fully rational strategy. The situations where best-response is known to lead to an attractive outcome are tightly connected to the concept of *potential games*. See [MS96, AAE⁺08, CS11]. For a damped version, *logit* dynamics, see [AFP⁺16]. For a general discussion of best-response and the related *fictitious play* dynamics, see [SLB09]. It is known that there are natural classes of games, in particular congestion games, where sequential best response dynamics require exponentially many steps to converge [FPT04]. In the stable matching problem, sequential best response dynamics can cycle [AGM⁺11].

2 Belief-based price update dynamics (BBPU)

The market model We consider a Fisher market with n perfectly divisible goods and m buyers. Each good is initially owned by a unique seller who controls its price, and its quantity is scaled to 1. The utility of

that seller is his/her revenue, which is equal to the quantity of the good sold (which is at most the available supply of 1 unit) times the price of the good. The collection of prices of all goods is denoted by a price vector $p = \{p_i\}_{i=1}^n \in \mathbb{R}^n$. For a price vector p , we write $p > 0$ to indicate that all the prices are strictly positive. Similarly for price vectors p, q , we write $p > q$ (resp. $p \geq q$) if $p_j > q_j$ (resp. $p_j \geq q_j$) for all j . We use $[n]$ to denote the set $\{1, \dots, n\}$.

The buyers respond instantly and myopically to price changes. The response is specified as a *demand function*

$$x_{ik}(p) = \text{Demand of buyer } i \text{ for good } k \text{ given the price vector } p$$

The demand functions are in one-one correspondence with *utilities* of the buyers: Suppose that buyer i has a strictly concave utility function $U_i(x_i)$ for a bundle of goods $x_i = \{x_{ik}\}_{k \in [n]}$ and a budget b . Then, the demand function $x_i(p)$ is given by

$$x_i(p) = \operatorname{argmax}_{x_i: \sum_k x_{ik} p_k \leq b} U_i(x_i)$$

where the argmax is unique due to the assumption of strict concavity. The demand for good k at prices p is $\sum_{i=1}^m x_{ik}(p)$. We write the overall demand function as $x(p) = \{x_{ik}(p)\}_{i \in [m], k \in [n]} \in \mathbb{R}^{m \times n}$ (note that this is a matrix-valued function of the vector p indexed by $i = 1, \dots, m, k = 1, \dots, n$).

Throughout the paper, we will use $\mathbf{1}$ to denote a vector with each coordinate equal to 1 (with the dimension clear from context).

Price updates In general, a market dynamic is based on an update rule for each seller that determines its new price. The rules can then be applied synchronously to all sellers, or serially to one seller at a time in some order. We will discuss these variations later. For now, we focus on the update rules. An update rule can take into account some or all of the dynamic history leading to the current state (including the current prices), and also some internal state of the seller that takes other factors into account. We are interested in update rules that depend on the current price vector (and any other parameters), and are *monotone*, *sub-homogeneous*, *price-bounded*, and *positive* with respect to that price vector. We begin by defining a *set of atomic price updates*:

Definition 1. A set of atomic price updates (APU) is a finite collection of mappings $F^k : \mathbb{R}_{++}^n \times [n] \mapsto \mathbb{R}_{++}, k = 1, \dots, l$. It represents the possible rules used by a seller to update prices given current prices of all the sellers.

The following definition characterizes a class of APU's (this is the class for which we will prove convergence/convergence rates of a price update dynamics):

Definition 2. (MONOTONICITY, SUB-HOMOGENEITY, PRICE-BOUNDEDNESS, POSITIVITY) We say that a set of functions $\mathcal{F} = \{F^k : \mathbb{R}_{++}^n \mapsto \mathbb{R}_{++}\}$ is stable monotone sub-homogeneous positive price-bounded with parameter $\theta \in (0, 1)$ (θ -MSPP) if:

- For every $F^k \in \mathcal{F}, i \in [n]$ and for all pairs of vectors $p, q \in \mathbb{R}_{++}^n$ such that $p \geq q$ coordinate-wise, $F^k(p, i) \geq F^k(q, i)$ (monotonicity).
- For every $F^k \in \mathcal{F}, i \in [n], p \in \mathbb{R}_{++}^n, \lambda \in (0, 1)$, $F^k(\lambda p, i) \geq \lambda^\theta F^k(p, i)$ (θ -subhomogeneity).
- $\exists 0 < p_{\min} < p_{\max} < \infty$ such that for all price vectors $p \in [p_{\min}, p_{\max}]^n$, $F^k(p, i) \in [p_{\min}, p_{\max}]$ for every $i \in [n]$ (positive price-boundedness).
- $\exists p^* \in [p_{\min}, p_{\max}]^n$ such that $F^k(p^*, i) = p_i^*$ for each $i \in [n], F^k \in \mathcal{F}$ (stability).

Definition 3. Given a set of mappings $\mathcal{F} = \{F : \mathbb{R}_{++}^n \times [n] \mapsto \mathbb{R}_{++}\}$, its look-ahead closure $\text{Cl}(\mathcal{F})$ is defined recursively as follows:

1. If $G \in \mathcal{F}$ then $G \in \text{Cl}(\mathcal{F})$.




Figure 1: Example recursion tree for a market with three sellers A, B, C

2. For every $\mathcal{I} \subset [n]$, if $F^1, \dots, F^{|\mathcal{I}|} \in \text{Cl}(\mathcal{F})$, $G \in \mathcal{F}$, then the following function is also in $\text{Cl}(\mathcal{F})$:

$$F'(p, i) = G(p', i) \text{ where } p'_k = \begin{cases} p_k & \text{if } k \notin \mathcal{I} \\ F^k(p, k) & \text{if } k \in \mathcal{I} \end{cases}$$

Interpretation of closure operation: belief formation We now interpret the definition of look-ahead closure (Definition 3) as a belief-based price update procedure that the sellers engage in. The set \mathcal{F} corresponds to the set of “atomic price updates” (APU), each of which defines a mapping from the set of prices of all sellers to a new price for a given seller. Concrete examples include best-response (BR) updates for given demand functions (see Section 3). We define \mathcal{F} to be the union of all update rules used by any player and assume that this is common knowledge.

We then augment this set of basic update rules with all rules from the look-ahead closure $\text{Cl}(\mathcal{F})$. We associate elements of this look-ahead closure with a look-ahead-based price update procedure followed by sellers in the market. In order to do this, we look at pairs (i, F) where $i \in [n]$ and $F \in \text{Cl}(\mathcal{F})$. The interpretation of this pair is that seller i updates his price given existing prices of all the sellers using the price update rule $F(\cdot, i)$. F can be interpreted in terms of its recursion trees. By definition, every element of $F' \in \text{Cl}(\mathcal{F})$ is either an element of \mathcal{F} or of the form $F'(p, i) = G(p', i)$ where p' is a function of p and \mathcal{I} . The former case corresponds to a single node tree (with node label i) and the latter corresponds to a rooted tree which is defined recursively: it has root i and children \mathcal{I} (with subtrees corresponding to $\text{Tree}(F^k)$):

$$\text{Tree}(F') = \text{FormTree}\left(i, \text{Tree}(F^1), \dots, \text{Tree}(F^{|\mathcal{I}|})\right)$$

Here $\text{FormTree}(l, T_1, \dots, T_k)$ takes an label l and a set of trees T_1, \dots, T_k and creates a new tree with a root l and children $\text{Root}(T_1), \dots, \text{Root}(T_k)$ with subtrees T_1, \dots, T_k beneath them. Suppose that the recursion tree $\text{Tree}(F')$ is given by the tree in Figure 2. Then the price update (A, F') can be interpreted as follows: A decides to update his price. In order to do this, he simulates the actions of all sellers a few steps into the future and computes a price update based on the simulated state of the game (i.e., prices of all the sellers). In this particular simulation, A assumes that B simply applies an APU rule to current prices, while C applies an APU rule to prices he generates from simulating A and B running some APU rules on current prices.

More general versions of this belief formation procedure are depicted in Figure 2.

We note in passing that beliefs thus formed, implicitly model sellers with epistemic assumptions that they are a bit smarter than their peers—every seller j updates with one extra step beyond the maximum number of steps used in j ’s mental model. Of course, such beliefs cannot possibly be consistent among sellers (unless they are children in Lake Wobegon¹).

The following lemma states the desired properties of belief formation (or equivalently the closure operation).

¹“Well, that’s the news from Lake Wobegon, where all the women are strong, all the men are good-looking, and all the children are above average.”- Monologue, G. Keillor




Figure 2: Various collective mental models for one round of play in a 3-seller market. A leaf (*level 0*) denotes responses to existing prices. In *level 1* dynamics everyone responds to everyone's response to current prices. Players' beliefs can be more complex as well as more heterogeneous, as depicted in the remaining rows of trees.

Lemma 4. *Let \mathcal{F} be a θ -MSPP APU. Then, $\text{Cl}(\mathcal{F})$ is also θ -MSPP.*

Proof. We prove this inductively using the recursive structure of $\text{Cl}(\mathcal{F})$. Let $F \in \text{Cl}(\mathcal{F})$. If $F \in \mathcal{F}$, then we know that F is positive-bounded, monotone, stable and θ -subhomogeneous. Otherwise

$$F(p, i) = G(p', i) \text{ where } p'_k = \begin{cases} p_k & \text{if } k \notin \mathcal{I} \\ F^k(p, k) & \text{if } k \in \mathcal{I} \end{cases}$$

Our inductive hypothesis is that F^k is positive-bounded, monotone and θ -subhomogeneous for each $k \in \mathcal{I}$. Let $p, q \in \mathbb{R}_+^n, p \geq q$ and $i \in [n]$. Then, by inductive hypothesis, we have that $F^k(p, i) \geq F^k(q, i)$, so that $p' \geq q'$. Further, since $G \in \mathcal{F}$, G is monotone, hence $G(p', i) \geq G(q', i)$. Hence, $F(p, i) \geq F(q, i)$ and F is monotone.

Let $\lambda \in (0, 1), i \in [n]$. Then, $F^k(\lambda p, i) \geq \lambda^\theta F^k(p, i) \geq \lambda F^k(p, i)$ (since $\theta \in (0, 1)$). Let $p'(\lambda)$ be the value of p' when p is replaced by λp , so that $F(p, i) = G(p'(1), i)$. We then have $p'(\lambda) \geq \lambda p'(1)$. Since $G \in \mathcal{F}$, we then have $F(\lambda p, i) = G(p'(\lambda), i) \geq G(\lambda p'(1), i) \geq \lambda^\theta G(p'(1), i) = \lambda^\theta F(p, i)$. Thus, F is θ -subhomogeneous as well.

Further, since G, F^1, \dots, F^k are positive and map $[p_{\min}, p_{\max}]^n \times [n] \mapsto [p_{\min}, p_{\max}]$, so does F .

Finally, since $F^k(p^*, i) = p_i^*, G(p^*, i) = p_i^*$, we have $F(p^*, i) = p_i^*$.

Thus, by induction, we have that every $F \in \text{Cl}(\mathcal{F})$ is θ -MSPP, and hence so is $\text{Cl}(\mathcal{F})$. \square

2.1 Convergence of price updates

We now define a synchronous update dynamic for prices given an APU \mathcal{F} .

Definition 5. (SYNCHRONOUS BELIEF-BASED PRICE DYNAMICS) Given an APU \mathcal{F} , the associated belief-based price-update (BBPU) dynamic is defined by the following update equation:

$$p_i^{t+1} = F^{i;t}(p^t, i), \quad i = 1, \dots, n \quad (1)$$

where $F^{1;t}, \dots, F^{n;t} \in \text{Cl}(\mathcal{F})$ are chosen arbitrarily

The interpretation is that every seller updates his price using a belief-based priced update $F^{i;t}$. All sellers update synchronously at the same time instant. The overall dynamic is written as $p^{t+1} = \mathbf{F}^t(p^t)$.

Similarly, we can also study dynamics where sellers update prices asynchronously. Here, we reason over epochs, i.e., periods of time over which every seller updates his price at least once.

Definition 6. (ASYNCHRONOUS BELIEF-BASED PRICE DYNAMICS) Given an APU \mathcal{F} , the associated belief-based price-update dynamics is defined by the following update equation:

$$p_i^{t+1} = F^{i;t}(p^t, i) \quad \forall i \in I^t, \quad p_j^{t+1} = p_j^t \quad \forall j \in [n] \setminus I^t \quad (2)$$

where $I^t \subseteq [n], F^{i;t} \in \text{Cl}(\mathcal{F})$ are chosen arbitrarily

An *epoch* is a period of time $[t_1, t_2]$ over which each seller $i \in [n]$ is chosen at least once for an update, and $[t_1, t]$ does not satisfy the same property for any $t < t_2$. When we speak of multiple epochs they must correspond to successive disjoint time intervals. The epoch-level dynamics can be written as

$$p^{\tau+1} = \mathbf{F}^{[t_1, t_2]}(p^\tau) \quad (3)$$

Clearly, synchronous belief-based price dynamics are a special case of asynchronous belief-based price dynamics, where $I^t = [n]$ for all t .

We are interested in the question of whether the dynamic converges. Before studying this question, we require a definition.

Definition 7. Consider the set $\mathbb{R}_{++}^n \subset \mathbb{R}^n$ of vectors with strictly positive coordinates. The Thompson metric d on \mathbb{R}_{++}^n ([Tho63], and see [LN12] and Lemma 19 for comparison with more common metrics) is defined as follows. For $x, y \in \mathbb{R}_{++}^n$,

$$d(x, y) = \max_i \left| \log \left(\frac{x_i}{y_i} \right) \right| = \|\log x - \log y\|_\infty,$$

where $\log x$ means the vector of logarithms of the entries of x .

We now state our main convergence result, which proves that any BBPU dynamic arising from a θ -MSPP APU converges to a unique set of “equilibrium prices.”

Theorem 8. *Suppose that \mathcal{F} is θ -MSPP for some $\theta \in (0, 1)$. Suppose further that the length of each epoch in the dynamic (3) is bounded uniformly. Then, there is a unique*

$$p^* \in [p_{\min}, p_{\max}]^n \text{ such that } F^k(p^*, i) = p_i^* \quad \forall F^k \in \mathcal{F}$$

Further, the dynamic (3) initialized at any point $p^0 \in [p_{\min}, p_{\max}]^n$ converges to the price vector p^ , and the rate of convergence can be bounded as*

$$d(p^\tau, p^*) \leq \theta^\tau d(p^0, p^*)$$

where p^τ is the price vector after τ epochs.

Proof. Consider an epoch $[t_1, t_2]$. Let u_k be the last time instant at which seller k updates his price. Let $p, q \in [p_{\min}, p_{\max}]^n$ be arbitrary and $d(p, q) = \eta$. Consider the sequence of updates in the epoch applied to the initial price vectors p and q . We have that $p \geq \exp(-\eta)q$, $q \geq \exp(-\eta)p$. By the monotonicity of the update functions, we have that $p^t \geq \exp(-\eta)q^t$ for every $t \in [t_1, t_2]$.

We know that

$$\begin{aligned} F^{u_k}(q^{u_k}, k) &\geq F^{u_k}(\exp(-\eta)p^{u_k}, k) \\ &\geq \exp(-\eta\theta)F^{u_k}(p^{u_k}, k) \\ F^{u_k}(p^{u_k}, k) &\geq F^{u_k}(\exp(-\eta)q^{u_k}, k) \\ &\geq \exp(-\eta\theta)F^{u_k}(q^{u_k}, k) \end{aligned}$$

Thus, we get

$$\begin{aligned} q_k^{t_2} &\geq \exp(-\eta\theta)p_k^{t_2} \\ p_k^{t_2} &\geq \exp(-\eta\theta)q_k^{t_2} \end{aligned}$$

for each $k \in [n]$. Thus, we have that

$$d(\mathbf{F}^{[t_1, t_2]}(p), \mathbf{F}^{[t_1, t_2]}(q)) \leq \theta d(p, q)$$

Also, let p^* be any vector such that $F(p^*, i) = p_i^* \quad \forall F \in \mathcal{F}$ (this is guaranteed to exist by the stability assumption in Definition 2). It is easy to see that $\mathbf{F}^{[t_1, t_2]}(p^*) = p^*$. We then have

$$d(p^{\tau+1}, p^*) = d(\mathbf{F}^{[t_1, t_2]}(p^\tau), \mathbf{F}^{[t_1, t_2]}(p^*)) \leq \theta d(p^\tau, p^*).$$

This establishes convergence. Since p^* was chosen arbitrarily among the set of vectors satisfying the stability assumption from Definition 2, this also establishes uniqueness of p^* . \square

This theorem allows the possibility of an arbitrary subset of sellers updating prices synchronously at any given time instant. Thus, it includes the following corollary as a special case.

Corollary 9. Suppose that \mathcal{F} is θ -MSPP for some $\theta \in (0, 1)$. Then, the dynamic (1) initialized at any point $p^0 \in [p_{\min}, p_{\max}]^n$ converges to the price vector p^* from Theorem 8. Further, the rate of convergence can be bounded as

$$d(p^t, p^*) \leq (\theta)^t d(p^0, p^*)$$

where p^t is the price after t time steps.

Remark 10. We later apply this theorem to particular update rules that sellers will strategically employ in case the demand functions are known to satisfy favorable conditions (a bit more than Weak Gross Substitutes). For such demand functions it is long known that the market has a unique equilibrium; in that case, the equilibrium generated by our strategic agents will of course be the same. But Theorem 8 and Corollary 9 are more general (and in particular not implied by the classic results) because the APU's need not be generated by a demand model.

3 Concrete markets with rationalizable MSPP updates: elastic demands

We now describe a concrete setting where the updates satisfy our axioms. In particular, we focus on price update rules that arise as best-responses (BR) to other sellers' prices and a fixed demand function.

Given a demand function $x(p)$ and two items j and k , we define the following quantity:

$$\frac{\partial \log \left(\sum_{i \in [m]} x_{ij}(p) \right)}{\partial \log(p_k)} \quad (\text{price elasticity})$$

(Throughout the paper we assume that all these partial derivatives exist, so henceforth we will not state this assumption explicitly.) If $j = k$ the price elasticity is called the *own price elasticity* (of the demand for j), and otherwise it is called the *cross price elasticity* (of the demand for j with respect to the price of k). Informally, price elasticity connects between the percentage of change in price and the percentage of change in demand. A demand function $x(p)$ is said to satisfy the *weak gross substitutes* (WGS) property if x is differentiable with respect to p and

$$\frac{\partial x_{ij}(p)}{\partial p_k} \geq 0 \quad \forall i \in [m], j \in [n], k \in [n] \setminus \{j\}.$$

Notice that if $x(p)$ satisfies WGS then this implies that all cross price elasticities are non-negative (but the reverse assertion does not necessarily hold).

We will discuss demand functions $x(p)$ that satisfy the following definition.

Definition 11. A demand function $x(p)$ is *elastic and bounded* (E&B) with parameter $\kappa \in (0, 1)$ iff it satisfies the following three properties for some $\mu \geq 1$:

- $\forall j \in [n], k \in [n] \setminus \{j\}, p \in \mathbb{R}_{++}^n$:

$$\frac{\partial \log \left(\sum_{i \in [m]} x_{ij}(p) \right)}{\partial \log(p_k)} \geq 0 \quad (5a)$$

(Cross price elasticities non-negative, implied by WGS)

- $\forall j \in [n], p \in \mathbb{R}_{++}^n$:

$$\frac{\partial \log \left(\sum_{i \in [m]} x_{ij}(p) \right)}{\partial \log(p_j)} \leq -\mu \quad (5b)$$

(Own price elasticities less than $-\mu$)

- $\forall j \in [n], p \in \mathbb{R}_{++}^n$:

$$\sum_{k \in [n]} \frac{\partial \log \left(\sum_{i \in [m]} x_{ij}(p) \right)}{\partial \log(p_k)} \geq -\kappa \mu \quad (5c)$$

(Net price elasticities greater than $-\kappa \mu$)

Remark 12. Property (5a) is implied by the Weak Gross Substitutes (WGS) property but does not require it (since (5a) is on the aggregate demand for a good while WGS is usually defined on every pair of goods). However, this does not suffice to ensure that sellers behaving strategically will remain at equilibrium prices even if market is already at equilibrium. For $\mu = 1$, the property (5b) is equivalent to the statement that the profit of a seller is always monotonically decreasing in the price of this seller's good (Corollary 14). Without this property, sellers may not have incentive to stay put at the market equilibrium (assuming one exists). They can increase their profit by increasing prices and reducing the aggregate demand for their good, thus leaving the market uncleared.

Lemma 13. *Consider a demand function x that satisfies the own price elasticity conditions of Definition 11. Fix a price vector p and a good j . Consider all price vectors p' with the property that for all $k \neq j$, $p'_k = p_k$. and such that $p'_j \geq p'_{\text{crit},j}$ where $p'_{\text{crit},j}$ is the price at which $\sum_{i \in [m]} x_{ij} = 1$. Among these price vectors, $\sum_{i \in [m]} x_{ij}(p') \cdot p'_j$ is monotonically decreasing in p'_j .*

Proof. Consider the aggregate spending $\sum_i x_{ij}(p') \cdot p'_j$ on good j . The derivative with respect to p'_j is

$$\begin{aligned} \sum_i \frac{\partial x_{ij}(p')}{\partial p'_j} p'_j + x_{ij}(p') &= \frac{\partial \sum_i x_{ij}(p')}{\partial p'_j} p'_j + \left(\sum_i x_{ij}(p') \right) \\ &= \left(\frac{\partial \log(\sum_i x_{ij}(p'))}{\partial \log(p'_j)} + 1 \right) \left(\sum_i x_{ij}(p') \right). \end{aligned}$$

This expression is negative because of the assumed own price elasticities. \square

Best-response updates In standard best-response dynamics, each seller updates its price to maximize its revenue given the current prices of the other players. In the particular setting of E&B demand functions, a seller j maximizes profit by setting the price p_j to clear the market for good j by the following corollary:

Corollary 14. *If the prices of all goods except j are fixed, the profit of seller j is maximized by setting p_j such that $\sum_{i \in [m]} x_{ij}(p) = 1$.*

Proof. As $p'_j > 0$ and $\sum_{i \in [m]} x_{ij}(p') > 0$, we have that

$$\frac{\partial \sum_{i \in [m]} x_{ij}(p')}{\partial p'_j} = \frac{\sum_{i \in [m]} x_{ij}(p')}{p'_j} \cdot \frac{\partial \log \left(\sum_{i \in [m]} x_{ij}(p') \right)}{\partial \log(p'_j)} < 0,$$

so the aggregate demand $\sum_i x_{ij}(p')$ for j decreases monotonically in p'_j . By Lemma 13, also the aggregate spending on j decreases monotonically in p'_j . Therefore, the profit is maximized at the lowest price for which the demand is at most 1 (lowering the price further will not increase the quantity sold beyond the initial endowment). \square

As a consequence of Corollary 14, if the current price vector is p , the seller chooses a new price $F_{br}(p)$ for good j , so that

$$\sum_{i=1}^m x_{ij}(p') = 1, \quad (6)$$

where $p'_j = F_{br}(p, j)$, and for all $k \neq j$, $p'_k = p_k$. The following lemma shows that under the conditions (5), F_{br} is well defined and equilibrium prices exist.

Lemma 15. Suppose that x satisfies the conditions (5). Then, F_{br} is well defined (ie (6) has a unique solution for every $j \in [n]$). Further, $\exists 0 < p_{\min} < p_{\max} < \infty$ such that $F_{br}(\cdot, j)$ maps $[p_{\min}, p_{\max}]^n$ to $[p_{\min}, p_{\max}]$. Further, $\exists p^* \in [p_{\min}, p_{\max}]^n$ such that

$$\sum_{i \in [m]} x_{ij}(p^*) = 1 \quad \forall i \in [n] \quad (p^* \text{ clears the market}) \quad (7a)$$

$$F_{br}(p^*, j) = p_j^* \quad (\text{Best-response updates leave } p^* \text{ unchanged}) \quad (7b)$$

(By Corollary 14 these consequences are equivalent.)

Proof. Using (5b), we know that $\sum_{i \in [m]} x_{ij}(p_j, p_{\sim j})$ is strictly monotone in p_j . Thus, (6) has a unique solution for each $j \in [n], p \in \mathbb{R}_{++}^n$ and hence F_{br} is well-defined. Using (5c), we know that $\forall p \in \mathbb{R}_{++}^n, \lambda \in (0, 1), j \in [n]$:

$$\left(\sum_{i \in [m]} x_{ij}(\lambda p) \right) \geq \frac{1}{\lambda^{\kappa\mu}} \left(\sum_{i \in [m]} x_{ij}(p) \right)$$

Let $\alpha = \min_{j \in [n]} \sum_{i \in [m]} x_{ij}(\mathbf{1})$. Choose $p_{\min} = (\min(\alpha, 1))^{\frac{1}{\kappa\mu}}$. Then we have $\forall j \in [n]$:

$$\sum_{i \in [m]} x_{ij}(p_{\min} \mathbf{1}) \geq \frac{1}{\min(\alpha, 1)} \left(\sum_{i \in [m]} x_{ij}(\mathbf{1}) \right) \geq 1$$

Therefore, using (5a),

$$\sum_{i \in [m]} x_{ij}(p_{\min}, p_{\sim j}) \geq 1 \quad \forall j \in [n], p_{\sim j} \in [p_{\min}, \infty)^{n-1} \quad (8)$$

Similarly, define $\beta = \max_{j \in [n]} \sum_{i \in [m]} x_{ij}(\mathbf{1})$. Choose $p_{\max} = (\max(\beta, 1))^{\frac{1}{\kappa\mu}}$. Then, we have $\forall j \in [n]$:

$$\sum_{i \in [m]} x_{ij}(p_{\max} \mathbf{1}) \leq \frac{1}{\max(\beta, 1)} \left(\sum_{i \in [m]} x_{ij}(\mathbf{1}) \right) \leq 1$$

Again, using (5a), we get

$$\sum_{i \in [m]} x_{ij}(p_{\max}, p_{\sim j}) \leq 1 \quad \forall j \in [n], p_{\sim j} \in (0, p_{\max}]^{n-1} \quad (9)$$

Combining (8) and (9), and using the fact that $\sum_{i \in [m]} x_{ij}$ is strictly decreasing in p_j , we get that

$$F_{br}(p, j) \in [p_{\min}, p_{\max}] \quad \forall j \in [n], p \in [p_{\min}, p_{\max}]^n$$

Further, $F_{br}(\cdot, j)$ is continuous in p for each j . Define the map $\mathbf{F} : \mathbb{R}_{++}^n \mapsto \mathbb{R}_{++}^n$ as

$$[\mathbf{F}(p)]_j = F_{br}(p, j) \quad \forall j \in [n]$$

\mathbf{F} is continuous and maps the compact convex set $[p_{\min}, p_{\max}]^n$ into itself. Hence it must have a fixed point $p^* \in [p_{\min}, p_{\max}]^n$ (by Brouwer's fixed point theorem). Thus, $\mathbf{F}(p^*) = p^*$, or

$$F_{br}(p^*, j) = p_j^* \quad \forall j \in [n]$$

Hence the lemma. \square

Lemma 16. *Let F_{br} be the best response update for an E&B demand function with parameter κ . Then, the set $\{F_{br}\}$ is a $(1 - \kappa)$ -MSPP APU.*

Proof. We begin with monotonicity. Denote by $x_{ij}(p_j, p_{\sim j})$ the demand of buyer i for good j when the price of good j is p_j and the price vector of all the other goods is $p_{\sim j}$.

Consider two price vectors $p \geq q$. We know that

$$\sum_{i \in [m]} x_{ij}(F_{br}(q, j), p_{\sim j}) \geq \sum_{i \in [m]} x_{ij}(F_{br}(q, j), q_{\sim j}) = 1,$$

which follows by the cross price elasticity conditions of Definition 11. By the own price elasticity conditions (5b), since $\sum_{i \in [m]} x_{ij}(F_{br}(p, j), p_{\sim j}) = 1$, we have that $F_{br}(p, j) \geq F_{br}(q, j)$.

Next we prove sub-homogeneity. Let $p > 0$ be a price vector and $\lambda \in (0, 1)$. We have

$$\sum_{i \in [m]} x_{ij}(\lambda F_{br}(p, j), \lambda p_{\sim j}) \geq \frac{1}{\lambda^{\kappa\mu}} \left(\sum_{i \in [m]} x_{ij}(F_{br}(p, j), p_{\sim j}) \right) = \frac{1}{\lambda^{\kappa\mu}},$$

where the first inequality uses (5c), and the second inequality uses the definition of F_{br} (6). Using (5b), we get

$$\frac{\sum_{i \in [m]} x_{ij}(\lambda F_{br}(p, j), \lambda p_{\sim j})}{\sum_{i \in [m]} x_{ij}(F_{br}(\lambda p, j), \lambda p_{\sim j})} \leq \left(\frac{F_{br}(\lambda p, j)}{\lambda F_{br}(p, j)} \right)^\mu$$

Since $\sum_i x_{ij}(F_{br}(\lambda p, j), \lambda p_{\sim j}) = 1$, this reduces to

$$\frac{1}{\lambda^{\kappa\mu}} \leq \sum_{i \in [m]} x_{ij}(\lambda F_{br}(p, j), \lambda p_{\sim j}) \leq \left(\frac{F_{br}(\lambda p, j)}{\lambda F_{br}(p, j)} \right)^\mu$$

Taking μ -th roots, we get

$$F_{br}(\lambda p, j) \geq \lambda^{1-\kappa} F_{br}(p, j)$$

which shows $(1 - \kappa)$ -subhomogeneity.

Stability and price-boundedness follows from Lemma 15. \square

Corollary 17. *The BBPU dynamics defined by $\text{Cl}(\{F_{br}\})$ converges to a unique set of equilibrium prices p^* as*

$$d(p^t, p^*) \leq (1 - \kappa)^t d(p^0, p^*)$$

where p^0 is the initial price vector and p^t is the price vector after t time steps (for synchronous updates) or t epochs (for asynchronous updates).

Proof. From Lemma 16 and Theorems 8 and 9. \square

3.1 WGS CES utilities

A particular class of utilities that give rise to E&B demand functions are CES (constant elasticity of substitution) utilities in the WGS regime. CES utility functions are of the following form (i indexes a buyer):

$$u_i(x) = \left(\sum_j (c_{ij} x_{ij})^\rho \right)^{\frac{1}{\rho}} \quad (10)$$

where $\rho \in (-\infty, 0) \cup (0, 1)$, $c_{ij} \geq 0$. CES utilities are WGS iff $\rho \geq 0$, as we subsequently assume they are. Define $\epsilon = \rho/(1 - \rho) > 0$. The demand function x is given by

$$x_i(p) = \operatorname{argmax}_{x_i: \sum_{k \in [n]} x_{ik} p_k \leq b_i} u_i(x_i).$$

where $b_i > 0$ is the budget of buyer i .

We have the following result about convergence of BBPU dynamics for the best response update with the above demand function:

Corollary 18. *Let F_{br} be the best-response update corresponding to the above demand function x . The BBPU dynamics associated with $\text{Cl}(\{F_{br}\})$ converges to a unique set of equilibrium prices $p^* \in [p_{\min}, p_{\max}]^n$ as*

$$d(p^t, p^0) \leq (1 - \kappa)^t d(p^0, p^*)$$

if initialized at any point $p^0 \in [p_{\min}, p_{\max}]^n$ and p^t is the price after t time steps (for synchronous updates) or t epochs (for asynchronous updates), where (for ϵ as specified above)

$$\begin{aligned} \kappa &= \frac{1}{1 + \epsilon \left(1 - \min_{j \in [n]} \max_{i \in [m]} \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_{\min}}{p_{\max}} \right)^\epsilon} \right)} \\ p_{\min} &= \min \left(\min_{j \in [n]} \sum_{i \in [m]} b_i \frac{c_{ij}}{\sum_{k \in [n]} c_{ik}}, 1 \right) \\ p_{\max} &= \max \left(\max_{j \in [n]} \sum_{i \in [m]} b_i \frac{c_{ij}}{\sum_{k \in [n]} c_{ik}}, 1 \right) \end{aligned}$$

Proof. Follows from Theorems 8 combined with Lemma 20 in the appendix. \square

We note here a brief comparison to the work [CFR10]. The assumptions made in the paper are closely related to the assumptions on the demand functions in Equation (5). However, [CFR10] analyzed tâtonnement with CES WGS utilities (instead of best response). They obtain a convergence rate (Theorem 2 in [CFR10]) that shows that the number of synchronous updates required to obtain a price vector a certain distance from equilibrium is linear in a parameter d (which can be seen as a one-sided upper bound on the Thompson-metric d from our paper). Because the measure of distance to equilibrium used in that paper is not shown to be a metric, a straightforward comparison is not possible, but it appears that the exponential convergence rate obtained here will almost always dominate.

4 Acknowledgements

KD was supported by a fellowship from the Caltech Center for the Mathematics of Information. YR was supported in part by ISF grant 956-15, BSF grant 2012333 and I-CORE Algo. LJS was supported in part by NSF grants 1319745, 1618795, 1909972, and BSF grant 2012333. An earlier version of this paper appeared in Proc. 28th SODA, 2017.

References

- [AAE⁺08] B. Awerbuch, Y. Azar, A. Epstein, V. S. Mirrokni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *Proc. 9th ACM Conf. Electronic Commerce*, pages 264–273, 2008.
- [ABH59] K. J. Arrow, H. D. Block, and L. Hurwicz. On the stability of the competitive equilibrium: II. *Econometrica*, 27(1):82–109, 1959.
- [AD54] K. J. Arrow and G. Debreu. Existence of equilibrium for a competitive economy. *Econometrica*, 22:265–290, 1954.

- [AFP⁺16] V. Auletta, D. Ferraoli, F. Pasquale, P. Penna, and G. Persiano. Convergence to equilibrium of logit dynamics for strategic games. *Algorithmica*, 76:110–142, 2016.
- [AGM⁺11] H. Ackermann, P. W. Goldberg, V. Mirrokni, H. Röglin, and B. Vöcking. Uncoordinated two-sided matching markets. *SIAM J. Comput.*, 40:92–106, 2011. (EC’08).
- [BCI⁺10] C. Borgs, J. T. Chayes, N. Immorlica, A. T. Kalai, V. S. Mirrokni, and C. H. Papadimitriou. The myth of the folk theorem. *Games and Economic Behavior*, 70(1):34–43, 2010.
- [BLNPL14] M. Babaioff, B. Lucier, N. Nisan, and R. Paes Leme. On the efficiency of the walrasian mechanism. In *Proc. 15th ACM Conference on Economics and Computation*, pages 783–800. ACM, 2014.
- [BPLS15] M. Babaioff, R. Paes Leme, and B. Sivan. Price competition, fluctuations and welfare guarantees. In *Proc. 16th ACM Conference on Economics and Computation*, pages 759–776, New York, NY, USA, 2015. ACM.
- [CCD13] Y. K. Cheung, R. Cole, and N. R. Devanur. Tatonnement beyond gross substitutes?: gradient descent to the rescue. In *Proc. 45th Ann. ACM Symp. Theory of Computing*, pages 191–200, 2013.
- [CCGI13] V. P. Crawford, M. A. Costa-Gomes, and N. Iriberri. Structural models of non-equilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1):5–62, 2013.
- [CCR12] Y. K. Cheung, R. Cole, and A. Rastogi. Tatonnement in ongoing markets of complementary goods. In *Proc. 13th Ann. ACM Conf. Electronic Commerce*, pages 337–354, 2012.
- [CF08] R. Cole and L. Fleischer. Fast-converging tatonnement algorithms for one-time and ongoing market problems. In *Proc. 40th Ann. ACM Symp. Theory of Computing*, pages 315–324, 2008.
- [CFR10] R. Cole, L. Fleischer, and A. Rastogi. Discrete price updates yield fast convergence in ongoing markets with finite warehouses. *ArXiv:1012.2124*, 2010.
- [CGC06] M. A. Costa-Gomes and V. P. Crawford. Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, 96(5):1737–1768, 2006.
- [CGCB01] M. A. Costa-Gomes, V. P. Crawford, and B. Broseta. Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235, 2001.
- [CHC04] C. F. Camerer, T.-H. Ho, and J.-K. Chong. A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119(3):861–898, 2004.
- [CI07a] V. P. Crawford and N. Iriberri. Fatal attraction: Salience, naïveté, and sophistication in experimental “hide-and-seek” games. *American Economic Review*, 97(5):1731–1750, 2007.
- [CI07b] V. P. Crawford and N. Iriberri. Level-k auctions: Can a non-equilibrium model of strategic thinking explain the winner’s curse and overbidding in private-value auctions? *Econometrica*, 75(6):1721–1770, 2007.
- [Cra03] V. P. Crawford. Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions. *American Economic Review*, 93(1):133–149, 2003.
- [CS11] S. Chien and A. Sinclair. Convergence to approximate Nash equilibria in congestion games. *Games and Economic Behavior*, 71(2):315–327, 2011.

- [dCSS14] G. de Clippel, R. Saran, and R. Serrano. Mechanism design with bounded depth of reasoning and small modeling mistakes. Working Papers 2014-7, Brown University, Dept. Economics, 2014. Downloaded Oct. 28, 2015.
- [Dix90] H. Dixon. Equilibrium and explanation. In J. Creedy, editor, *The Foundations of Economic Thought*, pages 356–394. Blackwell, 1990.
- [EFSW13] R. Engelberg, A. Fabrikant, M. Schapira, and D. Wajc. Best-response dynamics out of sync: Complexity and characterization. In *Proc. 14th ACM Conf. Electronic Commerce*, pages 379–396. ACM, 2013.
- [FFM12] A. Fanelli, M. Flammini, and L. Moscardelli. The speed of convergence in congestion games under best-response dynamics. *ACM Trans. Algorithms*, 8(3):25:1–25:15, July 2012.
- [FPT04] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proc. Thirty-sixth Ann. ACM Symp. Theory of Computing*, pages 604–612. ACM, 2004.
- [Gor18] O. Gorelkina. The expected externality mechanism in a level-k environment. *International Journal of Game Theory*, 47(1):103–131, Mar 2018.
- [HCW98] T.-H. Ho, C. Camerer, and K. Weigelt. Iterated dominance and iterated best response in experimental ‘p-beauty contests’. *American Economic Review*, 88(4):947–969, 1998.
- [Hil98] W. Hildenbrand. An exposition of Wald’s existence proof. In E. Dierker and K. Sigmund, editors, *Karl Menger*, pages 51–61. Springer, 1998.
- [HMC12] S. Hart and A. Mas-Colell. *Simple Adaptive Strategies*. World Scientific, 2012.
- [HPS14] J. Y. Halpern, R. Pass, and L. Seeman. Not just an empty threat: Subgame-perfect equilibrium in repeated games played by computationally bounded players. In *Proc. 10th Int’l Conf. on Web and Internet Economics*, pages 249–262, 2014.
- [JV10] K. Jain and V. V. Vazirani. Eisenberg-Gale markets: Algorithms and game-theoretic properties. *Games and Economic Behavior*, 70(1):84–106, 2010.
- [Kne15] T. Kneeland. Identifying higher-order rationality. *Econometrica*, 83(5):2065–2079, 2015.
- [LN12] B. Lemmens and R. Nussbaum. *Nonlinear Perron-Frobenius Theory*. Cambridge University Press, 2012. Cambridge Books Online.
- [McK54] L. McKenzie. On equilibrium in Graham’s model of world trade and other competitive systems. *Econometrica*, 22:147–161, 1954.
- [McK02] L. W. McKenzie. *Classical general equilibrium theory*. MIT Press, 2002.
- [MR90] P. Milgrom and J. Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277, 1990.
- [MR91] P. Milgrom and J. Roberts. Adaptive and sophisticated learning in normal form games. *Games and Economic Behavior*, 3:82–100, 1991.
- [MS96] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- [Muk02] A. Mukherji. *An introduction to general equilibrium analysis*. Oxford U Press, 2002.
- [MV04] V. S. Mirrokni and A. Vetta. Convergence issues in competitive games. In K. Jansen, S. Khanna, J. D. P. Rolim, and D. Ron, editors, *Proc. APPROX and RANDOM*, LNCS 3122, pages 183–194. Springer, 2004.

- [Nag95] R. Nagel. Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–1326, 1995.
- [NSVZ11] N. Nisan, M. Schapira, G. Valiant, and A. Zohar. Best-response mechanisms. In *Proc. Innovations in Computer Science - ICS*, pages 155–165, 2011.
- [Rou15] T. Roughgarden. Intrinsic robustness of the price of anarchy. *J. ACM*, 62(5):32:1–42, November 2015.
- [Sam41] P. A. Samuelson. The stability of equilibrium: Comparative statics and dynamics. *Econometrica*, 9:97–120, 1941.
- [SLB09] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009.
- [ST13] A. Saha and A. Tewari. On the nonasymptotic convergence of cyclic coordinate descent methods. *SIAM Journal on Optimization*, 23(1):576–601, 2013.
- [Str14] T. Strzalecki. Depth of reasoning and higher order beliefs. *Journal of Economic Behavior & Organization*, 108:108–122, 2014.
- [SW94] D. O. Stahl and P. W. Wilson. Experimental evidence on players’ models of other players. *Journal of Economic Behavior and Organization*, 25(3):309–327, 1994.
- [SW95] D. O. Stahl and P. W. Wilson. On players’ models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254, 1995.
- [Tho63] A. C. Thompson. On certain contraction mappings in a partially ordered vector space. *Proc. American Mathematical Society*, 14(3):438–443, 1963.
- [Tse01] P. Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *J. optimization theory and applications*, 109(3):475–494, 2001.
- [Wal74] L. Walras. *Éléments d’Economie Politique Pure*. Corbaz, 1874. (1st ed. 1874; revised ed. 1926; Transl. W. Jaffé, Elements of Pure Economics, Irwin, 1954).

5 Appendix

5.1 Thompson metric

The following lemma shows how this metric is related to the standard ℓ_2 and ℓ_∞ metrics.

Lemma 19. *Let $p^a, p^b \in [p_{\min}, p_{\max}]^n$ with $0 < p_{\min} < p_{\max}$. Then, we have*

$$\begin{aligned} \|p^a - p^b\|_\infty &\leq \frac{(p_{\max})^2}{p_{\min}} d(p^a, p^b) \\ \|p^a - p^b\|_2 &\leq \sqrt{n} \frac{(p_{\max})^2}{p_{\min}} d(p^a, p^b) \end{aligned}$$

Proof. For each i , we have

$$|p_i^a - p_i^b| \leq p_{\max} \left| \frac{p_i^a}{p_i^b} - 1 \right| \leq p_{\max} (\exp(d(p^a, p^b)) - 1)$$

The function $f(t) = \exp(t) - 1 - \kappa t$ is non-increasing on the interval $[0, \log(\kappa)]$ for every $\kappa > 0$ and evaluates to 0 at $t = 0$. Hence $\exp(t) - 1 \leq \kappa t$ for every $t \in [0, \log(\kappa)]$.

Since $d(p^a, p^b) \leq \log\left(\frac{p_{\max}}{p_{\min}}\right)$, we can choose $\kappa = \frac{p_{\max}}{p_{\min}}$ and conclude that

$$\exp(d(p^a, p^b)) - 1 \leq \frac{p_{\max}}{p_{\min}} d(p^a, p^b)$$

Thus, we have

$$|p_i^a - p_i^b| \leq p_{\max} \frac{p_{\max}}{p_{\min}} d(p^a, p^b) = \frac{(p_{\max})^2}{p_{\min}} d(p^a, p^b)$$

Since the bound holds for each i , it holds for the ∞ norm as well. The 2-norm bound simply uses the fact that the 2 norm is at most \sqrt{n} times the infinity norm. \square

5.2 WGS CES utilities

Lemma 20. *In a Fisher market where the buyers have CES utilities which are WGS ($0 < \rho < 1$, correspondingly $\epsilon = \rho/(1 - \rho) > 0$), the demand function is \mathcal{EB} with parameter*

$$\frac{1}{1 + \epsilon \left(1 - \min_{j \in [n]} \max_{i \in [m]} \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_{\min}}{p_{\max}} \right)^\epsilon} \right)}$$

over the set $[p_{\min}, p_{\max}]^n$ where

$$p_{\min} = \min \left(\min_{j \in [n]} \sum_{i \in [m]} b_i \frac{c_{ij}}{\sum_{k \in [n]} c_{ik}}, 1 \right) \quad (11)$$

$$p_{\max} = \max \left(\max_{j \in [n]} \sum_{i \in [m]} b_i \frac{c_{ij}}{\sum_{k \in [n]} c_{ik}}, 1 \right) \quad (12)$$

Proof. We start by writing down $x(p)$ in an explicit form. It can be shown that $\forall i \in [m], j \in [n]$

$$x_{ij}(p) = \frac{b_i}{p_j} \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_j}{p_k} \right)^\epsilon}$$

Then, we have $\forall j \in [n]$

$$f_j(p) = \sum_{i \in [m]} x_{ij}(p) = \sum_{i \in [m]} \frac{b_i}{p_j} \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_j}{p_k} \right)^\epsilon}$$

Clearly, f_j is differentiable at all points $p \in \mathbb{R}_{++}^n$. All conditions from Definition 11 are on derivatives of this quantity. First, it is easy to see that

$$f_j(\lambda p) = \frac{1}{\lambda} f_j(p) \geq \frac{1}{\lambda} f_j(p) \quad \forall \lambda \in (0, 1)$$

This establishes property (5c) with $\kappa\mu = 1$. Further, we have

$$\frac{\partial f_j}{\partial p_k} \geq 0 \quad \forall k \in [n] \setminus \{j\}$$

since the parameters b_i, c_{ik}, ϵ are positive. Finally, we can compute the derivative

$$\frac{\partial \log(f_j)}{\partial \log(p_j)} = -1 - \epsilon \left(\frac{\sum_{i \in [m]} b_i \frac{c_{ij} \left(\sum_{k \neq j} c_{ik} \left(\frac{p_j}{p_k} \right)^\epsilon \right)}{\left(c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_j}{p_k} \right)^\epsilon \right)^2}}{\sum_{i \in [m]} b_i \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_j}{p_k} \right)^\epsilon}} \right)$$

Define the probability distribution

$$q_i = \frac{b_i \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_j}{p_k}\right)^\epsilon}}{\sum_{i \in [m]} b_i \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_j}{p_k}\right)^\epsilon}}$$

Then the derivative can be written as

$$\begin{aligned} \frac{\partial \log(f_j)}{\partial \log(p_j)} &= -1 - \epsilon \left(\sum_{i \in [m]} q_i \left(1 - \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_j}{p_k}\right)^\epsilon} \right) \right) \\ &\leq -1 - \epsilon \min_{i \in [m]} \left(1 - \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_j}{p_k}\right)^\epsilon} \right) \\ &= -1 - \epsilon + \epsilon \max_{i \in [m]} \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_j}{p_k}\right)^\epsilon} \\ &\leq -1 - \epsilon + \epsilon \max_{i \in [m]} \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_{\min}}{p_{\max}}\right)^\epsilon} \quad \forall p \in [p_{\min}, p_{\max}]^n \end{aligned}$$

for any $0 < p_{\min} < p_{\max} < \infty$.

Following the argument of Lemma 16, it can be established that the corresponding best response update $F_{br}(\cdot, j)$ is well-defined and $F_{br}(p, j) \in [p_{\min}, p_{\max}]^n \forall p \in [p_{\min}, p_{\max}]^n$ where

$$\begin{aligned} p_{\min} &= \min \left(\min_{j \in [n]} f_j(\mathbf{1}), 1 \right) \\ p_{\max} &= \max \left(\max_{j \in [n]} f_j(\mathbf{1}), 1 \right) \end{aligned}$$

and that the equilibrium price vector $p^* \in [p_{\min}, p_{\max}]^n$ exists. Since the convergence argument from Theorem 8 only depends on properties of F_{br} for price vectors in $[p_{\min}, p_{\max}]^n$, we only need to establish property (5b) over this set. Thus, we have established properties (5a),(5b),(5c) with parameters

$$\begin{aligned} \mu &= 1 + \epsilon \left(1 - \min_{j \in [n]} \max_{i \in [m]} \frac{c_{ij}}{c_{ij} + \sum_{k \neq j} c_{ik} \left(\frac{p_{\min}}{p_{\max}}\right)^\epsilon} \right) > 1 \\ \kappa &= \frac{1}{\mu} < 1 \end{aligned}$$

where p_{\min}, p_{\max} are as given in Equation (11). □