Tracing Equilibrium in Dynamic Markets via Distributed Adaptation

Yun Kuen Cheung
Singapore University of Technology and Design
Singapore
yunkuen_cheung@sutd.edu.sg

Martin Hoefer
Institut für Informatik,
Goethe-Universität Frankfurt/Main
Frankfurt, Germany
mhoefer@cs.uni-frankfurt.de

Paresh Nakhe
Institut für Informatik,
Goethe-Universität Frankfurt/Main
Frankfurt, Germany
paresh.nakhe@gmail.com

ABSTRACT

In real-world decentralized systems, agents' actions are often coupled with changes in the environment which are out of the agents' control. Yet, in many important domains, the existing analyses presume static environments. The theme of our work is to bridge such a gap between existing work and reality, with a focus on markets.

Competitive (market) equilibrium is a central concept in economics with numerous applications beyond markets, such as scheduling, fair allocation of goods, or bandwidth distribution in networks. Natural and decentralized processes like tatonnement and proportional response dynamics (PRD) are known to converge quickly towards equilibrium in large classes of *static* Fisher markets. In contrast, many large real-world markets are subject to frequent and dynamic changes. We provide the first provable performance guarantees of discrete-time tatonnement and PRD in *dynamic markets*. We analyze the prominent class of CES (Constant Elasticity of Substitution) Fisher markets and quantify the impact of changes in supplies of goods, budgets of agents, and utility functions of agents on the convergences of the processes to equilibrium. Since the equilibrium becomes a dynamic object and will rarely be reached, we provide bounds expressing the distance to equilibrium that will be maintained. Our results indicate that in many cases, the processes trace the equilibrium rather closely and quickly recover conditions of approximate market clearing.

Our analyses proceed by quantifying the impact of variation in market parameters on several potential functions which guarantee convergences in static settings. This approach is captured in two general yet handy frameworks for Lyapunov dynamical systems. They are of independent interest, which we demonstrate with the analysis of load balancing in dynamic environment setting.

KEYWORDS

Fisher Markets; Peer-to-Peer Networks; Tatonnement; Proportional Response Dynamics; Dynamic Markets

Most of the work done while Yun Kuen Cheung was a postdoctoral fellow at Max Planck Institute for Informatics, Saarland Informatics Campus. He would like to acknowledge Singapore NRF 2018 Fellowship NRF-NRFF2018-07 and MOE AcRF Tier 2 Grant 2016-T2-1-170.

Martin Hoefer and Paresh Nakhe acknowledge funding by DFG Cluster of Excellence M2CI..

Full version of this paper is available at https://arxiv.org/abs/1804.08017.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

ACM Reference Format:

Yun Kuen Cheung, Martin Hoefer, and Paresh Nakhe. 2019. Tracing Equilibrium in Dynamic Markets via Distributed Adaptation. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 9 pages.

1 INTRODUCTION

A central concept to understand large economic systems is the notion of competitive or market equilibrium. The computational aspects of competitive equilibria have been a central theme at the intersection of game theory and computer science over the last decade, mainly for the prominent class of Fisher markets. In a Fisher market, there are a set of agents or buyers and a set of divisible goods. Each agent brings a budget of money to the market and wants to buy goods, for which she has an increasing and concave utility function. An equilibrium consists of a vector of prices and an allocation of goods and money such that (1) every agent purchases the most preferred bundle of goods that she can afford, and (2) market clears (supply equals demand).

There are successful approaches based on distributed adaptation processes for converging to competitive equilibria. For example, *tatonnement* is governed by the natural intuition that prices of over-demanded goods increase, while under-demanded goods become cheaper. It provides an explanation how decentralized price adjustment can lead a market into an equilibrium state, thereby providing additional justification for the concept. Recently, several works derived detailed analyses and proved fast convergence of discrete-time tatonnement in markets [6, 8, 9, 12, 13].

It is well-known that network rate control is closely related to Fisher competitive equilibria [20–22]. Towards this end, distributed market dynamics called *proportional response dynamics* (PRD) were proposed and analyzed in the context of peer-to-peer network [23, 28]. These dynamics avoid the usage of prices and work directly on the exchange and allocation of goods. PRD and its generalizations converge toward competitive equilibria in the full range of CES Fisher markets (see Section 2 for its definition) [5, 10, 30].

While tatonnement and PRD rely on dynamic changes of prices and allocations, the existing literature assumes that the market and its properties (agents, budgets, utilities, supplies of goods) remain static and unchanged over time. In fact, to the best of our knowledge, all of the existing work on computation of competitive equilibrium in algorithmic game theory assumes that the market is essentially a *static environment*. In contrast, in many (if not all) applications of markets, the market itself is subject to dynamic change, in the sense that supplies of goods changes over time, agents have different budgets at their disposal that they can spend, or the preferences of agents expressed via utility functions evolve over time. Analyzing and quantifying the impact of dynamic change in markets is critical to understand the robustness of competitive equilibrium in general, and of price adaptation dynamics like tatonnement in particular.

In this paper, we initiate the algorithmic study of *dynamic markets* in the form of dynamically evolving environments. Our interest lies in the performance of dynamic adaptation processes like tatonnement. We analyze a discrete-time process: in each round, the excess demands in the previous round are used to perform tatonnement updates. Simultaneously, the market is subject to (possibly adversarial) variations. This dynamic nature of markets gives rise to a number of interesting issues. Notably, even when in each round the market has a unique equilibrium, over time this equilibrium becomes a dynamic object. As such, exact competitive equilibria can rarely or never be reached. Instead, we consider how tatonnement can trace the equilibrium by maintaining a small distance (in terms of suitably defined notions of distance), which results in *approximate* clearing conditions. For PRD, we apply a similar approach based on adaptation of the allocation of goods.

More formally, we study the prominent class of Fisher markets, in which agent utilities exhibit constant elasticity of substitution (CES). We analyze the impact of changes in supply of goods, budgets of agents, and their utility parameters. The agent interaction approaches equilibrium conditions. Since equilibrium is moving, prices and allocations chase the equilibrium point over time. Our analyses provides distance bounds, which can be seen as a quantification of the extent of out-of-equilibrium trade due to the interplay of market variation and adaptation of agents.

Technically, the majority of our analyses is concerned with quantifying the impact of variation in market parameters on several potential functions that guarantee convergence of the dynamics. The results then follow by a combination with the convergence guarantees for static markets. In fact, this approach constitutes two powerful yet handy frameworks to analyze a large variety of protocols and dynamics that are well-understood in static systems, when these systems become subject to dynamic variation.

Contribution and Outline. After presenting necessary preliminaries in Section 2, we describe in Section 3 the general model for dynamic CES Fisher markets and a general convergence result. In the subsequent sections, we discuss the insightful case of CES markets with gross-substitute condition. In these markets, the total misspending (absolute excess demand times price) over all goods is a natural parameter to quantify the violation of market clearing conditions. Moreover, one round of tatonnement updates in static markets is known to decrease misspending by a multiplicative factor [9]. In Sections 3.1 and 3.2, we consider markets where the supply of goods, the budgets of agents, and the utility function of the agents are subject to dynamic variation, respectively. We quantify the impact of variation on the misspending in the market. These bounds reveal that the change is often a rather mild additive change in misspending. Together with the fact that tatonnement decreases the misspending multiplicatively, we see that the price adaptation is indeed able to incorporate and adapt to the changes quickly. Overall, the dynamics can trace the equilibrium point up to a distance that evolves from the changes in a small number of recent rounds.

We provide similar results for Fisher markets with *any* CES utilities (including those not satisfying gross-substitute condition) based on a convex potential function [8] in the full version [11]. A slight disadvantage is that this potential function does not have an equally intuitive interpretation as the misspending function.

The technique we apply for markets can be executed much more generally for a class of dynamical systems, which we formulate as a framework in Section 4. These systems have a set of control parameters (e.g., prices in markets, or strategic decisions in games) and system parameters (e.g., supplies or utilities in markets, or payoff values in games). Moreover, these systems admit a Lyapunov function, and a round-based adaptation process for the control parameters (e.g., tatonnement in markets, or best-response dynamics in classes of games) that multiplicatively decreases the Lyapunov function in a single round. Our results provide a bound on the value of the Lyapunov function when the system parameters are subject to dynamic changes. It seems likely that a similar analysis based on our techniques can be conducted for many more sophisticated systems with significantly more complex dynamics. To demonstrate this, we discuss an example of such system, network load balancing, in Section 7. Another example about minimization of strongly convex functions is deferred to the full version [11].

In Section 5, we use another generalization of the technique, which is based on Bregman divergence, to show that PRD can successfully trace equilibrium in gross-substitute CES Fisher markets. Again, we can extend this to a general framework of dynamical systems governed by progress in Bregman divergence; see Section 6.

All missing proofs can be found in the full version [11].

Related Work. Competitive equilibrium and tatonnement date back to Walras [27] in 1874. The existence of equilibrium was established in a non-constructive way for a general setting by Arrow and Debreu [2] in 1954. Computation of equilibrium has been a central subject in general equilibrium theory. In the past 15 years, there has been impressive progress on devising efficent algorithms for computing equilibria, e.g., using network-flow algorithms [3, 4, 15–17, 25], the ellipsoid method [19] or the interior point method [29].

Decentralized adaptation processes such as tatonnement are important due to their simple nature and plausible applicability in real markets. Tatonnement is broadly defined as a process that increases (resp. decreases) the price of a good if the demand for the good is more (resp. less) than the supply. The price updates are *distributed*, since the price adjustment for each good is based on its own excess demand, independent of the demands for other goods.

Arrow, Block and Hurwitz [1] showed that a continuous version of tatonnement converges to an equilibrium for markets satisfying the weak gross substitutes (WGS) property. The recent algorithmic advances provide new insights in analyzing tatonnement [8, 12]. Cole and Fleischer [13] proposed the *ongoing market model*, in which *warehouses* are introduced to allow out-of-equilibrium trade, and prices are updated in tatonnement-style *asynchronously*, to provide an *in-market* process which might capture how real markets work. There has been significant recent interest in further aspects of ongoing markets or asynchronous tatonnement [6, 7, 9, 14].

In contrast, proportional response dynamics are a class of distributed algorithms that originated in the literature on network bandwidth sharing. These dynamics work without prices and come with convergence guarantees in classes of static network exchange, where goods have a uniform value [28, 30]. For CES Fisher markets, these dynamics can be cast as a form of mirror descent [5, 10].

Notions of games and markets with perturbation and dynamic changes are receiving increased interest in algorithmic game theory. For example, recent work has started to quantify the average performance of simple auctions and regret-learning agents in combinatorial auctions with dynamic buyer population [18, 24]. In these scenarios, however, equilibria are probabilistic objects and convergence in the static case can only be shown in terms of regret on average in hindsight. Moreover, the main goal is to bound the price of anarchy.

2 PRELIMINARIES

Fisher Markets. In a Fisher market, there are n *goods* and m *agents* (buyers). Each agent i has an amount b_i of *budget*, and she has a *utility function* u_i representing her preference. For bundles $\mathbf{x}_i^1 = (x_{ij}^1)_{j=1,\dots,n}$ and $\mathbf{x}_i^2 = (x_{ij}^2)_{j=1,\dots,n}$, if $u_i(x_{i1}^1, \cdots, x_{in}^1) > u_i(x_{i1}^2, \cdots, x_{in}^2)$, then she prefers $\mathbf{x}_i^1$ to $\mathbf{x}_i^2$. We denote the vector of budgets by $\mathbf{b} = (b_i)_{i=1,\dots,m}$ and the vector of utility functions by $\mathbf{u} = (u_i)_{i=1,\dots,m}$. Let $B = \sum_i b_i$ be the total budget in the market.

Given a vector $\mathbf{p} = (p_j)_{j=1,\dots,n}$ of (per-unit) *prices* for each good, agent i requests a *demand bundle* of goods that maximizes her utility function subject to the budget constraint: $\hat{\mathbf{x}}_i = \arg\max\{u_i(\mathbf{x}_i) : \sum_{j=1}^n x_{ij} \cdot p_j \leq b_i\}$. In general, the arg max is a set of bundles. In this paper, we concern strictly concave utility function only, for which there is a *unique* demand bundle.

The sum of amount of good j purchased by all agents is the *demand for good j*, denoted by $x_j = \sum_{i=1}^m \hat{x}_{ij}$. The *supply of good j* is w_j, and we set $\mathbf{w} = (w_j)_{j=1,\dots,n}$. Let $\mathbf{z} = (z_j)_{j=1,\dots,n}$ be the vector of *excess demand*, i.e., demand minus supply: $z_j = x_j - w_j$.

A pair $(\mathbf{x}^*, \mathbf{p}^*)$ is a *competitive* or *market equilibrium* if (1) each vector $\mathbf{x}_i^*$ is a demand bundle of agent i at prices $\mathbf{p}^*$, (2) for each good j with $p_j^* > 0$, demand is equal to supply (i.e., $z_j = 0$), and (3) for each good j with $p_j^* = 0$, demand is at most supply (i.e., $z_j \leq 0$). $\mathbf{p}^*$ is often called a *market clearing price vector.*

CES Utility Functions. A prominent class of utility functions is the Constant Elasticity of Substitution (CES) utility functions. They have the form $u_i(\mathbf{x}_i) = \left(\sum_{j=1}^n a_{ij} \cdot (x_{ij})^\rho\right)^{1/\rho}$, where $1 \geq \rho > -\infty$ and all $a_{ij} \geq 0$.

For $\rho < 1$ and $\rho \neq 0$, agent i's demand for good j is

$$\hat{x}_{ij} = b_i \cdot \frac{(a_{ij})^{1-c}(p_j)^{c-1}}{\sum_{k=1}^n (a_{ik})^{1-c}(p_k)^c}, \text{ where } c = \frac{\rho}{\rho - 1} .$$

To avoid algebraic clutter, we assume the parameter ρ of each agent is the same, but we note that all our analyses can be easily extended to cover distinct ρ_i scenarios.

Dynamic Markets. For CES Fisher markets, tatonnement is known to converge quickly to equilibrium under static market conditions. We here consider a dynamic market where in the beginning of each round t the tatonnement dynamic proposes a price vector $\mathbf{p}^t$. Dynamic market parameters like budgets $\mathbf{b}^t$, supplies $\mathbf{w}^t$ and utility functions $\mathbf{u}^t$ are manifested, which can be different from their values in previous rounds. Each agent requests a demand bundle based on the price vector $\mathbf{p}^t$ and market $\mathcal{M}^t = (\mathbf{u}^t, \mathbf{b}^t, \mathbf{w}^t)$, which yields a vector of excess demands $\mathbf{z}^t$. Then the system proceeds to the next round $t + 1$.

We first provide a basic insight that lies at the core of the analysis and manages to lift convergence results for a class of static markets to a bound for dynamic markets from that class. Formally, assume that the following properties hold:

Potential: There is a non-negative potential function $\Phi(\mathcal{M}, \mathbf{p})$, for every market $\mathcal{M} = (\mathbf{u}, \mathbf{b}, \mathbf{w})$ and every price vector $\mathbf{p}$. It holds $\Phi(\mathcal{M}, \mathbf{p}) = 0$ if and only if $\mathbf{p}$ is a market clearing price vector for market $\mathcal{M}$.

Price-Improvement: The price dynamics satisfy $\Phi(\mathcal{M}, \mathbf{p}^t) \leq (1-\delta) \cdot \Phi(\mathcal{M}, \mathbf{p}^{t-1})$, for some $1 \geq \delta > 0$ and every market $\mathcal{M}$.

Market-Perturbation: The market dynamics satisfy $\Phi(\mathcal{M}^t, \mathbf{p}) \leq \Phi(\mathcal{M}^{t-1}, \mathbf{p}) + \Delta^t$, for some values $\Delta^t \geq 0$ and every $\mathbf{p}$.

Proposition 2.1. *Suppose the price and market dynamics satisfy the Potential, Price-Improvement, and Market-Perturbation properties.*

Then $$\Phi(\mathcal{M}^T, \mathbf{p}^T) \leq (1-\delta)^T \cdot \Phi(\mathcal{M}^0, \mathbf{p}^0) + \sum_{t=1}^{T} (1-\delta)^{T-t} \Delta^t .$$

Let $\Delta = \max_{t=1,\dots,T} \Delta^t$*, then it follows for any* $t \leq T$*,*

$$\Phi(\mathcal{M}^T, \mathbf{p}^T) \leq \sum_{\tau=t+1}^{T} (1-\delta)^{T-\tau} \Delta^\tau + \frac{(1-\delta)^{T-t}}{\delta} \cdot \Delta + (1-\delta)^T \cdot \Phi(\mathcal{M}^0, \mathbf{p}^0) .$$

The proof follows by a direct application of the three properties. We prove the proposition for a much more general class of dynamic systems with Lyapunov functions in Section 4.

Consider the three terms in the latter bound for Φ. The first term captures the impact of *recent* changes to the market. The second term bounds the effect of all *older* changes. The third term decays exponentially over time. Hence, when the process runs long enough, the potential is only affected by *recent changes* of the market, while all older changes can be accumulated into a constant term based on Δ and δ. Intuitively, the price dynamics follows the evolution of the equilibrium up to a "distance" of Δ/δ in the potential function value. Hence, if market perturbation Δ is small and price improvement δ is large, the process succeeds to maintain market clearing conditions almost exactly.

3 DYNAMIC MARKETS VIA MISSPENDING

We here describe our techniques for CES markets $\mathcal{M}$ with gross-substitutes property, i.e., when all buyers have CES utilities with $1 > \rho > 0$.

The tatonnement process we analyze here updates prices in each round based on the excess demand in the last round, i.e.,

$$p_j^t \;\leftarrow\; p_j^{t-1} \cdot \left[1 + \lambda \cdot \min\left(\frac{x_j^{t-1} - w_j}{w_j} \;,\; 1 \right) \right] \;, \qquad (1)$$

where $\lambda < 1$ is a parameter depending on ρ. The *misspending potential function* [9, 14] is

$$\Phi_{\mathsf{MS}}(\mathcal{M}, \mathbf{p}) \;=\; \sum_{j=1}^{n} p_j \cdot \left| z_j \right| .$$

The tatonnement process is known to have the Price-Improvement property based on the misspending potential function Φ_{MS} in CES markets with $1 > \rho > 0$. More formally, if $\lambda \leq \Theta(1 - \rho_{\max})$, then there exists $1 \geq \delta = \delta(\lambda) > 0$ such that $\Phi_{\mathsf{MS}}(\mathcal{M}, \mathbf{p}^t) \leq (1-\delta) \cdot \Phi_{\mathsf{MS}}(\mathcal{M}, \mathbf{p}^{t-1})$. [14]

In the subsequent subsections, we will present upper bounds on the misspending. We discuss how to interpret them. For a static market $\mathcal{M}$ and a given starting point $\mathbf{p}^0$, there exists a positive lower bound on all prices appeared throughout the tatonnement process. To see why, intuitively, when the price of a good is very low, its demand has to be very large[1] and hence its price must subsequently increase. See [13, 14] for how to explicitly derive such a lower bound. When the market becomes dynamic but all parameters stay within an appropriate range, a similar lower bound can be derived; we denote it by $P_{\min}$. Then $\Phi_{\mathsf{MS}}(\mathcal{M}, \mathbf{p}) \leq c$ implies

$$\text{for each good } j,\; |z_j| \;\leq\; c/p_j \;\leq\; c/P_{\min}.$$

Thus, an upper bound on the misspending can be converted to an upper bound on absolute excess demands via the above inequality.

3.1 Dynamic Supply and Budgets

Dynamic Supply. Let us first analyze the impact of changing supply on tatonnement dynamics and market clearing conditions. We normalize the initial supply $w_j^1 = 1$ for each good j. Suppose that the supplies are then changed additively[2] by $\boldsymbol{\varepsilon}^t = (\varepsilon_1^t, \varepsilon_2^t, \cdots, \varepsilon_n^t)$ at time t. We parametrize our bounds using the maximum supply change $\kappa = \max_t \|\boldsymbol{\varepsilon}^t\|_1$.

Assumption 1. Every price is universally bounded by some time-independent constant P, i.e., for any j and any time t, we have $p_j^t \leq P$.

Assumption 1 is made for technical reasons, but it is simple to satisfy by constant parameters of the market. For example, if all initial prices are at most B, then since $\lambda < 1$ Assumption 1 holds with $P = 2B$. The main result in this section is as follows.

Proposition 3.1. *For any $t \leq T$,*

$$\Phi_{\mathsf{MS}}(\mathcal{M}^T, \mathbf{p}^T) \;\leq P \cdot \left(\sum_{\tau=t+1}^{T} (1-\delta)^{T-\tau} \|\boldsymbol{\varepsilon}^\tau\|_1 \;+\; \frac{(1-\delta)^{T-t}}{\delta} \cdot \kappa \right) + (1-\delta)^T \cdot \Phi_{\mathsf{MS}}(\mathcal{M}^0, \mathbf{p}^0) \;.$$

Proof. Consider the misspending potential Φ_{MS}. Tatonnement satisfies the Price-Improvement property. Hence, to show the result, we establish the Market-Perturbation property.

Note that the misspending potential can be written as

$$\Phi_{\mathsf{MS}}(\mathcal{M}^t, \mathbf{p}^t) \;=\; \sum_{j=1}^{n} p_j^t \cdot \left| x_j^t - 1 - \sum_{\tau=1}^{t} \varepsilon_j^\tau \right| \;.$$

Hence, by the triangle inequality and Assumption 1,

$$\begin{aligned} \Phi_{\mathsf{MS}}(\mathcal{M}^t, \mathbf{p}^t) \;&=\; \Phi_{\mathsf{MS}}(\mathcal{M}^{t-1}, \mathbf{p}^t) \;+\; \sum_{j=1}^{n} p_j^t \cdot \left|\varepsilon_j^t\right| \\ &\leq\; \Phi_{\mathsf{MS}}(\mathcal{M}^{t-1}, \mathbf{p}^t) \;+\; P \cdot \|\boldsymbol{\varepsilon}^t\|_1 \;. \end{aligned}$$

We are done by Proposition 2.1, with $\Delta^t = P \cdot \|\boldsymbol{\varepsilon}^t\|_1$ and $\Delta = P\kappa$. ◻

A Remark. If the supplies of all goods shrink *multiplicatively* by the same factor of $(1-\beta)$, then in CES markets, the equilibrium price of every good increases by a factor of $(1-\beta)^{-1}$. However, the tatonnement update rule allows the current price to be raised by a factor of at most $(1+\lambda)$ per time step. Thus, for plausible *tracing* of equilibrium, λ must satisfy $(1+\lambda) > (1-\beta)^{-1}$.

Dynamic Budgets. We now analyze the impact of changing buyer budgets on tatonnement dynamics and market clearing conditions. Starting from the initial budgets, the budgets are then changed additively by $\boldsymbol{\varepsilon}^t = (\varepsilon_1^t, \varepsilon_2^t, \cdots, \varepsilon_m^t)$ at time t. We parametrize our bounds using the maximum budget change $\eta = \max_t \|\boldsymbol{\varepsilon}^t\|_1$.

Proposition 3.2. *For any $t \leq T$,*

$$\Phi_{\mathsf{MS}}(\mathcal{M}^T, \mathbf{p}^T) \;\leq\; \sum_{\tau=t+1}^{T} (1-\delta)^{T-\tau} \|\boldsymbol{\varepsilon}^\tau\|_1 \;+\; \frac{(1-\delta)^{T-t}}{\delta} \cdot \eta \;+\; (1-\delta)^T \cdot \Phi_{\mathsf{MS}}(\mathcal{M}^0, \mathbf{p}^0) \;.$$

3.2 Dynamic Buyer Utility

We analyze the impact of changing the parameters a_{ij} in the CES utility functions on tatonnement dynamics and market clearing conditions. Starting from the initial utility values, each a_{ij} can in each round t be changed by some multiplicative factor γ_{ij}^t. Let $\gamma^t = \max_{i,j}((\gamma_{ij}^t)^{\frac{1}{1-\rho}}, (1/\gamma_{ij}^t)^{\frac{1}{1-\rho}})$ and $\gamma = \max_t \gamma^t$.

Proposition 3.3. *For any $t \leq T$,*

$$\Phi_{\mathsf{MS}}(\mathcal{M}^T, \mathbf{p}^T) \;\leq\; (1-\delta)^T \cdot \Phi_{\mathsf{MS}}(\mathcal{M}^0, \mathbf{p}^0) + B \cdot \left(\sum_{\tau=t+1}^{T} (1-\delta)^{T-\tau} \cdot \frac{2(\gamma^\tau - 1)}{\gamma^\tau + 1} \;+\; \frac{(1-\delta)^{T-t}}{\delta} \cdot \frac{2(\gamma - 1)}{\gamma + 1} \right) \;.$$

Proof. To show the result, we establish the Market-Perturbation property. Note that the misspending potential can be given by

$$\begin{aligned} \Phi_{\mathsf{MS}}(\mathcal{M}^t, \mathbf{p}^t) \;&=\; \sum_{j=1}^{n} p_j^t \cdot \left| x_j^t - w_j \right| \\ &=\; \sum_{j=1}^{n} p_j^t \cdot \left| \sum_{i=1}^{m} b_i \cdot \frac{(a_{ij} \prod_{\tau=1}^{t} \gamma_{ij}^\tau)^{1-c} (p_j)^{c-1}}{\sum_{k=1}^{n} (a_{ik} \prod_{\tau=1}^{t} \gamma_{ik}^\tau)^{1-c} (p_k)^c} - w_j \right| \;. \end{aligned}$$

[1] This is not true in general, but it holds for the markets we concern.
[2] We here study additive change for mathematical convenience. The bounds can be adjusted to hold accordingly for multiplicative change.

Using $a_{ij}^{t-1} = a_{ij} \prod_{\tau=1}^{t-1} \gamma_{ij}^{\tau}$ we derive

$$
\begin{aligned}
\Delta^t &= \Phi_{\mathrm{MS}}(\mathcal{M}^t, \mathbf{p}) - \Phi_{\mathrm{MS}}(\mathcal{M}^{t-1}, \mathbf{p}) \\
&= \sum_{j=1}^{n} p_j \cdot \left(\left| \sum_{i=1}^{m} b_i \cdot \frac{(a_{ij}^{t-1}\gamma_{ij}^t)^{1-c}(p_j)^{c-1}}{\sum_{k=1}^{n}(a_{ik}^{t-1}\gamma_{ik}^t)^{1-c}(p_k)^c} - w_j \right| \right. \\
&\qquad\qquad \left. - \left| \sum_{i=1}^{m} b_i \cdot \frac{(a_{ij}^{t-1})^{1-c}(p_j)^{c-1}}{\sum_{k=1}^{n}(a_{ik}^{t-1})^{1-c}(p_k)^c} - w_j \right| \right) \\
&\le \sum_{j=1}^{n} p_j \cdot \sum_{i=1}^{m} b_i \cdot \left| \frac{(a_{ij}^{t-1}\gamma_{ij}^t)^{1-c}(p_j)^{c-1}}{\sum_{k=1}^{n}(a_{ik}^{t-1}\gamma_{ik}^t)^{1-c}(p_k)^c} - \frac{(a_{ij}^{t-1})^{1-c}(p_j)^{c-1}}{\sum_{k=1}^{n}(a_{ik}^{t-1})^{1-c}(p_k)^c} \right| \\
&= \sum_{i=1}^{m} b_i \cdot \sum_{j=1}^{n} \left| \frac{(a_{ij}^{t-1}\gamma_{ij}^t)^{1-c}\, p_j^c}{\sum_{k=1}^{n}(a_{ik}^{t-1}\gamma_{ik}^t)^{1-c}\, p_k^c} - \frac{(a_{ij}^{t-1})^{1-c}\, p_j^c}{\sum_{k=1}^{n}(a_{ik}^{t-1})^{1-c}\, p_k^c} \right| .
\end{aligned}
$$

For the rest of the proof, we construct an upper bound on the difference of two fractions. Fix a buyer i, set $\alpha_j = \frac{(a_{ij}^{t-1})^{1-c}\, p_j^c}{\sum_{k=1}^{n}(a_{ik}^{t-1})^{1-c}\, p_k^c}$, $\beta_j = (\gamma_{ij}^t)^{1-c}$, $\mu = \gamma^t$ and observe $\mu \ge \beta_j \ge 1/\mu$. We define

$$\Delta s_j = \left| \frac{\alpha_j \beta_j}{\sum_k \alpha_k \beta_k} - \alpha_j \right| .$$

Lemma 3.4. *There exists a vector* $(\beta_1', \ldots, \beta_n')$ *with*

$$\beta_j' = \begin{cases} \mu & \text{if } \frac{\alpha_j \beta_j'}{\sum_k \alpha_k \beta_k'} \ge \alpha_j \\ 1/\mu & \text{otherwise.} \end{cases}$$

such that $$\sum_j \Delta s_j \le \sum_j \left| \frac{\alpha_j \beta_j'}{\sum_k \alpha_k \beta_k'} - \alpha_j \right| .$$

Now let β' be the vector defined in the above lemma, let $S = \{j : \beta_j' = \mu\}$ and $R = G \setminus S$. Using $\alpha_S = \sum_{j \in S} \alpha_j$, we obtain

$$
\begin{aligned}
\sum_j \Delta s_j &\le \left(\sum_{j \in S} \frac{\alpha_j \mu}{\sum_{k \in S} \alpha_k \mu + \sum_{i \in R} \alpha_i/\mu} - \sum_{j \in S} \alpha_j \right) \\
&\qquad + \left(\sum_{j \in R} \alpha_j - \sum_{j \in R} \frac{\alpha_j/\mu}{\sum_{k \in S} \alpha_k \mu + \sum_{i \in R} \alpha_i/\mu} \right) \\
&= \left(\frac{\mu \alpha_S}{\mu \alpha_S + (1-\alpha_S)/\mu} - \alpha_S \right) + \left(1 - \alpha_S - \frac{(1-\alpha_S)/\mu}{\mu \alpha_S + (1-\alpha_S)/\mu} \right) \\
&= 1 - 2\alpha_S + \frac{\left(\mu + \frac{1}{\mu}\right)\alpha_S - \frac{1}{\mu}}{\left(\mu - \frac{1}{\mu}\right)\alpha_S + \frac{1}{\mu}} .
\end{aligned}
$$

The RHS is maximized at $\alpha_S = \frac{1}{\mu+1}$, attaining value of $\frac{2(\mu-1)}{\mu+1}$. We are done by Proposition 2.1, with $\Delta^t \le B \cdot \frac{2(\gamma^t - 1)}{\gamma^t + 1}$ and $\Delta \le B \cdot \frac{2(\gamma-1)}{\gamma+1}$. □

4 PARAMETRIZED LYAPUNOV DYNAMICAL SYSTEMS

In this section, we prove a general theorem, which includes as special case the bound shown for markets in Proposition 2.1. Our focus here are dynamical systems, in which time is discrete and represented by non-negative integers.

We assume that the dynamical system can be described by two sets of parameters. There is a set of *control variables* that can be adjusted by an algorithm or a protocol. In addition, there is a set of *system parameters* that can change in each round in an adversarial way. For example, in our analyses of markets, the control variables are prices, whereas system parameters are supplies of goods, budgets of agents, or utility parameters. As another example, in games the control variables could be the strategy choices of agents, whereas system parameters are payoff values of states. A further example: control variables could also be bird headings in a bird flock, while system parameters are wind direction and velocity.

The classical theory of dynamical systems often studies the behaviour of systems with static system parameters. However, dynamical systems with varying system parameters often arise in practice (see Section 7). Here, we propose a simple framework to analyze Lyapunov dynamical systems with varying system parameters. More formally, the dynamical system L is described by an initial *control variable vector* $\mathbf{p}^0 \in \mathbb{R}^n$ and an *evolution rule* $F : \mathbb{R}^n \to \mathbb{R}^n$, which specifies how the control variables are adjusted. For each time $t \ge 1$, we have $\mathbf{p}^t = F(\mathbf{p}^{t-1})$.

The system L is called a *Lyapunov dynamical system* (LDS) if it admits a Lyapunov function $G : \mathbb{R}^n \to \mathbb{R}^+$ such that

(a) for every fixed point (equilibrium) $\mathbf{p}$ of F, i.e., $F(\mathbf{p}) = \mathbf{p}$, it holds $G(\mathbf{p}) = 0$;
(b) for every $\mathbf{p} \in \mathbb{R}^n$, it holds $G(F(\mathbf{p})) \le G(\mathbf{p})$, while the equality holds if and only if $\mathbf{p}$ is a fixed point.

An LDS L is called *linearly converging* (LCLDS) if it further satisfies

(c) there exists a *decay parameter* $\delta = \delta(L) > 0$ such that for any $\mathbf{p} \in \mathbb{R}^n$, $G(F(\mathbf{p})) \le (1-\delta) \cdot G(\mathbf{p})$.

Let $\mathcal{L}$ be a family of dynamical systems, while each dynamical system $L_{\mathbf{s}} \in \mathcal{L}$ is specified by a *system parameter* vector $\mathbf{s} \in \mathbb{R}^d$. The family $\mathcal{L}$ is called a family of *parametrized, linearly converging LDS* (PLCLDS) if each $L_{\mathbf{s}} \in \mathcal{L}$ is an LCLDS and $\delta(\mathcal{L}) = \inf_{L_{\mathbf{s}} \in \mathcal{L}} \delta(L_{\mathbf{s}}) > 0$. For each $L_{\mathbf{s}}$, we denote its evolution rule by $F_{\mathbf{s}}$ and its Lyapunov function by $G_{\mathbf{s}}$.

In many scenarios, particularly in agent-based dynamical systems, the control variables $\mathbf{p}$ change by the evolution rule that expresses, e.g., the sequential decisions of the agents, but the system parameters $\mathbf{s}$ can change in an exogenous (or even adversarial) way. However, in many cases the impact of changes in a single time step is rather mild. The following theorem states our *recovery result* by relating the Lyapunov value to the magnitude of changes in each step. Intuitively, it characterizes the "distance" that the evolution rule maintains to a fixed point over the course of the dynamics.

Theorem 4.1. *Let $\mathcal{L}$ be a PLCLDS with $\delta \equiv \delta(\mathcal{L}) > 0$, let $\mathbf{s}^0, \mathbf{s}^1, \ldots, \mathbf{s}^T$ denote the system parameter vectors at times $0, 1, \cdots, T$, respectively, and let $\Phi(\mathbf{s}^t, \mathbf{p}^t) = G_{\mathbf{s}^t}(\mathbf{p}^t)$. Suppose that for every $t = 1, \ldots, T$ the system parameters $\mathbf{s}^{t-1}, \mathbf{s}^t \in \mathbb{R}^d$ invoke a change such that for every $\mathbf{p} \in \mathbb{R}^n$, we have $\Phi(\mathbf{s}^t, \mathbf{p}) \le \Phi(\mathbf{s}^{t-1}, \mathbf{p}) + \Delta^t$. The initial control variable vector is denoted by $\mathbf{p}^0$, and the system* *evolves such that for every $t \geq 1$ we have $\mathbf{p}^t = F_{\mathbf{s}_{t-1}}(\mathbf{p}^{t-1})$. Then*

$$\Phi(\mathbf{s}^T, \mathbf{p}^T) \leq (1-\delta)^T \cdot \Phi(\mathbf{s}^0, \mathbf{p}^0) + \sum_{t=1}^{T} (1-\delta)^{T-t} \cdot \Delta^t .$$

Let $\Delta = \max_{t=1,\ldots,T} \Delta^t$, then for any $t \leq T$,

$$\Phi(\mathbf{s}^T, \mathbf{p}^T) \leq \sum_{\tau=t+1}^{T} (1-\delta)^{T-\tau} \Delta^\tau + \frac{(1-\delta)^{T-t}}{\delta} \cdot \Delta + (1-\delta)^T \cdot \Phi(\mathbf{s}^0, \mathbf{p}^0) .$$

Proof. For any time $t \geq 1$,

$$\begin{aligned}\Phi(\mathbf{s}^t, \mathbf{p}^t) &= G_{\mathbf{s}_t}(\mathbf{p}^t) \leq G_{\mathbf{s}_{t-1}}(\mathbf{p}^t) + \Delta^t \\ &= G_{\mathbf{s}_{t-1}}(F_{\mathbf{s}_{t-1}}(\mathbf{p}^{t-1})) + \Delta^t \\ &\leq (1-\delta) \cdot G_{\mathbf{s}_{t-1}}(\mathbf{p}^{t-1}) + \Delta^t = (1-\delta) \cdot \Phi(\mathbf{s}^{t-1}, \mathbf{p}^{t-1}) + \Delta^t.\end{aligned}$$

Iterating the above recurrence yields the first result. The second result follows from the calculation below:

$$\begin{aligned}\sum_{\tau=1}^{t} (1-\delta)^{T-\tau} \Delta^\tau &\leq \Delta (1-\delta)^T \cdot \sum_{\tau=1}^{t} \left(\frac{1}{1-\delta}\right)^\tau \\ &< \Delta \cdot \frac{(1-\delta)^T}{\delta} \cdot \left(\frac{1}{1-\delta}\right)^t . \qquad \square\end{aligned}$$

In the scenarios where $\sum_{t=1}^{T} \Delta^t = O(T^\alpha)$ for small constant α, we have the following corollary.

Corollary 4.2. *In the setting of Theorem 4.1, if $\sum_{t=1}^{T} \Delta^t = O(T^\alpha)$ for some constant $\alpha > 0$, then for any constant $\beta > 0$,*

$$\Phi(\mathbf{s}^T, \mathbf{p}^T) \leq \sum_{\tau=T-\left\lceil \frac{\alpha+\beta}{\delta} \log T \right\rceil + 1}^{T} \Delta^\tau + O(T^{-\beta}) + (1-\delta)^T \cdot \Phi(\mathbf{s}^0, \mathbf{p}^0).$$

As $T \to \infty$, the last two terms of the above inequality diminish. The bound is dominated by the first term, which describes the impact of the changes in the *recent* $O\left(\frac{\log T}{\delta}\right)$ steps.

5 PROPORTIONAL RESPONSE DYNAMICS

In the Fisher market setting, the general protocol of proportional response dynamics (PRD) is as follows. In each round, each buyer i splits her budget b_i among the n goods according to some rule, and sends the bids to the sellers of the corresponding goods. Based on the bids gathered from all buyers, the seller of each good j send back (simple) signals to buyers, which are then used by buyers for updating their bids in the next round. We summarize the notation and results we need from [10] below. When buyer i splits her budget b_i among the n goods, let b_{ij} denote the spending by her on good j. Let $\mathbf{B}$ denote $\{b_{ij}\}_{i\in[m], j\in[n]}$. Let $p_j := \sum_i b_{ij}$.

Consider the substitute domain, i.e., when the ρ_i parameters of all buyers are strictly between 0 and 1. In each round, the seller of good j distributes the good among buyers in proportion to the bids received, and then after receiving the goods, each buyer splits her budget in proportion to the utility generated from the quantity of each good received. More formally, let $p_j^t = \sum_i b_{ij}^t$, then the update rule is

$$b_{ij}^{t+1} = b_i \cdot a_{ij} \left(\frac{b_{ij}^t}{p_j^t}\right)^{\rho_i} \Bigg/ \left(\sum_k a_{ik} \left(\frac{b_{ik}^t}{p_k^t}\right)^{\rho_i}\right). \qquad (2)$$

The *Kullback-Leibler (KL) divergence* is similar to a distance measure. For vectors $\mathbf{x}$ and $\mathbf{y}$ such that $\sum_j x_j = \sum_j y_j$, the explicit formula is $\mathsf{KL}(\mathbf{x}||\mathbf{y}) := \sum_j x_j \cdot \ln \frac{x_j}{y_j}$. The above update rule is equivalent to mirror descent w.r.t. the KL divergence (but with different step sizes for different buyers) of the function:

$$g(\mathbf{B}) = -\sum_{ij} \frac{b_{ij}}{\rho_i} \log \frac{a_{ij}(b_{ij})^{\rho_i - 1}}{(p_j)^{\rho_i}} , \qquad (3)$$

defined on the domain $C = \left\{\mathbf{B} \,\middle|\, \forall i,\ \sum_j b_{ij} = b_i \text{ and } \forall i,j,\ b_{ij} \geq 0\right\}$.

For our purpose, it suffices to know that any equilibrium $\mathbf{B}^* \in C$ of PRD corresponds to a minimum point of g. The market potential with proportional response dynamics will be defined as:

$$G(\mathbf{B}) = g(\mathbf{B}) - g(\mathbf{B}^*) \qquad (4)$$

Cheung, Cole and Tao [10] show that for positive constants $q_1 < q_2$ (which depend on the maximum and minimum values of ρ_i) the market potential in a static market is bounded by

$$G(\mathbf{B}^{t+1}) \leq q_1 \cdot \mathsf{KL}(\mathbf{B}^*, \mathbf{B}^t) - q_2 \cdot \mathsf{KL}(\mathbf{B}^*, \mathbf{B}^{t+1}),$$

and hence the following holds due to telescoping on the RHS:

$$\left(\frac{q_2}{q_1}\right)^{T-1} G(\mathbf{B}^T) \leq \sum_{t=0}^{T-1} \left(\frac{q_2}{q_1}\right)^t G(\mathbf{B}^{t+1}) \leq q_1 \cdot \mathsf{KL}(\mathbf{B}^*, \mathbf{B}^\circ).$$

Dividing both sides by $(q_2/q_1)^{T-1}$ shows that $G(\mathbf{B}^T)$ converges linearly with T.

In the rest of the section, we analyze the impact of changing utility functions and supplies on the convergence properties of proportional response dynamics. For the varying budgets case, the domain C varies too, prohibiting a similar analysis.

Dynamic Buyer Utilities. Starting with the initial utility parameters, suppose that each a_{ik} changes by a factor within $[e^{-\varepsilon}, e^{\varepsilon}]$. For a given budget allocation $\mathbf{B}$, let $G(\mathcal{M}^t, \mathbf{B})$ denote the market potential for the utility of the buyers in round t, and $\mathbf{B}^{t,*} \in C$ the allocation that minimizes $G(\mathcal{M}^t, \mathbf{B})$.

Proposition 5.1. *After T rounds, it holds that*

$$G(\mathcal{M}^T, \mathbf{B}^T) \leq q_1 \left(\frac{q_1}{q_2}\right)^{T-1} \cdot \mathsf{KL}(\mathbf{B}^{\circ,*}, \mathbf{B}^\circ) + \frac{q_2}{q_2 - q_1} \cdot \Delta ,$$

where

$$\Delta = \sum_i \left(\frac{b_i(e^{\kappa_i} - 1)}{1-\rho_i} \cdot \left|\rho_i \log\left(\frac{B}{b_i}\right) - \log\left(\min_{t,j} \frac{a_{ij}^t}{\sum_k a_{ik}^t}\right)\right| + \frac{2 b_i \varepsilon}{\rho_i}\right) ,$$

and $\kappa_i = 2\varepsilon\,(1 - c_i\,(3 - 2\min_i c_i))$, where $c_i = \rho_i/(\rho_i - 1)$.

Claim 5.2. *For any round $t \leq T$ it holds*

$$\begin{aligned}G(\mathcal{M}^{t+1}, B^{t+1}) &\leq q_1 \cdot \mathsf{KL}(\mathbf{B}^{t,*}, \mathbf{B}^t) - q_2 \cdot \mathsf{KL}(\mathbf{B}^{t+1,*}, \mathbf{B}^{t+1}) \\ &\quad + 2\sum_i \frac{b_i \varepsilon}{\rho_i} + \sum_i b_i(e^{\kappa_i} - 1) \cdot |\log C_i - \log \Pi_i| ,\end{aligned}$$

where $C_i = \left(\frac{B}{b_i}\right)^{\frac{\rho_i}{1-\rho_i}}$ *is a constant and* $\Pi_i = \left(\min_{t,j} \frac{a^t_{ij}}{\sum_k a^t_{ik}}\right)^{\frac{1}{1-\rho_i}}$.

Proof of Proposition 5.1. Now suppose Δ is given as in the proposition, then with the Claim 5.2 it follows that:

$$q_2 \cdot \mathsf{KL}(\mathbf{B}^{t+1,*}, \mathbf{B}^{t+1}) \leq q_1 \cdot \mathsf{KL}(\mathbf{B}^{t,*}, \mathbf{B}^{t}) + \Delta. \quad (5)$$

The potential of the market at round T can be bounded by

$$\begin{aligned} G(\mathcal{M}^T, \mathbf{B}^T) &\leq q_1 \cdot \mathsf{KL}(\mathbf{B}^{T-1,*}, \mathbf{B}^{T-1}) + \Delta \\ &\leq q_1 \left(\frac{q_1}{q_2} \cdot \mathsf{KL}(\mathbf{B}^{T-2,*}, \mathbf{B}^{T-2}) + \frac{\Delta}{q_2}\right) + \Delta \\ &\leq q_1 \cdot \left(\frac{q_1}{q_2}\right)^{T-1} \cdot \mathsf{KL}(\mathbf{B}^{0,*}, \mathbf{B}^{0}) + \frac{q_2}{q_2 - q_1} \cdot \Delta, \end{aligned}$$

where the inequalities follow by recursive application of (5). □

Dynamic Supplies. It turns out that the case with varying supplies can be reduced to the case with varying utility functions. To see why, note that the function g defined in (3) assumes that the supply of each good is normalized to be one unit. When the supply of good j is changed from 1 to e^{ϵ}, by performing a re-normalization of the supply, it is equivalent to changing a_{ij} to $a_{ij} \cdot e^{\epsilon \rho_i}$.

6 GENERALIZATION TO PLCLDS USING BREGMAN DIVERGENCE

In the last section, we analyse PRD using KL divergence by adapting the analyses by Cheung, Cole and Tao [10]. Indeed, they proposed a far more general framework for demonstrating linear convergence when the underlying function is *strongly Bregman convex*, a new generalization of the standard notion of strong convexity, in the context of mirror descent. We propose a variant of PLCLDS in dynamic environment setting based on their framework.

Let C be a compact and convex set. Given a differentiable convex function $h(\mathbf{x})$ with domain of C, the *Bregman divergence* generated by the *kernel h* is denoted by d_h, defined as

$$d_h(\mathbf{x}, \mathbf{y}) = h(\mathbf{x}) - [\, h(\mathbf{y}) + \langle \nabla h(\mathbf{y}) , \mathbf{x} - \mathbf{y} \rangle \,],$$

for all $\mathbf{x} \in C$ and $\mathbf{y} \in \mathsf{rint}(C)$, where $\mathsf{rint}(C)$ is the relative interior of C. A convex function f is (σ, L)-strongly Bregman convex w.r.t. Bregman divergence d_h if, $0 < \sigma \leq L$, and for any $\mathbf{y} \in \mathsf{rint}(C)$ and $\mathbf{x} \in C$,

$$\begin{aligned} & f(\mathbf{y}) + \langle \nabla f(\mathbf{y}) , \mathbf{x} - \mathbf{y} \rangle + \sigma \cdot d_h(\mathbf{x}, \mathbf{y}) \\ & \quad \leq f(\mathbf{x}) \leq f(\mathbf{y}) + \langle \nabla f(\mathbf{y}) , \mathbf{x} - \mathbf{y} \rangle + L \cdot d_h(\mathbf{x}, \mathbf{y}). \end{aligned}$$

Note that the KL divergence used in the analysis of proportional response dynamics in Section 5 is an instance of Bregman divergence where the kernel function is $h(\mathbf{x}) = \sum_j (x_j \cdot \ln x_j - x_j)$. If the system is static, the variant of PLCLDS satisfies properties (a) and (b), with property (c) replaced by the following new property: there exists positive numbers $q_1 < q_2$ such that for any $\mathbf{p}, \mathbf{p}^* \in C$,

$$G(F(\mathbf{p})) \leq q_1 \cdot d_h(\mathbf{p}^*, \mathbf{p}) - q_2 \cdot d_h(\mathbf{p}^*, F(\mathbf{p})). \quad (6)$$

The above property holds when, for instance, G is a (σ, L)-strongly Bregman convex function with minimum value zero, and F is a mirror descent update:

$$F(\mathbf{p}) = \arg\min_{\mathbf{p}'} \{ \langle \nabla G(\mathbf{p}) , \mathbf{p}' - \mathbf{p} \rangle + L \cdot d_h(\mathbf{p}', \mathbf{p}) \},$$

for which $q_1 = L - \sigma$ and $q_2 = L$. [10]

By a suitable telescoping with (6), it is easy to show that

$$G(\mathbf{p}^T) \leq q_1 \cdot (q_1/q_2)^{T-1} \cdot d_h(\mathbf{p}^*, \mathbf{p}^0),$$

where $\mathbf{p}^*$ is *any* fixed point (equilibrium) of the Lyapunov system. However, for system which is dynamic, we will need a modification of the above property, presented in the theorem below. The style of its proof is similar to those appear in Section 5.

Theorem 6.1. *Let $\mathcal{L}$ be a PLCLDS with $\delta \equiv \delta(\mathcal{L}) > 0$. Let $\mathbf{s}^0, \mathbf{s}^1, \ldots, \mathbf{s}^T$ and $\mathbf{p}^{*,0}, \mathbf{p}^{*,1}, \ldots, \mathbf{p}^{*,T}$ denote the sequence of system parameters and fixed points at times $0, 1, \cdots, T$, respectively and let $\Phi(\mathbf{s}^t, \mathbf{p}^t) = G_{\mathbf{s}^t}(\mathbf{p}^t)$. Suppose that for every $t = 1, \ldots, T$ the system parameters $\mathbf{s}^{t-1}, \mathbf{s}^t \in \mathbb{R}^d$ invoke a change such that the fixed points change from $\mathbf{p}^{*,t-1}$ to $\mathbf{p}^{*,t}$, and:*

$$\Phi(\mathbf{s}^t, \mathbf{p}^t) \leq q_1 \cdot d_h(\mathbf{p}^{*,t-1}, \mathbf{p}^{t-1}) - q_2 \cdot d_h(\mathbf{p}^{*,t}, \mathbf{p}^t) + \Delta^t,$$

If the initial control variable vector is denoted by $\mathbf{p}^0$, and the system evolves such that for every $t \geq 1$ we have $\mathbf{p}^t = F_{\mathbf{s}_{t-1}}(\mathbf{p}^{t-1})$, then

$$\Phi(\mathbf{s}^T, \mathbf{p}^T) \leq q_1 \cdot \left(\frac{q_1}{q_2}\right)^{T-1} d_h(\mathbf{p}^{*,0}, \mathbf{p}^0) + \sum_{i=0}^{T-1} \left(\frac{q_1}{q_2}\right)^i \Delta^{T-i}.$$

7 APPLICATION: LOAD BALANCING WITH DYNAMIC MACHINE SPEED

Consider a setting with n distinct machines all connected to each other to form an arbitrary network. For ease of notation, we label the machines as m_i for $i = 1$ to n. Each machine m_i can process jobs at speed s_i. Jobs/tasks, assumed to be infinitely divisible, of total weight M are arbitrarily distributed over the network. Our goal is to design a decentralized load balancing algorithm with the objective that the total processing time over all machines is minimized.

Algorithm 1: Diffusion

```
for t = 1 to T do
    for each machine m_i do
        f_i^(t) ← total processing time on m_i ;
        broadcast f_i^(t) to all j ∈ nbd(m_i) ;
        forall j ∈ nbd(m_i) do
            if f_i^(t) > f_j^(t) then
                send P_ij(f_i^(t) − f_j^(t))s_i load to j.;
```

Before proceeding, we set up some notation. $\mathbf{s}$ denotes the vector of machine speeds. $\boldsymbol{\ell}^{(t)} = (\ell_i^{(t)})_i$ denotes the vector of loads and $\mathbf{f}^{(t)} = (f_i^{(t)})_i$ the corresponding finishing times at round t. We assume throughout that the total load stays constant i.e. $\sum_i \ell_i^{(t)} = M$. For machine speed $\mathbf{s}$, $f^{*,\mathbf{s}}$ denotes the corresponding vector of finishing times in the balanced state, i.e. a state where the finishing time of all machines is the same.

Algorithm 1 is based on the diffusion principle [26], where if a machine has more jobs than its neighbours, then some jobs diffuse to the neighbour. In our context, since the goal is to equalize the finishing times of all machines, the number of jobs that diffuse is proportional to the difference in the finishing times. The proportionality constant depends on the connecting edge. Specifically, in the algorithm that follows we use a diffusivity matrix $\mathbf{P}$ satisfying the following conditions: (a) $P_{ii} \geq 1/2$ (b) $P_{ij} > 0$ iff (i, j) is an edge in G. (c) $\mathbf{P}$ is symmetric and stochastic, i.e., for every machine m_i, $\sum_j P_{ij} = 1$.

If each machine m_i uses the load balancing protocol as described above, then the finishing time of machine m_i at time $t+1$ is known to satisfy $\mathbf{f}^{(t+1)} = \mathbf{P}\mathbf{f}^{(t)}$. Further, for the balanced state $\mathbf{f}^*$ (when the finishing times of all machines are equal), $\mathbf{P}\mathbf{f}^* = \mathbf{f}^*$. If we denote the *error* in round $t+1$ by $\mathbf{e}^{(t+1)} := \mathbf{f}^{(t+1)} - \mathbf{f}^*$, then:

$$\mathbf{e}^{(t+1)} \;=\; \mathbf{f}^{(t+1)} - \mathbf{f}^* \;=\; \mathbf{P}(\mathbf{f}^{(t)} - \mathbf{f}^*) \;=\; \mathbf{P}\mathbf{e}^{(t)},$$

i.e., the same transformations apply to the error vector as well. Since $\mathbf{P}$ is a symmetric matrix, it has n eigenvalues $\lambda_1, \lambda_2 \cdots \lambda_n$ and linearly independent corresponding eigenvectors. By theory of Markov chains, it is also known that $1 = |\lambda_1| \geq |\lambda_2| \geq \cdots |\lambda_n|$. Since $\mathbf{P}$ scales the length of $\mathbf{e}^{(t)}$ by a factor of $\leq |\lambda_2|$:

$$\left\|\mathbf{e}^{(t+1)}\right\| \;=\; \left\|\mathbf{P}\mathbf{e}^{(t)}\right\| \;\leq\; |\lambda_2| \left\|\mathbf{e}^{(t)}\right\| \;\Rightarrow\; \left\|\mathbf{e}^{(t+1)}\right\| \;\leq\; |\lambda_2|^t \left\|\mathbf{e}^{(0)}\right\|. \tag{7}$$

For a given speed vector $\mathbf{s}$, one can define the "potential" as the normed distance: $\left\|\mathbf{f}^{(t)} - \mathbf{f}^{*,\mathbf{s}}\right\|_1$. This measures the imbalance in the network in terms of the finishing times. From (7), since the error vector $\mathbf{e}$ converges to zero linearly, the potential at balanced state is zero. Note that this load balancing setting is identical to the Lyapunov dynamical system introduced in Section 4. Specifically, the speed vector $\mathbf{s}$ is the *system parameter*, the evolution function $F(\boldsymbol{\ell}^{(t)})$ is the diffusion process as described in Algorithm 1 and the potential as mentioned above corresponds to the Lyapunov function $G_{\mathbf{s}}(\boldsymbol{\ell}^{(t)}) = G_{\mathbf{s}}^t$. Note that by (7) it follows that $G_{\mathbf{s}}^{t+1} \;\leq\; |\lambda_2|^t\, G_{\mathbf{s}}^{\circ}$. In the following, all norms are L1 norms.

Lemma 7.1. *For a speed vector $\mathbf{s}$ and an arbitrary load profile vector $\boldsymbol{\ell}$, let $\mathbf{f}$ denote the corresponding finishing time vector. For a Lyapunov function defined as $G_{\mathbf{s}} = \|\mathbf{f} - \mathbf{f}^{*,\mathbf{s}}\|$, if the speed vector changes to $\mathbf{s}'$ for the same load profile, then:*

$$G_{\mathbf{s}'} \;\leq\; G_{\mathbf{s}} \;+\; Mn\left|\frac{1}{\|\mathbf{s}'\|} - \frac{1}{\|\mathbf{s}\|}\right|.$$

Proof. For speed vector changes $\mathbf{s}'$ and the same load profile, the Lyapunov function is given by:

$$G_{\mathbf{s}'} \;=\; \left\|\mathbf{f} - \mathbf{f}^{*,\mathbf{s}'}\right\| \;\leq\; \|\mathbf{f} - \mathbf{f}^{*,\mathbf{s}}\| + \left\|\mathbf{f}^{*,\mathbf{s}'} - \mathbf{f}^{*,\mathbf{s}}\right\| \;=\; G_{\mathbf{s}} + \left\|\mathbf{f}^{*,\mathbf{s}'} - \mathbf{f}^{*,\mathbf{s}}\right\|.$$

Let ℓ_i denote the load on machine m_i. The optimal load on the machines in a balanced state can be characterized using the following optimization problem:

$$\min \sum_i \frac{\ell_i}{s_i} \qquad \text{s.t.} \sum_i \ell_i = M.$$

Using the underlying symmetry, we can claim that the load on any machine m_i in the balanced state and its corresponding finishing time are $\ell_i^* = \frac{s_i \cdot M}{\sum_k s_k}$ and $f_i^{*,\mathbf{s}} = \frac{\ell_i^*}{s_i} = \frac{M}{\sum_k s_k}$ respectively. It then follows that:

$$\left\|\mathbf{f}^{*,\mathbf{s}'} - \mathbf{f}^{*,\mathbf{s}}\right\| \;=\; \sum_i \left|\frac{\ell_i'^*}{s_i'} - \frac{\ell_i^*}{s_i}\right| \;=\; \sum_i \left|\frac{M}{\sum_k s_k'} - \frac{M}{\sum_k s_k}\right|. \qquad \square$$

To formalize the problem, let $\mathcal{LB}(N, M)$ be a family of load balancing environments where N denotes the network of underlying machines and M the total weight of jobs. Each individual environment $LB_{\mathbf{s}} \in \mathcal{LB}(N, M)$ is parameterized by the machine-speed vector $\mathbf{s}$. The corresponding potential function is denoted by $G_{\mathbf{s}}$. The theorem below follows directly from the above lemma and Theorem 4.1.

Theorem 7.2. *Let $\mathcal{LB}(N, M)$ be a family of load balancing environments on n machines with the corresponding diffusivity matrix being P_N. Let $\mathbf{s}^0, \mathbf{s}^1, \cdots, \mathbf{s}^T$ denote the vector of machine speeds at times $0, 1 \cdots T$ respectively. If we denote by λ_2 the second largest eigenvalue of P_N and $\Phi(\mathbf{s}^t, \boldsymbol{\ell}^t) := G_{\mathbf{s}^t}(\boldsymbol{\ell}^t)$, then*

$$\Phi(\mathbf{s}^T, \boldsymbol{\ell}^T) \;\leq\; |\lambda_2|^T \cdot \Phi(\mathbf{s}^0, \boldsymbol{\ell}^0) \;+\; Mn \sum_{t=1}^{T} |\lambda_2|^{T-t} \cdot \left|\frac{1}{\|\mathbf{s}^t\|} - \frac{1}{\|\mathbf{s}^{t-1}\|}\right| \;.$$

Since Φ is a measure of load imbalance in the network in terms of finishing times, the above theorem implies that if the change in the speed vectors across rounds is small, then the imbalance at time T is small and depends largely on the most recent changes.

8 DISCUSSION

A canonical approach to analysing multi-agent (dynamical) system is by designing a Lyapunov or potential function. In this paper, we provide two general yet handy frameworks to generalize static analyses of dynamical systems with linearly converging Lyapunov functions to dynamic environment settings. Given the vast literature on dynamical systems (and iterative algorithms) under this category and the immense desire to demonstrate their robustness against environment variations, our frameworks can serve as a textbook technique for this important aspect of multi-agent dynamical systems.

A major open problem is whether a similar framework is admissible when the Lyapunov function convergence rate is slower than linear, say $O(1/T)$. Our intuition is if this happens, a variation near the fixed point (equilibrium) might take much longer time to be recovered than the same level of variation when afar from the fixed point. It is not clear to us how to design an analysis for such systems which captures such a distinction and provides clean performance guarantees (as we have done in the paper).

REFERENCES

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

[2] K. Arrow and G. Debreu. 1954. Existence of an equilibrium for a competitive economy. *Econometrica* 22, 3 (1954), 265–290.

[3] Xiaohui Bei, Jugal Garg, and Martin Hoefer. 2016. Ascending-Price Algorithms for Unknown Markets. In *Proc. 17th ACM Conf. Economics and Computation (EC)*. 699.

[4] Xiaohui Bei, Jugal Garg, Martin Hoefer, and Kurt Mehlhorn. 2016. Computing Equilibria in Markets with Budget-Additive Utilities. In *Proc. 24th European Symposium on Algorithms (ESA)*. 8:1–8:14.

[5] Benjamin Birnbaum, Nikhil Devanur, and Lin Xiao. 2011. Distributed algorithms via gradient descent for fisher markets. In *Proc. 12th ACM Conf. Electronic Commerce (EC)*. 127–136.

[6] Yun Kuen Cheung and Richard Cole. 2016. A Unified Approach to Analyzing Asynchronous Coordinate Descent and Tatonnement. *CoRR* abs/1612.09171 (2016).

[7] Yun Kuen Cheung and Richard Cole. 2018. Amortized Analysis of Asynchronous Price Dynamics. In *Proc. 26th European Symposium on Algorithms (ESA)*. 18:1–18:15.

[8] Yun Kuen Cheung, Richard Cole, and Nikhil Devanur. 2013. Tatonnement beyond gross substitutes? Gradient descent to the rescue. In *Proc. 45th Symp. Theory of Computing STOC*. 191–200.

[9] Yun Kuen Cheung, Richard Cole, and Ashish Rastogi. 2012. Tatonnement in ongoing markets of complementary goods. In *Proc. 13th ACM Conf. Electronic Commerce (EC)*. 337–354.

[10] Yun Kuen Cheung, Richard Cole, and Yixin Tao. 2018. Dynamics of Distributed Updating in Fisher Markets. In *Proc. 19th ACM Conf. Economics and Computation (EC)*.

[11] Yun Kuen Cheung, Martin Hoefer, and Paresh Nakhe. 2018. Tracing Equilibrium in Dynamic Markets via Distributed Adaptation. *CoRR* abs/1804.08017 (2018). http://arxiv.org/abs/1804.08017

[12] Bruno Codenotti, Benton McCune, and Kasturi R. Varadarajan. 2005. Market equilibrium via the excess demand function. In *Proc. 37th Symp. Theory of Computing (STOC)*. 74–83.

[13] Richard Cole and Lisa Fleischer. 2008. Fast-converging tatonnement algorithms for one-time and ongoing market problems. In *Proc. 40th Symp. Theory of Computing (STOC)*. 315–324.

[14] Richard Cole, Lisa Fleischer, and Ashish Rastogi. 2010. Discrete Price Updates Yield Fast Convergence in Ongoing Markets with Finite Warehouses. *CoRR* abs/1012.2124 (2010). http://arxiv.org/abs/1012.2124

[15] Nikhil R. Devanur, Christos H. Papadimitriou, Amin Saberi, and Vijay V. Vazirani. 2008. Market equilibrium via a primal–dual algorithm for a convex program. *J. ACM* 55, 5 (2008).

[16] Ran Duan, Jugal Garg, and Kurt Mehlhorn. 2016. An Improved Combinatorial Polynomial Algorithm for the Linear Arrow-Debreu Market. In *Proc. 27th Symp. Discrete Algorithms (SODA)*. 90–106.

[17] Ran Duan and Kurt Mehlhorn. 2015. A combinatorial polynomial algorithm for the linear Arrow-Debreu market. *Inf. Comput.* 243 (2015), 112–132.

[18] Dylan Foster, Zhiyuan Li, Thodoris Lykouris, Karthik Sridharan, and Éva Tardos. 2016. Learning in Games: Robustness of Fast Convergence. In *Proc. Conf. Neural Information Processing Systems (NIPS)*. 4727–4735.

[19] Kamal Jain. 2007. A Polynomial Time Algorithm for Computing an Arrow-Debreu Market Equilibrium for Linear Utilities. *SIAM J. Comput.* 37, 1 (2007), 303–318.

[20] F. Kelly. 1997. Charging and rate control for elastic traffic. *European Transactions on Telecommunications* 8, 1 (1997), 33–37.

[21] F. Kelly, A. Maulloo, and D. Tan. 1998. Rate control in communication networks. *Journal of Operational Research Society* 49 (1998), 237–252.

[22] F. Kelly and V. Vazirani. 2002. Rate Control as a Market Equilibrium. (2002). Unpublished manuscript.

[23] Dave Levin, Katrina LaCurts, Neil Spring, and Bobby Bhattacharjee. 2008. Bittorrent is an Auction: Analyzing and Improving Bittorrent's Incentives. In *Proc. ACM Conf. Data Communication (SIGCOMM)*. 243–254.

[24] Thodoris Lykouris, Vasilis Syrgkanis, and Éva Tardos. 2016. Learning and Efficiency in Games with Dynamic Population. In *Proc. 27th Symp. Discrete Algorithms (SODA)*. 120–129.

[25] James Orlin. 2010. Improved algorithms for computing fisher's market clearing prices: computing fisher's market clearing prices. In *Proc. 42nd Symp. Theory of Computing (STOC)*. 291–300.

[26] Raghu Subramanian and Isaac Scherson. 1994. An analysis of diffusive load-balancing. In *Proc. 6th Symp. Parallelism in Algorithms and Architectures (SPAA)*. 220–225.

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

[28] Fang Wu and Li Zhang. 2007. Proportional response dynamics leads to market equilibrium. In *Proc. 39th Symp. Theory of Computing (STOC)*. 354–363.

[29] Yinyu Ye. 2008. A path to the Arrow-Debreu competitive market equilibrium. *Math. Program.* 111, 1-2 (2008), 315–348.

[30] Li Zhang. 2011. Proportional response dynamics in the Fisher market. *Theor. Comput. Sci.* 412, 24 (2011), 2691–2698.