



King's Research Portal

DOI:

[10.24963/ijcai.2021/16](https://doi.org/10.24963/ijcai.2021/16)

Document Version

Publisher's PDF, also known as Version of record

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Cheung, Y. K., Leonardos, S., & Piliouras, G. (2021). Learning in Markets: Greed Leads to Chaos but Following the Price is Right. In Z.-H. Zhou (Ed.), *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21* (pp. 111-117). International Joint Conferences on Artificial Intelligence Organization. <https://doi.org/10.24963/ijcai.2021/16>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Learning in Markets: Greed Leads to Chaos but Following the Price is Right

Yun Kuen Cheung¹, Stefanos Leonardos², Georgios Piliouras²

¹Royal Holloway University of London

²Singapore University of Technology and Design

yunkuen.cheung@rhul.ac.uk, {stefanos_leonardos; georgios}@sutd.edu.sg

Abstract

We study learning dynamics in distributed production economies such as blockchain mining, peer-to-peer file sharing and crowdsourcing. These economies can be modelled as multi-product Cournot competitions or all-pay auctions (Tullock contests) when individual firms have market power, or as Fisher markets with quasi-linear utilities when every firm has negligible influence on market outcomes. In the former case, we provide a formal proof that Gradient Ascent (GA) can be Li-Yorke chaotic for a step size as small as $\Theta(1/n)$, where n is the number of firms. In stark contrast, for the Fisher market case, we derive a Proportional Response (PR) protocol that converges to market equilibrium. The convergence result of the PR dynamics is obtained in full generality, in the sense that it holds for Fisher markets with *any* quasi-linear utility functions. Conversely, the chaos results for the GA dynamics are established in the simplest possible setting of two firms and one good, and hold for a wide range of price functions with different demand elasticities. Our findings suggest that as multi-agent interactions grow larger, the ensuing market (instead of game-theoretic) conditions allow us to formally derive natural and stable learning protocols which converge to effective outcomes rather than being chaotic.

1 Introduction

Multi-agent learning in production economies is an important yet underexplored domain. Production economies are classically modelled as Cournot competitions [Varian, 2010] or imperfectly discriminating all-pay auctions (Tullock contests) [DiPalantino and Vojnovic, 2009]. In these models, participating firms have *market power*, and they can significantly influence aggregate outcomes (prices or total exerted effort) with their decisions. However, the advancement of the internet has prompted a rapid paradigm shift in economic competition. Blockchain mining [Arnosti and Weinberg, 2018; Fiat *et al.*, 2019], peer-to-peer file sharing [Levin *et al.*, 2008] and crowdsourcing [Horton and Chilton, 2010], among oth-

ers, all constitute *distributed* production economies with large numbers of small competitors (individuals or firms). In contrast to the classic Cournot or Tullock models, firms in these economies typically engage in multiple concurrent competitions. Moreover, due to their relative small sizes, each firm has negligible influence on prices and hence becomes *price-taker*. As a result, this form of competition more closely resembles the economic model of Fisher markets in which firms take prices as independently given signals, and purchase optimal bundles of goods (or invest on optimal portfolios to produce goods) given their budget (or capital) constraints.

The question of which adaptive or learning protocols behave well in these economies is largely open and actively researched. In both Cournot competition¹ and Fisher markets, firms repeatedly observe the aggregate production, and adjust their production outputs over time to improve their own profits. However, empirical results regarding Cournot competition suggest that standard adaptive algorithms, e.g., best response, can lead to unstable and irregular adjustments, even in very simple instances (e.g., when there are only two firms and one good) [Theocharis, 1960; Puu, 1991; Wärneryd, 2018]. In contrast, when firms ignore their market power and act as price-takers, the outcomes can be more stable. A line of recent works [Wu and Zhang, 2007; Zhang, 2011; Birnbaum *et al.*, 2011; Cheung *et al.*, 2012; Cheung *et al.*, 2020; Cheung *et al.*, 2018; Brânzei *et al.*, 2019; Cheung *et al.*, 2019; Gao and Kroer, 2020] showed that natural adaptive algorithms, including tâtonnement and proportional response (PR), lead to stable adjustments in different markets, where they converge to *market equilibria*.²

1.1 Model and Contribution

Motivated by the above, our aim is to study the behavior of learning dynamics in production economies from a theoretical perspective. Our research goals are 1) to establish formal

¹The mathematical equivalence between Cournot competition with isoelastic demand and imperfectly discriminating all-pay auctions with proportional success functions or simply Tullock contests is documented in [Szidarovszky and Okuguchi, 1997; Wärneryd, 2018] (among others). We elaborate on this relation in Section 2.

²Another relevant and interesting result suggests that when firms trade resources using PR, the underlying production economy grows in the long term under mild conditions [Brânzei *et al.*, 2018].

mathematical arguments that explain the irregular behavior of *greedy* learning rules, such as Gradient Ascent and Best Response dynamics, and 2) to seek protocols that behave well under general conditions.

Concerning the first goal, we present the first rigorous mathematical proof that the constant step-size *Gradient Ascent* (GA) algorithm can exhibit *Li-Yorke chaos* [Li and Yorke, 1975] in Cournot competition (equivalently, in all-pay auctions or Tullock contests) even when the firms are homogeneous. This provides a formal explanation for the *unpredictable* evolution of these systems that is frequently observed in practice. To derive this result, we leverage Sharkovsky’s theorem which provides a tractable way to verify the conditions in Li-Yorke’s characterization of chaos [Palaioanos *et al.*, 2017]. In the case of GA, our findings are robust in two aspects: first, chaos emerges for a large family of price functions induced by different demand elasticities, and second, chaos emerges even when the step-size is as small as $\Theta(1/n)$. Our results in this direction contribute to the growing literature that studies various forms of chaos in game dynamics [Sato *et al.*, 2002; Galla and Farmer, 2013; Cheung and Piliouras, 2019; Cheung and Piliouras, 2020; Cheung and Tao, 2021; Chotibut *et al.*, 2021; Leonardos and Piliouras, 2021].

Informally, a dynamical system is Li-Yorke chaotic if there are uncountably many pairs of trajectories which get arbitrarily close together (but never intersect) and move apart indefinitely. When two trajectories are very close to each other, they become essentially indistinguishable due to the precision limitation inherent with the environment or computer. In other words, we cannot tell which of the two trajectories will be realized in the future — this is exactly what *unpredictable* means. A primary reason for the chaos to arise is that each firm uses its own market power to strategically influence the price. When all firms make such strategic manipulations simultaneously, they aggregately drive prices up and down without proper control.

While the previous technique does not lead to a formal proof of Li-Yorke chaos in the case of Best Response (BR) dynamics, we formalize the (in)-stability properties of the latter via eigenvalue analysis of the non-linear dynamical system. Here, instability refers to abrupt changes in the long term behavior of the dynamics in response to small perturbations of the systems’ parameters (e.g., firms costs).³

Since robustness is an essential property in distributed production economies both from a normative and a descriptive perspective, the above results provide a convincing argument against the use of *game-theoretically motivated* protocols. This brings us to our second goal which is to seek learning protocols that result in stable outcomes.

Our main result in this direction is to propose a *market-motivated* Proportional Response (PR) algorithm and show that it is stable and robust: from any initial condition, the

PR update rule converges to the market equilibrium of an ensuing Fisher market that captures production economies, namely Fisher market with quasi-linear utility functions. The protocol is simple and can be run by each firm independently using only local and observable (market level) information, which makes it particularly suitable for these distributed settings. It can be interpreted as a naturally motivated adaptive algorithm from a firm’s perspective: in each round, each firm appropriates a certain amount of money, and invests it to the productions of different goods in proportion to the revenues received from selling them in the previous round.

One necessary assumption to establish this result is that as economies grow larger, firms have a negligible influence on aggregate outputs. However, we formally argue that in the distributed production economy setting, market equilibria are approximate Nash equilibria. This finding is in line with the largeness concept in [Cole and Tao, 2016], who showed that when markets grow *large*, they become asymptotically efficient even under agents’ strategic behaviors. This implies that the assumption of diminished influence on outcomes does not significantly affect the equilibrium outcome of the system. However, it does have important implications from a technical perspective. In particular, by modeling production economies as Fisher markets, we can leverage their *Eisenberg-Gale convex-program formulation* [Eisenberg and Gale, 1959] to draw a direct analogue between our PR algorithm and standard optimization methods like mirror descent. This allows us to apply tools from optimization theory and provides a principled approach to derive convergence proofs.

Paper outline. In Section 2, we present our three models: Cournot competition with multiple-goods, Tullock contests and Fisher Markets, and discuss their mathematical connections. Section 3 contains our main results: convergence of PR dynamics and chaos and instabilities of GA and BR dynamics. We discuss the techniques we use in Sections 4 and 5; detailed proofs can be found in the full version.

2 Models and Definitions

In this section, we describe the Cournot competition and Fisher market models. In their classical descriptions, *quantities of goods produced* are used as the driving variables to define the notions of Nash and market equilibria. However, it will be more convenient to use *spendings/investments on the production of a good* as the driving variables here, since this is the domain of the PR algorithm. In all models, $N = \{1, 2, \dots, n\}$ is the set of firms (agents) and $M = \{1, 2, \dots, m\}$ is the set of goods.

Multi-good Cournot competition (CC) with isoelastic demands. Each firm i invests an amount $b_{ij} \geq 0$ on producing good j . We write $\mathbf{b}_i := (b_{ij})_{j \in M}$ and $\mathbf{b} := (\mathbf{b}_i)_{i \in N}$. Each firm i has only finite amount of capital, K_i , to invest, thus it is subject to a capital constraint $\sum_j b_{ij} \leq K_i$. We assume that the marginal cost of producing good j is the same for all firms, which we denote by α_j . Thus, the quantity of good j produced by firm i is b_{ij}/α_j . Each good j has isoelastic demand, i.e., the total sales revenue of the good is con-

³This formalization closely mirrors existing empirical results on BR dynamics [Puu, 1991; Wärneryd, 2018]. Hence, we only present some indicative visualizations (Figure 3), and defer the formal statement to the full version (<https://arxiv.org/abs/2103.08529>).

stant, denoted by v_j . Thus, the price function⁴ for good j is $P_j(\mathbf{b}) := v_j / (\sum_i b_{ij} / \alpha_j)$, and the revenue of firm i received from the sales of good j is $P_j(\mathbf{b}) \cdot (b_{ij} / \alpha_j) := v_j \cdot y_{ij}$, where y_{ij} denotes the *market share* of firm i on good j :

$$y_{ij} := b_{ij} / \sum_k b_{kj}. \quad (1)$$

The profit of firm i is its revenue from the sales of all goods minus its total investment: $\sum_j v_j y_{ij} - \sum_j b_{ij}$.

Tullock contest (TC). The above setting admits a correspondence to *multiple Tullock contests*. According to this interpretation, each firm i invests an amount of $b_{ij} \geq 0$ on producing good j , but now the goods are considered as *prizes*, and the probability that firm i wins good j is y_{ij} as defined in Eqn. (1). This probabilistic interpretation is natural in the applications of blockchain mining and imperfectly discriminating all-pay auctions (crowdsourcing). Now, different firms can have different valuations on the prize, so the parameter v_j in CC may be distinct for different firms; we let v_{ij} denote the valuation of firm i on good j . The expected profit of firm i is

$$u_i(\mathbf{b}_i) := \sum_j v_{ij} y_{ij} - \sum_j b_{ij}. \quad (2)$$

While CC and TC have differences in their rationales, they admit a correspondence in mathematical terms, by replacing deterministic profit in CC with expected profit in TC, and v_j with v_{ij} for different firms i . Accordingly, we will henceforth refer to this model as CC/TC or simply TC.

Definition 1 (Nash equilibrium). For any $\delta \geq 0$, we say that \mathbf{b}^* is a δ -Nash Equilibrium (δ -NE) of a CC/TC if for each agent $i \in N$, $\max_{\mathbf{b}_i: \sum_j b_{ij} \leq K_i} u_i(\mathbf{b}_i, \mathbf{b}_{-i}^*) \leq (1 + \delta) \cdot u_i(\mathbf{b}_i^*, \mathbf{b}_{-i}^*)$. In other words, agent i cannot improve her utility by more than a δ fraction at \mathbf{b}^* by unilaterally changing her own investment portfolio. We call a 0-NE simply a NE.

Fisher market (FM). In a Fisher market, each good j has a supply which is normalized to one unit. Again, b_{ij} denotes the spending of firm i on good j , and each firm i has a budget of K_i , so the constraint $\sum_j b_{ij} \leq K_i$ applies. Let $\mathbf{p} = (p_j)_{j \in M}$, where p_j denotes the price of good j . At \mathbf{b}_i , firm i gets b_{ij}/p_j units of good j and has a *quasi-linear* utility function, $u_i(\mathbf{b}_i | \mathbf{p})$, which takes the form

$$u_i(\mathbf{b}_i | \mathbf{p}) = \sum_j v_{ij} \cdot (b_{ij}/p_j) - \sum_j b_{ij}, \quad (3)$$

where v_{ij} denotes firm i 's valuation of one unit of good j . At price vector \mathbf{p} , each firm i select an *optimal budget allocation* $\mathbf{b}_i^\#$ in $\arg \max_{\mathbf{b}_i} u_i(\mathbf{b}_i | \mathbf{p})$ which maximizes its utility subject to the constraint $\sum_j b_{ij} \leq K_i$. At an optimal budget vector $\mathbf{b}_i^\#$, a vector $\mathbf{x}_i^\# := (b_{ij}^\# / p_j)_{j \in M}$ is called a *production bundle* of agent i at price vector \mathbf{p} .

Definition 2 (Market equilibrium). A price vector $\mathbf{p}^\# = (p_j^\#)_{j \in M}$ is a *market equilibrium (ME)* if there exists an optimal budget allocation $\mathbf{b}^\# = (\mathbf{b}_i^\#)_{i \in N}$ at $\mathbf{p}^\#$, such that for

each good j , $\sum_i b_{ij}^\# = p_j^\#$. The vector $\mathbf{b}^\#$ is called a *market equilibrium spending*.⁵

2.1 Connection between TC and FM

The crucial difference between TC and FM is that in TC, prices are determined endogenously as a function of \mathbf{b} , whereas in FM, prices are viewed as independent inputs that do not explicitly depend on \mathbf{b} . Thus, while both models require each firm i to make an allocation \mathbf{b}_i that is subject to the same budget constraint $\sum_j b_{ij} \leq K_i$, the methods to determine outcomes differ.

However, if (\mathbf{p}, \mathbf{b}) are market equilibrium and market equilibrium spending respectively of an FM, then $\sum_i b_{ij}/p_j = 1$ for each good j . Thus, we can translate b_{ij}/p_j , which is the quantity of good j that firm i gets at the market equilibrium, to the probability that firm i wins good j in the corresponding TC. Under this translation, the outcome in the FM is the same as the outcome in the TC. Due to the well-known properties of Fisher markets, this outcome is Pareto-optimal, and it is envy-free if K_i is identical for all i .

The above suggest that if there is an algorithm that converges to the market equilibrium spending (our Theorem 4 establishes this) of the FM, then it yields a feasible solution of the corresponding TC. The remaining question is the quality of this feasible solution, i.e., how close it is to a Nash equilibrium of the TC. It turns out that if the underlying distributed production economy satisfies a natural *largeness* property, then the market equilibrium spending is also a δ' -NE for some small $\delta' > 0$. In particular, as we show in Proposition 3 below, this is the case if the budget of each firm is small compared to *any* market equilibrium price, i.e., if $\max_{i,j} \{K_i/p_j^\#\} \leq \delta$ for a small $\delta > 0$. We may view δ as a parameter that describes the largeness of the economy: the smaller δ is, the *larger* the economy is. (We also need the *bang-per-buck ratio* $\beta_i := \max_j \{v_{ij}/p_j^\#\}$ to be sufficiently high for all firms i , because otherwise a firm might invest nothing thus attain zero utility, forcing δ' to be $+\infty$.)

Proposition 3. Suppose that $\mathbf{b}^\#$ is a market equilibrium spending vector of a quasi-linear FM, and $\mathbf{p}^\#$ is the corresponding market equilibrium price vector. For every $i \in N$, let $\beta_i := \max_j \{v_{ij}/p_j^\#\}$. If $\max_{i,j} \{K_i/p_j^\#\} \leq \delta$, then $\mathbf{b}^\#$ is also a δ' -NE of the corresponding TC, where

$$\delta' = \max_{i: \beta_i > 1} \left\{ \left(\frac{\beta_i}{1 - \delta} - 1 \right) / (\beta_i - 1) \right\} - 1,$$

provided that there is no firm i with $1 - \delta < \beta_i < 1$.

It is easy to see that if $\min_i \{\beta_i\}$ grows, then δ' tends toward $\delta/(1 - \delta)$.

3 Our Main Results

We present our two main results here. We discuss the methodology of proving them in Sections 4 and 5.

⁵The last condition is same as $\sum_i b_{ij}^\# / p_j^\# = 1$, which is the classical definition of market equilibrium.

⁴We also consider more general price functions induced by different demand elasticities in Section 5.

Algorithm 1 PR-QLIN Learning Protocol

Input: $(K_i, v_{i1}, v_{i2}, \dots, v_{im}, \mathbf{b}_i^\circ)$ for each firm i
Output: market equilibrium spending $\mathbf{b}^\#$.

```

1: for  $t = 0, 1, 2, \dots$  do
2:   for every firm  $i$  and good  $j$  do
3:      $y_{ij}^t \leftarrow b_{ij}^t / \sum_k b_{kj}^t$ 
4:   for every firm  $i$  do
5:      $S_i^t \leftarrow \sum_j v_{ij} y_{ij}^t$ 
6:     if  $S_i^t > K_i$  then
7:       for every good  $j$  do
8:          $b_{ij}^{t+1} \leftarrow (v_{ij} y_{ij}^t / S_i^t) \cdot K_i$ 
9:     else for every good  $j$  do
10:       $b_{ij}^{t+1} \leftarrow v_{ij} y_{ij}^t$  // same as  $(v_{ij} y_{ij}^t / S_i^t) \cdot S_i^t$ 
    
```

3.1 Proportional Response in Quasi-Linear FM

In a quasi-linear Fisher market, our PR protocol starts with each firm $i \in N$ investing an arbitrary portfolio \mathbf{b}_i° which is positive, i.e., $b_{ij}^\circ > 0$ for all $j \in M$. In each round, firms update their portfolios simultaneously according to the PR-QLIN protocol in Algorithm 1.

The PR-QLIN protocol can be naturally interpreted. After all firms update their investment portfolios in round t , one unit of each good j is allocated to the firms in proportion to their investments on the good. Thus, firm i gets y_{ij}^t units of good j (line 3). Then each firm i computes its attained utility, S_i^t , without subtracting investment cost (line 5). If $S_i^t > K_i$, then firm i will appropriate all of its capital, K_i , for investment in round $t + 1$; otherwise it will only appropriate an amount of S_i^t for investment. Then each firm invests its appropriated capital on each good in proportion to the utility attained from that good in the previous round, i.e., firm i invests a fraction of $v_{ij} y_{ij}^t / S_i^t$ of its appropriated capital on good j . Our main result is stated below.

Theorem 4. *Given any positive starting point \mathbf{b}° , the algorithm PR-QLIN converges to the set of market equilibrium spending vectors of the quasi-linear Fisher market.*

3.2 Gradient Ascent (GA) and Li-Yorke Chaos

To establish our chaos results of the GA dynamics in CC (hence, also in TC), we consider a CC with one good and n firms. Since there is only one good, we omit the subscript $j = 1$ and use the shorthand $\alpha \equiv \alpha_1$ to denote the marginal cost of producing the good (recall from Section 2 that this is equal for all firms). In this setting, it is more convenient to use the quantities of the good produced, i.e., the variables $x_i = b_{i1}/\alpha$, as the driving variables. Without loss of generality, let $v_1 = 1$. Then the utility of firm i is $u_i(\mathbf{x}) = x_i / (\sum_k x_k) - \alpha x_i$. The *Gradient Ascent (GA)* update rule is given by $x_i^{t+1} \leftarrow x_i^t + \eta \cdot \nabla_i u_i(\mathbf{x}^t)$, where η is the step-size.

Assuming that the initial point is *symmetric*, i.e., that x_i° is identical for all i , then in each round $t > 0$, the x_i^t 's remain identical for all i . Thus, a symmetric GA dynamic is essen-

tially one-dimensional, and its trajectory can be represented by the sequence $\{x_1^t\}_{t \geq 0}$ generated by the GA update rule:

$$x_1^{t+1} \leftarrow x_1^t + \eta \cdot \left(\frac{n-1}{n^2 x_1^t} - \alpha \right). \quad (4)$$

Our main result states that even for such an apparently simple one-dimensional dynamical system, *chaos* occurs with step-size η as small as $\Theta(1/n)$. Here, we refer to *Li-Yorke chaos* which is formally defined below.

Definition 5 (Li-Yorke Chaos). A discrete time dynamical system $(x^t)_{t \in \mathbb{N}}$ such that $x^t := f^t(x^\circ)$ for a continuous update rule $f : X \rightarrow X$ on a compact set $X \subseteq \mathbb{R}$ is called *Li-Yorke chaotic*, if (i) for each $k \in \mathbb{N}$, there exists a periodic point $\hat{x} \in X$ with period k , and (ii) there is an uncountably infinite set $S \subset X$ that is *scrambled*, i.e., if for each $x \neq x' \in S$ it holds that $\liminf_{t \rightarrow \infty} |f^t(x) - f^t(x')| = 0 < \limsup_{t \rightarrow \infty} |f^t(x) - f^t(x')|$.

Theorem 6 (Li-Yorke Chaos in n -Player CC/TC). *Consider a symmetric GA dynamic with n firms and marginal cost $\alpha > 0$. Then for any step-size $\eta \geq 3(n-1)/n^2 \alpha^2$, the essentially-one-dimensional dynamical system (4) is Li-Yorke chaotic.*

This theorem applies with isoelastic price function. In Section 5, we consider a larger family of price functions and show that Li-Yorke chaos also occurs in the corresponding symmetric GA dynamics. We also present theoretical and empirical evidences that instability arises when the GA rule is replaced by the Best Response rule.


Remark. In practice, firms may choose to use a large step-size in a myopic, greedy approach to profit maximization. Given that chaos occurs with a vanishingly small step-size $\Theta(1/n)$ as the number of firms increases (cf. Theorem 6), our result is practically relevant for distributed production economies in which many small firms are involved. Stability results should be possible for smaller step sizes, however, such step sizes are not particularly interesting from a practical perspective. Finally, the presence of a centralised planner who may enforce small step sizes is a rather unnatural assumption for the settings and applications that we consider.

4 Proportional Response (PR) Dynamics

Our proof of Theorem 4 consists of two major steps. In the first step, we derive a convex program that captures the market equilibrium (ME) spending of the quasi-linear Fisher market via the approach of [Birnbaum *et al.*, 2011; Cole *et al.*, 2017; Cheung *et al.*, 2018]. In the second step, we show that a general Mirror Descent (MD) algorithm converges to the optimal solution of this convex program; PR-QLIN is an instantiation of this MD algorithm.

Convex program framework. We first utilize a convex optimization framework to derive a convex program that captures the ME spendings of any quasi-linear FM. The ensuing framework is summarized in Figure 1. In short, via duality and variable transformations, the market equilibria of a FM can be captured by various convex programs, each with a different domain.⁶ Our starting point is a convex program pro-

⁶For linear Fisher markets, i.e. markets in which each agent has a



Program	Description	Variables
(EG)	Eisenberg-Gale	x_{ij} = allocations $i \in N, j \in M$
(D)	Dual	p_j = prices $j \in M$
(TD)	Transformed dual	$q_j = \ln(p_j)$ $j \in M$
(SH)	Shmyrev-type	b_{ij} = spending $i \in N, j \in M$

Figure 1: Derivation of the PR-QLIN protocol via the Mirror Descent (MD) protocol. Starting from the convex program (D), which is the dual of a generalized Eisenberg-Gale (EG) program, we move to the transformed dual (TD) and by convex duality to a Shmyrev-type primal program (SH) which is, hence, equivalent to the initial program (EG). The objective function of (SH) for quasi-linear utilities is 1-Bregman convex which implies convergence of the MD protocol.

posed by [Cole *et al.*, 2017] that captures ME prices of quasi-linear Fisher market (which belongs to type (D) in Figure 1). From this, we derive a new convex program with captures the ME spendings of the market (which belongs to type (SH)); see the full version for the details. The convex program is

$$\begin{aligned}
 & \min_{\mathbf{b}, \mathbf{w}, \mathbf{p}} F(\mathbf{b}, \mathbf{w}, \mathbf{p}) \\
 & \text{s.t. } \sum_{i=1}^n b_{ij} = p_j, \quad \forall j \in M, \\
 & \quad \sum_{j=1}^m b_{ij} + w_i = K_i, \quad \forall i \in N, \\
 & \quad b_{ij}, w_i \geq 0, \quad \forall i \in N, j \in M,
 \end{aligned} \tag{SH}$$

where $F(\mathbf{b}, \mathbf{w}, \mathbf{p}) := -\sum_{i=1}^n \sum_{j=1}^m b_{ij} \ln v_{ij} + \sum_{i=1}^n w_i + \sum_{j=1}^m p_j \ln p_j$. For brevity, we will write $F(\mathbf{z})$. Observe that the first and second constraints determine the values of \mathbf{w}, \mathbf{p} in terms of b_{ij} 's. Thus, we can rewrite the convex program to have variables \mathbf{b} only, and the remaining constraints are $b_{ij} \geq 0$ and $\sum_{j=1}^m b_{ij} \leq K_i$.

4.1 From Mirror Descent to PR

After having the convex program with variables \mathbf{b} only, we can compute a ME spending by the optimization algorithm of Mirror Descent (MD). To begin, we recap a general result about MD [Chen and Teboulle, 1993; Birnbaum *et al.*, 2011].

Definition 7 (KL-divergence and L -Bregman convexity). Let C be a compact and convex set and let h be a convex function on C . Then, for any $\mathbf{z}' \in C, \mathbf{z} \in \text{rint}(C)$ where $\text{rint}(C)$ is the relative interior of C , the *Bregman divergence*, $d_h(\mathbf{z}', \mathbf{z})$, generated by h is defined by

$$d_h(\mathbf{z}', \mathbf{z}) := h(\mathbf{z}') - [h(\mathbf{z}) + \langle \nabla h(\mathbf{z}), \mathbf{z}' - \mathbf{z} \rangle].$$

The *Kullback-Leibler (KL) divergence* between \mathbf{z}' and \mathbf{z} is defined by $\text{KL}(\mathbf{z}' \| \mathbf{z}) := \sum_j z'_j \cdot \ln \frac{z'_j}{z_j} - \sum_j z'_j + \sum_j z_j$, which is the same as the Bregman divergence d_h with regularizer $h(\mathbf{z}) := \sum_j (z_j \cdot \ln z_j - z_j)$. A function F is *L -Bregman convex* w.r.t. the Bregman divergence d_h if for any $\mathbf{z}' \in C$ and $\mathbf{z} \in \text{rint}(C)$, $F(\mathbf{z}) + \langle \nabla F(\mathbf{z}), \mathbf{z}' - \mathbf{z} \rangle \leq f(\mathbf{z}') \leq f(\mathbf{z}) + \langle \nabla f(\mathbf{z}), \mathbf{z}' - \mathbf{z} \rangle + L \cdot d_h(\mathbf{z}', \mathbf{z})$.

utility similar to a quasi-linear utility, but without the subtraction of investment cost, [Eisenberg and Gale, 1959] derived a convex program which captures the ME allocation, where the driving variables are quantities of goods allocated to the agents. Subsequent works established that by considering suitable duals and transformations of Eisenberg and Gale's convex program, new convex programs can be derived which capture the ME prices and ME spendings.

Algorithm 2 MD protocol w.r.t. KL-divergence

Input: A convex set C , a function F defined on C , a parameter Γ and a point $\mathbf{z}^0 \in C$.

Output: $\mathbf{z}^* = \arg \min_{\mathbf{z} \in C} F(\mathbf{z})$.

```

1: for  $t = 0, 1, 2, \dots$  do
2:    $g(\mathbf{z}, \mathbf{z}^t) \leftarrow \langle \nabla F(\mathbf{z}^t), \mathbf{z} - \mathbf{z}^t \rangle + \Gamma \cdot \text{KL}(\mathbf{z} \| \mathbf{z}^t)$ 
3:    $\mathbf{z}^{t+1} \leftarrow \arg \min_{\mathbf{z} \in C} \{g(\mathbf{z}, \mathbf{z}^t)\}$ 

```

For the problem of minimizing a convex function $F(\mathbf{z})$ subject to $\mathbf{z} \in C$, the MD protocol w.r.t. the KL divergence is presented in Algorithm 2. In the protocol, $1/\Gamma$ is the step-size, which may vary with t (and typically diminishes with t). However, in the current application of distributed dynamics, a time-varying step-size is undesirable or even impracticable, since it requires firms to keep track of a global clock.

Theorem 8. Suppose F is an L -Bregman convex function w.r.t. the Bregman divergence d_h , and \mathbf{z}^t is the point reached after t applications of the MD update rule in Algorithm 2 with parameter $\Gamma = L$. Then $F(\mathbf{z}^t) - F(\mathbf{z}^*) \leq L \cdot d(\mathbf{z}^*, \mathbf{z}^0)/t$, where $\mathbf{z}^* = \arg \min_{\mathbf{z} \in C} F(\mathbf{z})$.

Proof sketch: We first prove Lemma 9 below. Then we show that PR-QLIN is an instantiation of Algorithm 2 with $\Gamma = 1$. This is achieved by identifying the variables \mathbf{b} in PR-QLIN as the variables \mathbf{z} in Algorithm 2, and the domain $\{\mathbf{b} \mid b_{ij} \geq 0 \text{ and } \sum_{j=1}^m b_{ij} \leq K_i\}$ as the convex set C in Algorithm 2. Thus, Theorem 8 guarantees the updates of PR-QLIN converge to an optimal solution of the convex program (SH), and hence Theorem 4 follows.

Lemma 9. The objective function $F(\mathbf{z})$ of (SH) is a 1-Bregman convex function w.r.t. the KL-divergence.

5 GA and Best Response Dynamics

To establish Theorem 6 about the GA dynamics in Eqn. (4) for $n = 2$ (the technique is similar for any $n > 2$), let $f(x) := x + \eta(\frac{1}{4x} - \alpha)$; note that $x_1^{t+1} = f(x_1^t)$ by Eqn. (4). To prove that Li-Yorke chaos occurs, we use a seminal theorem of [Li and Yorke, 1975], which states that if f has two easy-to-verify properties, then the dynamical system is Li-Yorke chaotic. The two properties are: (i) an invariant set of f that includes a fixed point x^* , i.e., an interval $I = [L, U]$ such that $f(I) \subseteq I$ with a point $L < x^* < U$ satisfying $f(x^*) = x^*$, and (ii) a point $x' \in I$ other than x^* with period 3, i.e., $f^{(3)}(x') = x'$, where $f^{(3)}(x) := (f \circ f \circ f)(x)$.




Figure 2: Li-Yorke chaos of the Gradient Ascent (GA) dynamics with constant step-size in n -firm Cournot competition with isoelastic inverse demand function (equivalently, Tullock contest with proportional success function). First panel: chaotic trajectories (light to dark lines) and their planar projections (blue dots) of the output pairs of the output pairs in $[10^{-5}, 1]$ and step-size $\eta = 5 \cdot 10^{-4}$. Second panel: chaotic aggregate output in a market with $n = 10^4$ firms with randomly selected costs in $[10^{-5}, 1]$ and step-size $\eta = 5 \cdot 10^{-4}$. Third panel: Minimum step-size for which chaos provably occurs in a 2-firm Cournot competition with inverse demand function $(x_1 + x_2)^{-\gamma}$, $\gamma > 0$. Chaotic behavior is more likely when demand is inelastic.




Figure 3: Best response dynamics: firms' outputs (x^t, y^t) (horizontal planes) with respect to time $t \in [10, 250]$ (vertical axis) for different values of the cost asymmetry parameter, r , between the two firms. In line with the theoretical predictions (see full version), the attractors of the dynamics may change significantly even for small perturbations in the system parameters.

These properties are formally established in the full version. A visualization of Theorem 6 is provided in the first two panels of Figure 2. It can be seen that chaos may emerge even for small step-size and for asymmetric marginal costs.

General price functions. The two conditions that are required in the theorem of Li and Yorke can be also verified numerically (via computer software) in the case of the parametric price function $X^{-\gamma}$, where $X := \sum_i x_i$ and $\gamma > 0$ is the inverse of the demand elasticity $\varepsilon_{(p)}$, see e.g., [López and Vives, 2019]. The lower bound of the step-size η at which chaos emerges (in the symmetric case) is decreasing as demand becomes more elastic, cf. third panel of Figure 2.

Best response dynamics. We conclude by revisiting the well-studied Best Response (BR) dynamics and formally establish that they can be unstable even in the simplest setting of two firms and one good. The general BR update rule is $x_i^{t+1} \leftarrow \arg \max_{x_i} u_i(x_i, x_{-i}^t)$. For TC with isoelastic demand, the BR dynamics take the form $x_i^{t+1} \leftarrow (x_{-i}^t / \alpha_i)^{1/2} - x_{-i}^t$, for $i = 1, 2$, where α_i is the marginal cost of firm i . BR dynamics in Cournot duopoly with isoelastic functions have been studied by [Puu, 1991] and, in the framework of contests, by [Wärneryd, 2018]. Both papers suggest that the stability of the unique fixed point, (x_1^*, x_2^*) , depends on the degree of asymmetry between the two firms, captured by the ratio $r := \alpha_1 / \alpha_2$ with instabilities emerging as the asymmetry increases. While our previous technique does not lead to a formal proof of Li-Yorke chaos in BR dynamics, we for-

malize the (in)-stability properties of the latter via eigenvalue analysis of the non-linear system, cf. full version. The result is visualized in Figure 3 which shows how the trajectories of the dynamics may change dramatically in response to even small perturbations of the model parameters (firms' costs).

6 Conclusions

The current work brings together multi-agent learning with optimization, market theory and chaos theory. Our findings suggest that by considering production economies from a market rather than a game-theoretic perspective, we can formally derive a natural learning protocol (PR) which is stable and converges to effective outcomes rather than being chaotic (GA). Due to its simple form and mild informational requirements, PR can be used to study real-world multi-agent settings from an AI perspective. Since distributed production economies capture many important applications (blockchain, peer-to-peer networks, crowdsourcing), our contributions are significant both for theoretical and practical purposes.

Acknowledgements

This research is supported in part by NRF2019-NRF-ANR095 ALIAS grant, grant PIE-SGP-AI-2018-01, NRF 2018 Fellowship NRF-NRFF2018-07, AME Programmatic Fund (Grant No. A20H6b0151) from the Agency for Science, Technology and Research (A*STAR) and the National Research Foundation, Singapore under its AI Singapore Program (AISG Award No: AISG2-RP-2020-016).

References

- [Arnosti and Weinberg, 2018] N. Arnosti and S. M. Weinberg. Bitcoin: A Natural Oligopoly. In Avrim Blum, editor, *10th ITCS*, volume 124, pages 5:1–5:1, 2018.
- [Birnbaum *et al.*, 2011] B. Birnbaum, N. R. Devanur, and L. Xiao. Distributed Algorithms via Gradient Descent for Fisher Markets. In *EC’11*, pages 127–136. ACM, 2011.
- [Brânzei *et al.*, 2018] S. Brânzei, R. Mehta, and N. Nisan. Universal Growth in Production Economies. In *NeurIPS 2018*, volume 31, pages 1973–1973, 2018.
- [Brânzei *et al.*, 2019] S. Brânzei, N. R. Devanur, and Y. Rabani. Proportional Dynamics in Exchange Economies. *CoRR*, abs/1907.05037, 2019.
- [Chen and Teboulle, 1993] G. Chen and M. Teboulle. Convergence Analysis of a Proximal-Like Minimization Algorithm Using Bregman Functions. *SIAM J. Optim.*, 3(3):538–543, 1993.
- [Cheung and Piliouras, 2019] Y. K. Cheung and G. Piliouras. Vortices Instead of Equilibria in MinMax Optimization: Chaos and Butterfly Effects of Online Learning in Zero-Sum Games. In *COLT*, pages 807–834, 2019.
- [Cheung and Piliouras, 2020] Y. K. Cheung and G. Piliouras. Chaos, Extremism and Optimism: Volume Analysis of Learning in Games. In *NeurIPS (Pre-proceedings)*, 2020.
- [Cheung and Tao, 2021] Y. K. Cheung and Y. Tao. Chaos of Learning Beyond Zero-Sum and Coordination via Game Decompositions. In *ICLR*, 2021.
- [Cheung *et al.*, 2012] Y. K. Cheung, R. Cole, and A. Rastogi. Tatonnement in ongoing markets of complementary goods. In *EC’12*, pages 337–354, 2012.
- [Cheung *et al.*, 2018] Y. K. Cheung, R. Cole, and Y. Tao. Dynamics of Distributed Updating in Fisher Markets. In *EC’18*, pages 351–368, 2018.
- [Cheung *et al.*, 2019] Y. K. Cheung, M. Hoefer, and P. Nakhe. Tracing equilibrium in dynamic markets via distributed adaptation. In *AAMAS*, pages 1225–1233, 2019.
- [Cheung *et al.*, 2020] Y. K. Cheung, R. Cole, and N. R. Devanur. Tatonnement beyond gross substitutes? Gradient descent to the rescue. *Games and Economic Behavior*, 123:295–326, 2020.
- [Chotibut *et al.*, 2021] T. Chotibut, F. Falniowski, M. Misiurewicz, and G. Piliouras. Family of chaotic maps from game theory. *Dynamical Systems*, 36(1):48–63, 2021.
- [Cole and Tao, 2016] R. Cole and Y. Tao. Large Market Games with Near Optimal Efficiency. In *EC’16*, pages 791–808, New York, NY, USA, 2016. ACM.
- [Cole *et al.*, 2017] R. Cole, N. R. Devanur, V. Gkatzelis, K. Jain, T. Mai, V. V. Vazirani, and S. Yazdanbod. Convex Program Duality, Fisher Markets, and Nash Social Welfare. In *EC’17*, pages 459–460, 2017.
- [DiPalantino and Vojnovic, 2009] D. DiPalantino and M. Vojnovic. Crowdsourcing and All-Pay Auctions. In *EC ’09*, pages 119–128, 2009.
- [Eisenberg and Gale, 1959] E. Eisenberg and D. Gale. Consensus of Subjective Probabilities: The Pari-Mutuel Method. *Ann. Math. Statist.*, 30(1):165–168, 1959.
- [Fiat *et al.*, 2019] A. Fiat, A. Karlin, E. Koutsoupias, and C. Papadimitriou. Energy Equilibria in Proof-of-Work Mining. In *EC’19*, pages 489–502, 2019.
- [Galla and Farmer, 2013] T. Galla and J. D. Farmer. Complex Dynamics in Learning Complicated Games. *PNAS*, 110(4):1232–1236, 2013.
- [Gao and Kroer, 2020] Yuan Gao and Christian Kroer. First-order methods for large-scale market equilibrium computation. In *NeurIPS 2020*, 2020.
- [Horton and Chilton, 2010] J. J. Horton and L. B. Chilton. The Labor Economics of Paid Crowdsourcing. In *EC’10*, pages 209–218, 2010.
- [Leonardos and Piliouras, 2021] S. Leonardos and G. Piliouras. Exploration-Exploitation in Multi-Agent Learning: Catastrophe Theory Meets Game Theory. In *AAAI*, 2021.
- [Levin *et al.*, 2008] D. Levin, K. LaCurts, N. Spring, and B. Bhattacharjee. Bittorrent is an Auction: Analyzing and Improving Bittorrent’s Incentives. *SIGCOMM Comput. Commun. Rev.*, 38(4):243–254, 2008.
- [Li and Yorke, 1975] T.-Y. Li and J. A. Yorke. Period Three Implies Chaos. *The American Mathematical Monthly*, 82(10):985–992, 1975.
- [López and Vives, 2019] Á. L. López and X. Vives. Overlapping Ownership, R&D Spillovers, and Antitrust Policy. *Journal of Political Economy*, 127(5):2394–2437, 2019.
- [Palaiojanos *et al.*, 2017] G. Palaiojanos, I. Panageas, and G. Piliouras. Multiplicative Weights Update with Constant Step-Size in Congestion Games: Convergence, Limit Cycles and Chaos. In *NIPS’17*, pages 5874–5884, 2017.
- [Puu, 1991] T. Puu. Chaos in duopoly pricing. *Chaos, Solitons & Fractals*, 1(6):573–581, 1991.
- [Sato *et al.*, 2002] Y. Sato, E. Akiyama, and J. D. Farmer. Chaos in Learning a Simple Two-person Game. *PNAS*, 99(7):4748–4751, 2002.
- [Szidarovszky and Okuguchi, 1997] F. Szidarovszky and K. Okuguchi. On the Existence and Uniqueness of Pure Nash Equilibrium in Rent-Seeking Games. *Games and Economic Behavior*, 18(1):135–140, 1997.
- [Theocharis, 1960] R. D. Theocharis. On the Stability of the Cournot Solution on the Oligopoly Problem1. *The Review of Economic Studies*, 27(2):133–134, 02 1960.
- [Varian, 2010] H. R. Varian. *Intermediate Microeconomics: A Modern Approach*. W.W. Norton & Co., New York, eighth edition, 2010.
- [Wärneryd, 2018] K. Wärneryd. Chaotic dynamics in contests. *Economic Inquiry*, 56(3):1486–1491, 2018.
- [Wu and Zhang, 2007] F. Wu and L. Zhang. Proportional Response Dynamics Leads to Market Equilibrium. In *STOC ’07*, pages 354–363, 2007.
- [Zhang, 2011] L. Zhang. Proportional response dynamics in the Fisher market. *Theor. Comput. Sci.*, 412(24):2691–2698, 2011.