
Follow-the-Regularized-Leader Routes to Chaos in Routing Games

Jakub Bielawski¹ Thiparat Chotibut² Fryderyk Falniowski¹ Grzegorz Kosiorowski¹ Michał Misiurewicz³
Georgios Piliouras⁴

Abstract

We study the emergence of chaotic behavior of Follow-the-Regularized Leader (FoReL) dynamics in games. We focus on the effects of increasing the population size or the scale of costs in congestion games, and generalize recent results on unstable, chaotic behaviors in the Multiplicative Weights Update dynamics (Chotibut et al., 2020; 2021; Palaiopoulos et al., 2017) to a much larger class of FoReL dynamics. We establish that, even in simple linear non-atomic congestion games with two parallel links and *any* fixed learning rate, unless the game is fully symmetric, increasing the population size or the scale of costs causes learning dynamics to become unstable and eventually chaotic, in the sense of Li-Yorke and positive topological entropy. Furthermore, we prove the existence of novel non-standard phenomena such as the coexistence of stable Nash equilibria and chaos in the same game. We also observe the simultaneous creation of a chaotic attractor as another chaotic attractor gets destroyed. Lastly, although FoReL dynamics can be strange and non-equilibrating, we prove that the time average still converges to an *exact* equilibrium for any choice of learning rate and any scale of costs.

1. Introduction

We study the dynamics of online learning in a non-atomic repeated congestion game. Namely, every iteration of the game presents a population (i.e., a continuum of players)

¹Department of Mathematics, Cracow University of Economics, Rakowicka 27, 31-510 Kraków, Poland. ²Chula Intelligent and Complex Systems, and Department of Physics, Faculty of Science, Chulalongkorn University, Bangkok 10330, Thailand. ³Department of Mathematical Sciences, Indiana University-Purdue University Indianapolis, 402 N. Blackford Street, Indianapolis, IN 46202, USA. ⁴Engineering Systems and Design, Singapore University of Technology and Design, 8 Somapah Road, Singapore 487372.. Correspondence to: Georgios Piliouras <georgios@sutd.edu.sg>, Thiparat Chotibut <thiparat.c@chula.ac.th>.

with a choice between two strategies, and imposes on them a cost which increases with the fraction of population adopting the same strategy. In each iteration, the players update their strategy accommodating for the outcomes of previous iterations. The structure of cost function here concerns that of the congestion games, which are introduced by Rosenthal (Rosenthal, 1973) and are amongst the most studied classes of games. A seminal result of (Monderer & Shapley, 1996) shows that congestion games are isomorphic to potential games; as such, numerous learning dynamics are known to converge to Nash equilibria (Berenbrink et al., 2014; Even-Dar & Mansour, 2005; Fischer et al., 2006; Fotakis et al., 2008; Kleinberg et al., 2009; 2011; Mertikopoulos & Sandholm, 2018).

A prototypical class of online learning dynamics is Follow the Regularized Leader (FoReL) (Hazan, 2016; Shalev-Shwartz, 2012). FoReL algorithm includes as special cases ubiquitous meta-algorithms, such as the Multiplicative Weights Update (MWU) algorithm (Arora et al., 2012). Under FoReL, the strategy in each iteration is chosen by minimizing the weighted (by the learning rate) sum of the total cost of all actions chosen by the players and the regularization term. FoReL dynamics are known to achieve optimal regret guarantees (i.e., be competitive with the best fixed action with hindsight), as long as they are executed with a highly optimized learning rate; i.e., one that is decreasing with the steepness of the online costs (inverse to the Lipschitz constant of the online cost functions) as well as decreasing with time T at a rate $1/\sqrt{T}$ (Shalev-Shwartz, 2012).

Although precise parameter tuning is perfectly reasonable from the perspective of algorithmic design, it seems implausible from the perspective of behavioral game theory and modeling. For example, experimental and econometric studies based on a behavioral game theoretic learning model known as Experienced Weighted Attraction (EWA), which includes MWU as a special case, have shown that agents can use much larger learning rates than those required for the standard regret bounds to be meaningfully applicable (Camerer, 2011; Ho & Camerer, 1998; 1999; Ho et al., 2007). In some sense, such a tension is to be expected because small and optimized learning rates are designed with system stability and asymptotic optimality in mind,

whereas selfish agents care more about short-term rewards which result in larger learning rates and more aggressive behavioral adaptation. Interestingly, recent work on learning in games study exactly such settings of FoReL dynamics with large, fixed step-sizes, showing that vanishing and even constant regret is possible in some game settings (Bailey & Piliouras, 2019; Bailey et al., 2019).

For congestion games, it is reasonable to expect that increased demands and thus larger daily costs should result in steeper behavioral responses, as agents become increasingly agitated at the mounting delays. However, to capture this behavior we need to move away from the standard assumption of effective scaling down of the learning rate. Then, the costs increase and allow more general models that can incorporate non-vanishing regret. Thus, in this regime, FoReL dynamics in congestion games cannot be reduced to standard regret based analysis (Blum et al., 2006), or even Lyapunov function arguments (e.g., Panageas et al., 2019) and more refined techniques are needed.

In a recent series of papers (Chotibut et al., 2020; 2021; Palaiopanos et al., 2017), the special case of MWU was analyzed under arbitrary population, demands. In a nutshell, for any fixed learning rate, MWU becomes unstable/chaotic even in small congestion games with just two strategies/paths as long as the total demand exceeds some critical threshold, whereas for small population sizes it is always convergent. *Can we extend our understanding from MWU to more general FoReL dynamics? Moreover, are the results qualitatively similar showing that the dynamic is either convergent for all initial conditions or non-convergent for almost all initial conditions, or can there be more complicated behaviors depending on the choice of the regularizer of FoReL dynamics?*

Our model & results. We analyze FoReL-based dynamics with steep regularizers¹ in non-atomic linear congestion games with two strategies. This seemingly simple setting will suffice for the emergence of highly elaborate and unpredictable behavioral patterns. For any such game G and an arbitrarily small but fixed learning rate ϵ , we show that there exists a system capacity $N_0(G, \epsilon)$ such that the system is unstable when the total demand exceeds this threshold. In such case, we observe *complex non-equilibrating dynamics: periodic orbits of any period and chaotic behavior of trajectories* (Section 7). A core technical result is that almost all such congestion games (i.e. unless they are fully symmetric), given sufficiently large demand, will exhibit Li-Yorke chaos and positive topological entropy (Section 7.1). In the case of games with asymmetric equilibrium flow, the bifurcation diagram is very complex (see

Section 8). Li-Yorke chaos implies that there exists an uncountable set of initial conditions that gets scrambled by the dynamics. Formally, given any two initial conditions $x(0), y(0)$ in this set, $\liminf_t \text{dist}(x(t), y(t)) = 0$ while $\limsup_t \text{dist}(x(t), y(t)) > 0$ as t goes to infinity, meaning trajectories come arbitrarily close together infinitely often but also then move away again. In the special case where the two edges have symmetric costs (equilibrium flow is the 50 – 50% split), the system will still become unstable given large enough demand, but chaos is not possible. Instead, in the unstable regime, all but a measure zero set of initial conditions gets attracted by periodic orbits of period two which are symmetric around the equilibrium. Furthermore, we construct formal criteria for when the Nash equilibrium flow is globally attracting. For such systems we can prove their equilibration and thus social optimality even when standard regret bounds are not applicable (Section 6.1). Also, remarkably, whether the system is equilibrating or chaotic, we prove that the time-average flows of FoReL dynamics exhibit regularity and *always converge exactly to the Nash equilibrium* (Section 4).

In Section 8, for the first time, to our knowledge, we report strange dynamics arising from FoReL in congestion games. Firstly, we numerically show that for FoReL dynamics a locally attracting Nash equilibrium and chaos can coexist, see Figure 2. This is also formally proven in Section 6.2. Given the prominence of local stability analysis to equilibria for numerous game theoretic settings which are widely used in Artificial Intelligence, such as Generative Adversarial Networks (GANs), e.g., (Goodfellow et al., 2014; Liang & Stokes, 2019; Mescheder et al., 2018; Nagarajan & Kolter, 2017; Yazici et al., 2018), we believe that this result is of general interest as it reveals that local stability analysis is not sufficient to guard against chaotic behaviors even in a trivial game with just one (locally stable) Nash equilibrium. Secondly, Figure 4 reveals that chaotic attractors can be non-robust. Specifically, we show that mild perturbations in the parameter can lead to the destruction of one complex attractor while another totally distinct complex attractor is born! To the best of our knowledge, these phenomena have never been reported before, and thus expand our understanding of the range of possible behaviors in game dynamics. Several more examples of complex phenomena are provided in Section 8. Finally, further calculations for entropic regularizers can be found in Appendix A.

Our findings suggest that the chaotic behavior of players using Multiplicative Weights Update algorithm in congestion games (see results from (Chotibut et al., 2020; 2021; Palaiopanos et al., 2017)) is not an exception but the rule. Chaos is robust and can be seen for a vast subclass of online learning algorithms. In particular, our results apply to an important subclass of regularizers, of generalized entropies, which are widely used concepts in information theory, com-

¹Steepness of the regularizer guarantees that the dynamics will be well-defined as a function of the current state of the congestion game. For details, see Section 2.

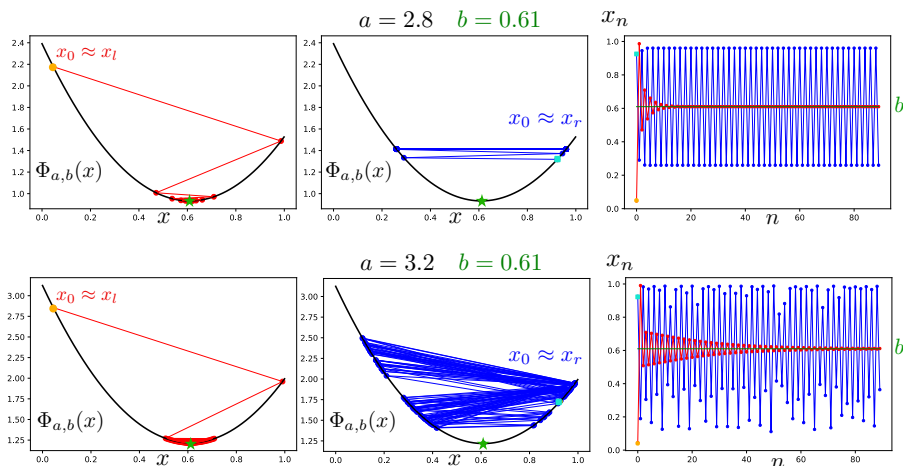


Figure 1. Coexistence of locally attracting Nash equilibrium (green), limit cycles, and chaos in the same congestion game. Since congestion game has an associated convex potential (cost) function $\Phi_{a,b}(x) = \frac{a^2}{2} ((1-b)x^2 + b(1-x)^2)$ with a unique global minimum at the Nash equilibrium b , standard learning algorithms such as gradient-like update with a small step size will converge to the equilibrium. However, here we highlight the unusual coexistence of the attracting Nash equilibrium, limit cycles, and chaos for FoReL dynamics with log-barrier regularizer $r(x) = (1-x)\log(1-x) + x\log(x) - \frac{5}{12}\log(-x^2+x+0.11)$. The right column shows that FoReL dynamics x_n depends on the initial conditions (cyan and orange colors.) Red color encodes the dynamics initialized near the left critical point x_l , which converges to the Nash equilibrium b . Blue color encodes the dynamics initialized near the right critical point x_r , which converge to the limit cycle of period 2 (top), and to chaotic attractors (bottom). Convergence to the Nash equilibrium arises through dynamics that lower the cost function at every successive steps (left column), while convergence to a limit cycle or a chaotic attractor incur large cost, bouncing around in the cost landscape away from the Nash equilibrium. Remarkably, despite being periodic or chaotic, we prove that the time-average of the dynamics converges *exactly* to the Nash equilibrium b , independent of the interior initial conditions. The bifurcation diagram associated with $b = 0.61$ that demonstrates coexistence of multiple attractors in the same game is shown in Fig. 2.

plexity theory, and statistical mechanics (Csiszár, 2008; Słomczyński et al., 2000; Tsallis, 1988). Steep functions (Mertikopoulos & Sandholm, 2016; 2018; Mertikopoulos & Zhou, 2019) and generalized entropies are also often used as regularizers in game-theoretic setting (Bomze et al., 2019; Coucheny et al., 2015; Mertikopoulos & Sandholm, 2018). In particular, Havrda-Charvát-Tsallis entropy-based dynamics was studied, for instance, in (Harper, 2011; Karev & Koonin, 2013). Lastly, the emergence of chaos is clearly a hardness type of result. Such results only increase in strength the simpler the class of examples is. Complicated games are harder to learn and it is harder for players to coordinate on an equilibrium. Thus, in more complicated games one should expect even more complicated, unpredictable behaviors.

2. Model

We consider a two-strategy *congestion game* (see (Rosenthal, 1973)) with a continuum of players (agents), where all of the players apply the *Follow the Regularized Leader* (FoReL) algorithm to update their strategies (Shalev-Shwartz, 2012). Each of the players controls an infinitesimally small fraction of the flow. We assume that the total flow of all the players is equal to N . We denote the fraction

of the players adopting the first strategy at time n as x_n . The second strategy is then chosen by $1 - x_n$ fraction of the players. This model encapsulates how a large population of commuters selects between the two alternative paths that connect the initial point to the end point. When a large fraction of the players adopt the same strategy, congestion arises, and the cost of choosing the same strategy increases.

Linear congestion games: We focus on linear cost functions. Specifically, the cost of each path (link, route, or strategy) is proportional to the *load*. By denoting c_j the cost of selecting the strategy j (when x relative fraction of the agents choose the first strategy),

$$c_1(x) = \alpha Nx, \quad c_2(1-x) = \beta N(1-x), \quad (1)$$

where $\alpha, \beta > 0$ are the coefficients of proportionality. Without loss of generality we will assume throughout the paper that $\alpha + \beta = 1$. Therefore, the values of α and $\beta = 1 - \alpha$ indicate how different the path costs are from each other.

A quantity of interest is the value of the equilibrium split; i.e., the relative fraction of players using the first strategy at equilibrium. The first benefit of this formulation is that the fraction of agents using each strategy at equilibrium is independent of the flow N . The second benefit is that, independent of α, β and N , playing Nash equilibrium results in the optimal social cost, which is the point of contact

with the Price of Anarchy research (Chotibut et al., 2020; Koutsoupias & Papadimitriou, 1999).

2.1. Learning in congestion games with FoReL algorithms

We assume that the players at time $n + 1$ know the cost of the strategies at time n (equivalently, the realized flow (split) $(x_n, 1 - x_n)$) and update their choices according to the *Follow the Regularized Leader* (FoReL) algorithm. Namely, in the period $n + 1$ the players choose the first strategy with probability x_{n+1} such that:

$$\begin{aligned} x_{n+1} = \arg \min_{x \in (0,1)} & \left(\varepsilon \sum_{j \leq n} [c_1(x_j) \cdot x + c_2(1 - x_j) \cdot (1 - x)] \right. \\ & \left. + R(x, 1 - x) \right) = \arg \min_{x \in (0,1)} \left(\varepsilon \sum_{j \leq n} [\alpha N \cdot x_j \cdot x \right. \\ & \left. + \beta N \cdot (1 - x_j) \cdot (1 - x)] + R(x, 1 - x) \right), \end{aligned} \quad (2)$$

where $c_1(x_j) \cdot x + c_2(1 - x_j) \cdot (1 - x)$ is a total cost that is inflicted on the population of agents playing against the mix $(x, 1 - x)$ in period j , while $R: (0, 1)^2 \mapsto \mathbb{R}$ is a *regularizer* which represents a “risk penalty”: namely, that term would penalize abrupt changes of strategy based on a small amount of data from previous iterations of the game. The existence of a regularizer rules out strategies that focus too much on optimizing with respect to the history of our game. A weight coefficient $\varepsilon > 0$ of our choosing is used to balance these two terms and may be perceived as a propensity to learn and try new strategies based on new information: the larger ε is, the faster the players learn and the more eager they are to update their strategies. Commonly adopted as a standard assumption, the learning rate ε can be regarded as a small, fixed constant in the following analysis but its exact value is not of particular interest. Our analysis/results holds for any fixed choice of ε .

Note that FoReL can also be regarded as an instance of an exploration-exploitation dynamics under the multi-armed bandits framework in online learning (Zhao, 2019). In the limit $\varepsilon \gg 1$ such that (2) is well approximated by the minimization of the cumulative expected cost

$$\begin{aligned} & \sum_{j \leq n} [c_1(x_j) \cdot x + c_2(1 - x_j) \cdot (1 - x)] \\ & = \sum_{j \leq n} c_2(1 - x_j) + \left(\sum_{j \leq n} [c_1(x_j) - c_2(1 - x_j)] \right) x, \end{aligned}$$

the minimization yields

$$x_{n+1} = \begin{cases} 0, & \sum_{j \leq n} [c_1(x_j) - c_2(1 - x_j)] > 0, \\ 1, & \sum_{j \leq n} [c_1(x_j) - c_2(1 - x_j)] < 0. \end{cases}$$

Namely, the strategy that incurs the *least* cumulative cost in the past time horizon is selected with probability 1. This term thus represents *exploitation* dynamics in reinforcement learning and multi-armed bandits framework. In the opposite limit when $\varepsilon \ll 1$, (2) is well approximated by the minimization of the regularizer $R(x, 1 - x)$. For the Shannon entropy regularizer $R(x, 1 - x) = -H_S(x, 1 - x) = x \log x + (1 - x) \log(1 - x)$ that results in the Multiplicative Weight Update algorithm (see the details in Appendix A and Sec. 3), its minimization yields

$$x_{n+1} = (1 - x_{n+1}) = 1/2.$$

The entropic regularization term tends to explore every strategy with equal probabilities, neglecting the information of the past cumulative cost. Thus, this regularization term corresponds to *exploration* dynamics. Therefore, ε adjusts the tradeoff between *exploration* and *exploitation*. The continuous time variant of (2) with the Shannon entropy regularizer has been studied as models of collective adaption (Sato & Crutchfield, 2003; Sato et al., 2005), also known as Boltzmann Q learning (Kianercy & Galstyan, 2012), in which the exploitation term is interpreted as behavioral *adaptation* whereas the exploration term represents *memory loss*. More recent continuous-time variants study generalized entropies as regularizers, leading to a larger class of dynamics called Escort Replicator Dynamics (Harper, 2011) which was analyzed extensively in (Mertikopoulos & Sandholm, 2016; 2018).

Motivated by the continuous-time dynamics with generalized entropies, we extend FoReL discrete-time dynamics (2) to a larger class of regularizers. For a given regularizer R , we define an auxiliary function:

$$r: (0, 1) \ni x \mapsto R(x, 1 - x) \in \mathbb{R}. \quad (3)$$

We restrict the analysis to a FoReL class of regularizers for which the dynamics implied by the algorithm is well-defined. Henceforth, we assume that R is a steep symmetric convex regularizer, namely $R \in \mathbf{SSC}$, where:

$$\begin{aligned} \mathbf{SSC} = \{ & R \in \mathcal{C}^2((0, 1)^2) : \forall (x, y) \in (0, 1)^2 R(y, x) = R(x, y); \\ & \forall x \in (0, 1) r''(x) > 0; \lim_{x \rightarrow 0^+} r'(x) = -\infty \}. \end{aligned}$$

These conditions on regularizers are not overly restrictive: the assumptions for convexity and symmetry of the regularizer are natural, and if $\lim_{x \rightarrow 0} r'(x)$ is finite, then the dynamics of x_n from (2) will not be well-defined.

Many well-known and widely used regularizers like (negative) Arimoto entropies (Shannon entropy, Havrda-Charvát-Tsallis (HCT) entropies and log-barrier being most famous ones) and (negative) Rényi entropies, under mild assumptions, belong to \mathbf{SSC} (see Appendix A). A standard non-example is the square of the Euclidean norm $R(x, 1 - x) = x^2 + (1 - x)^2$.

3. The dynamics introduced by FoReL

Let $R \in \text{SSC}$. Assume that up to the iteration $n > 0$ the trajectory $(x_0, x_1, \dots, x_{n-1})$ was established by (2). Then

$$x_n = \arg \min_{x \in (0,1)} \left(N\varepsilon \sum_{j \leq n-1} [\alpha x_j x + \beta(1-x_j)(1-x)] + r(x) \right).$$

First order condition yields

$$r'(x_n) = -N\varepsilon \sum_{j \leq n-1} [\alpha x_j - \beta(1-x_j)].$$

We know that r is convex, therefore the sufficient and necessary condition for x_{n+1} to satisfy (2) takes form:

$$\begin{aligned} r'(x_{n+1}) &= -N\varepsilon \sum_{j \leq n} [\alpha x_j - \beta(1-x_j)] \\ &= r'(x_n) - N\varepsilon [\alpha x_n - \beta(1-x_n)] \\ &= r'(x_n) - N\varepsilon [x_n - \beta]. \end{aligned} \quad (4)$$

We define $\Psi: (0, 1) \ni x \mapsto -r'(x) \in \mathbb{R}$. Table 1 depicts functions Ψ for different entropic regularizers². Before proceeding any further, we need to establish crucial properties of the function Ψ .

Proposition 3.1. *Let Ψ be a function derived from a regularizer from SSC. Then*

- i) $\Psi(1-x) = -\Psi(x)$ for $x \in (0, 1)$.
- ii) Ψ is a homeomorphism, $\lim_{x \rightarrow 0^+} \Psi(x) = \infty$, and $\lim_{x \rightarrow 1^-} \Psi(x) = -\infty$.

By Proposition 3.1.ii, Ψ is a homeomorphism between $(0, 1)$ and \mathbb{R} .

After substituting

$$a = N\varepsilon, \quad b = \beta \quad (5)$$

we obtain from (4) a general formula for the dynamics

$$x_{n+1} = \Psi^{-1}(\Psi(x_n) + a(x_n - b)), \quad (6)$$

where $a > 0, b \in (0, 1)$. Thus, we introduce $f_{a,b}: [0, 1] \mapsto [0, 1]$ as

$$f_{a,b}(x) = \begin{cases} 0, & x = 0 \\ \Psi^{-1}(\Psi(x) + a(x - b)), & x \in (0, 1) \\ 1, & x = 1 \end{cases} \quad (7)$$

²By substituting the negative Shannon entropy as r in (4), that is $r(x) = R(x, 1-x) = x \log x + (1-x) \log(1-x)$, we obtain the Multiplicative Weights Update algorithm.

By the properties of Ψ , $f_{a,b}: [0, 1] \mapsto [0, 1]$ is continuous, and (7) defines a discrete dynamical system emerging from the FoReL algorithm for the pair of parameters (a, b) .

Lemma 3.2. *The following properties hold:*

- i) $f_{a,b}(x) > x$ if and only if $x < b$ and $f_{a,b}(x) < x$ if and only if $x > b$.
- ii) If $\varphi: (0, 1) \mapsto (0, 1)$ is given by $\varphi(x) = 1-x$, then $\varphi \circ f_{a,b} = f_{a,1-b} \circ \varphi$.
- iii) Under the dynamics defined by (7), there exists a closed invariant and globally attracting interval $I \subset (0, 1)$.

4. Average behavior — Nash equilibrium is Cesáro attracting

We start by studying asymptotic behavior by looking on the average behavior of orbits. We will show that the orbits of our dynamics exhibit regular average behavior known as Cesáro attraction to the Nash equilibrium b .

Definition 4.1. *For an interval map f a point p is Cesáro attracting if there is a neighborhood U of p such that for every $x \in U$ the averages $\frac{1}{n} \sum_{k=0}^{n-1} f^k(x)$ converge to p .*

Theorem 4.2 (Cesáro attracting). *For every $a > 0, b \in (0, 1)$ and $x_0 \in (0, 1)$ we have*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f_{a,b}^k(x_0) = b. \quad (8)$$

Corollary 4.3. *The center of mass of any periodic orbit $\{x_0, x_1, \dots, x_{n-1}\}$ of $f_{a,b}$ in $(0, 1)$, namely $\frac{x_0 + x_1 + \dots + x_{n-1}}{n}$, is equal to b .*

Applying the Birkhoff Ergodic Theorem, we obtain:

Corollary 4.4. *For every probability measure μ , invariant for $f_{a,b}$ and such that $\mu(\{0, 1\}) = 0$, we have*

$$\int_{[0,1]} x \, d\mu = b.$$

In the following sections, we will show that, despite the regularity of the average trajectories which converge to the Nash equilibrium b , the trajectories themselves typically exhibit complex and diverse behaviors.

5. Two definitions of chaos

In this section we introduce two notions of chaotic behavior: Li-Yorke chaos and (positive) topological entropy. Most definitions of chaos focus on complex behavior of trajectories, such as Li-Yorke chaos or fast growth of the number of distinguishable orbits of length n , detected by positivity of the topological entropy.

Table 1. Homeomorphisms Ψ for regularizers from SSC.

regularizer	$r(x)$	$\Psi(x)$
Shannon	$x \log x + (1-x) \log(1-x)$	$\log \frac{1-x}{x}$
Havrda-Charvát-Tsallis, $q \in (0, 1)$	$\frac{1}{1-q} (1-x^q - (1-x)^q)$	$\frac{q}{1-q} (x^{q-1} - (1-x)^{q-1})$
Rényi, $q \in (0, 1)$	$\frac{1}{q-1} \log(x^q + (1-x)^q)$	$\frac{q}{1-q} \cdot \frac{x^{q-1} - (1-x)^{q-1}}{x^q + (1-x)^q}$
log-barrier	$-\log x - \log(1-x)$	$\frac{1}{x} - \frac{1}{1-x}$

Definition 5.1 (Li-Yorke chaos). *Let (X, f) be a dynamical system and $(x, y) \in X \times X$. We say that (x, y) is a Li-Yorke pair if $\liminf_{n \rightarrow \infty} \text{dist}(f^n(x), f^n(y)) = 0$, and $\limsup_{n \rightarrow \infty} \text{dist}(f^n(x), f^n(y)) > 0$. A dynamical system (X, f) is Li-Yorke chaotic if there is an uncountable set $S \subset X$ (called scrambled set) such that every pair (x, y) with $x, y \in S$ and $x \neq y$ is a Li-Yorke pair.*

Intuitively orbits of two points from the scrambled set have to gather themselves arbitrarily close and spring aside infinitely many times but (if X is compact) it cannot happen simultaneously for each pair of points. Obviously the existence of a large scrambled set implies that orbits of points behave in unpredictable, complex way.

A crucial feature of the chaotic behavior of a dynamical system is also exponential growth of the number of distinguishable orbits. This happens if and only if the topological entropy of the system is positive. In fact positivity of topological entropy turned out to be an essential criterion of chaos (Glasner & Weiss, 1993). This choice comes from the fact that the future of a deterministic (zero entropy) dynamical system can be predicted if its past is known (see (Weiss, 2000)) and positive entropy is related to randomness and chaos. For every dynamical system over a compact phase space, we can define a number $h(f) \in [0, \infty]$ called the *topological entropy* of transformation f . This quantity was first introduced by Adler, Konheim and McAndrew (Adler et al., 1965) as the topological counterpart of a metric (and Shannon) entropy. In general, computing topological entropy is not an easy task. However, in the context of piecewise monotone interval maps, topological entropy is equal to the exponential growth rate of the minimal number of monotone subintervals for f^n .

Theorem 5.2 ((Misiurewicz & Szlenk, 1980)). *Let f be a piecewise monotone interval map and, for all $n \geq 1$, let m_n be the minimal cardinality of a monotone partition for f^n . Then $h(f) = \lim_{n \rightarrow \infty} \frac{1}{n} \log m_n = \inf_{n \geq 1} \frac{1}{n} \log m_n$.*

6. Asymptotic stability of Nash equilibria

6.1. Asymptotic stability of Nash equilibria

The dynamics induced by (7) admits three fixed points: 0, 1 and b . By Lemma 3.2.iii we know that all orbits starting from $(0, 1)$ eventually fall into a globally attracting interval

I . Thus, the points 0 and 1 are repelling. When does the Nash equilibrium b attract all point from $(0, 1)$? First, we look when b is an attracting and when it is a repelling fixed point. With this aim, we study the derivative of $f_{a,b}$:

$$f'_{a,b}(x) = (\Psi^{-1})'(\Psi(x) + a(x-b)) \cdot (\Psi'(x) + a).$$

Then,

$$f'_{a,b}(b) = (\Psi^{-1})'(\Psi(b)) \cdot (\Psi'(b) + a) = \frac{\Psi'(b) + a}{\Psi'(b)}. \quad (9)$$

The fixed point b is attracting if and only if $|f'_{a,b}(b)| < 1$, which is equivalent to the condition:

$$|\Psi'(b) + a| < -\Psi'(b). \quad (10)$$

Thus, the fixed point b is attracting if and only if $a \in (0, -2 \cdot \Psi'(b))$ and repelling otherwise. We will answer when b is globally attracting on $(0, 1)$. First we will show the following auxiliary lemma.

Lemma 6.1. *Let a function $g: I \mapsto \mathbb{R}$ be such that $g''' < 0$. Then*

$$g' \left(\frac{x+y}{2} \right) > \frac{g(x) - g(y)}{x-y}$$

for every $x, y \in I$.

The following theorem answers, whether b is globally attracting on $(0, 1)$.

Theorem 6.2. *Let Ψ be a homeomorphism derived from a regularizer from SSC. Suppose that b is an attracting fixed point of $f_{a,b}$. If $\Psi''' < 0$, then trajectories of all points from $(0, 1)$ converge to b .*

Corollary 6.3. *Let $\Psi''' < 0$. Then the Nash equilibrium b attracts all points from the open interval $(0, 1)$ if and only if $a \in (0, -2\Psi'(b))$.*

Functions Ψ derived from Shannon entropy, HCT entropy or log-barrier satisfy the inequality $\Psi''' < 0$. Nevertheless, this additional condition is needed, because for an arbitrary Ψ derived from SSC attracting orbits of any period may exist together with the attracting Nash equilibrium b . In the next section we will discuss thoroughly an example of such behavior. This shows that even for the well-known class of FoReL algorithms knowledge of local behavior (even attraction) of the Nash equilibrium may not be enough to properly describe behavior of agents.




Figure 2. Coexistence of the attracting Nash equilibrium and chaos. The bifurcation diagrams for $f_{a,b}$ where the dynamics is induced by the regularizer $r(x) = (1-x)\log(1-x) + x\log x - 0.4167\log(-x^2 + x + 0.11)$ for $b = 0.61$. On the horizontal axis the parameter a is between 2.6 and 3.4, and on the vertical axis values of $f_{a,b}$ are shown. As starting points for bifurcation diagrams two critical points of $f_{a,b}$ are taken — red refers to the critical point in $(0, 0.5)$ and blue the critical point in $(0.5, 1)$. Each critical point is iterated 4000 times, visualizing the last 200 iterates. On the top picture first red and then blue trajectories are drawn, and on the bottom one the order is reversed. Function $\Psi(x) = -r'(x) = \log(1-x) - \log x + 0.4167\left[\frac{1}{1.1-x} - \frac{1}{x+0.1}\right]$ fulfills all assumptions of Theorem 6.2 excluding $\Psi''' < 0$. Although for $a < -2\Psi'(b) \approx 3.28$ the unique Nash equilibrium is attracting we can observe chaotic behavior already for $a > 3.15$. The picture suggests that in the coexistence region we have an interval which is invariant for $f_{a,b}^2$, and in it we see the usual evolution of unimodal maps. This means that sometimes we see an attracting periodic orbit, sometimes a chaotic attractor.

6.2. Coexistence of attracting Nash equilibrium and chaos

In this section we will describe an example of the regularizer from SSC, which introduces game dynamics in which attracting Nash equilibrium coexist with chaos, see example in Figure 2. This phenomenon is observed by replacing the Shannon entropic regularizer by the log-barrier regularizer. Namely, we take $\Psi(x) = \log(1-x) - \log x + 0.4167 \cdot \left(\frac{1}{1.1-x} - \frac{1}{x+0.1}\right)$.

We will show that there exist $a > 0, b \in (0, 1)$ such that $f_{a,b}$ has an attracting fixed point (which is the Nash equilibrium) yet the map can be chaotic!

Proposition 6.4. *There exist $a > 0, b \in (0, 1)$ such that $f_{a,b}$ has an attracting fixed point (Nash equilibrium), positive topological entropy and is Li-Yorke chaotic.ⁱ*

Corollary 6.5. *There exist FoReL dynamics such that when applied to symmetric linear congestion games with only two strategies/paths the resulting dynamics have i) a set of positive measure of initial conditions that converge to the unique and socially optimum Nash equilibrium, ii) an uncountable scrambled set for which trajectories exhibit Li-Yorke chaos, and iii) periodic orbits of all possible even periods. Thus, the (long-term) social cost depends critically on the initial condition.*

FoReL dynamics induced by this regularizer manifests drastically different behaviors that depend on the initial condition.

7. Behavior for sufficiently large a

7.1. Non-convergence for sufficiently large a

In this subsection, we study what happens as we fix b and let a be arbitrarily large³. First, we study the asymmetric case, namely $b \neq 1/2$. We show chaotic behavior of our dynamical system for a sufficiently large, that is we will show that if a is sufficiently large then $f_{a,b}$ is Li-Yorke chaotic, has periodic orbits of all periods and positive topological entropy. The crucial ingredient of our analysis is the existence of periodic orbit of period 3.

Theorem 7.1. *If $b \in (0, 1) \setminus \{\frac{1}{2}\}$, then there exists a_b such that if $a > a_b$ then $f_{a,b}$ has a periodic point of period 3.*

By the Sharkovsky Theorem (Sharkovsky, 1964), existence of a periodic orbit of period 3 implies existence of periodic orbits of all periods, and by the result of (Li & Yorke, 1975), period 3 implies Li-Yorke chaos. Moreover, because $f_{a,b}$ has a periodic point of period that is not a power of 2, the topological entropy $h(f_{a,b})$ is positive (see (Misiurewicz,

³By (5) it reflects the case when we fix cost functions (and learning rate ε) and increase the total demand N .

1979)). Thus:

Corollary 7.2. *If $b \in (0, 1) \setminus \{1/2\}$, then there exists a_b such that if $a > a_b$ then $f_{a,b}$ has periodic orbits of all periods, has positive topological entropy and is Li-Yorke chaotic.*

This result has an implication in non-atomic routing games. Recall that the parameter a expresses the normalized total demand. Thus, Corollary 7.2 implies that when the cost functions of paths are different, then increasing the total demand of the system will inevitably lead to chaotic behavior.

Now we consider the symmetric case, when $b = \frac{1}{2}$, which corresponds to equal coefficients of the cost functions, $\alpha = \beta$. To simplify the notation we denote $f_a = f_{a,1/2}$.

Theorem 7.3. *If the parameter a is small enough, then all trajectories of f_a starting from $(0, 1)$ converge to the attracting fixed point $1/2$. There exists a_b such that if $a > a_b$, then all points from $(0, 1)$ (except countably many points, whose trajectories eventually fall into the repelling point $1/2$) are attracted by periodic attracting orbits of the form $\{\sigma_a, 1 - \sigma_a\}$, where $0 < \sigma_a < 1/2$. Moreover, if there exists $\delta > 0$ such that Ψ is convex on $(0, \delta)$, then there exists a unique attracting orbit $\{\sigma_a, 1 - \sigma_a\}$, which attracts trajectories of all points from $(0, 1)$, except countably many points, whose trajectories eventually fall into the repelling fixed point $1/2$.*

Theorem 7.3 implies that if the cost functions are equal, there is a threshold such that if the total demand exceeds it, then starting from almost any initial state the system will converge to a symmetric periodic orbit of period 2.

8. Experimental results

In this section we report complex behaviors in bifurcation diagrams of FoReL dynamics. All figures can be found in the Appendix. We investigate the structures of the attracting periodic orbits and chaotic attractors associated with the interval map $f_{a,b} : [0, 1] \mapsto [0, 1]$ defined by (7). In the asymmetric case, that is when b differs from 0.5, the standard equilibrium analysis applies when the fixed point b is stable, which is when $|f'_{a,b}(b)| \leq 1$, or equivalently when $a \leq -2\Psi'(b)$. Therefore, as we argued in the previous section, in this case the dynamics will converge toward the fixed point b whenever $a < -2\Psi'(b)$. However, when $a \geq -2\Psi'(b)$ there is no attracting fixed point. Moreover, a chaotic behavior of trajectories emerges when a is sufficiently large, as the period-doubling bifurcations route to chaos is guaranteed to arise.

In particular, we study the attractors of the map $f_{a,b}$ generated by the log-barrier regularizer (see Example A.2 with $\eta(x) = \log x$) and by the Havrda-Charvát-Tsallis regularizer for $q = 0.5$ (see Example A.3). Note that for both of

these regularizers, we have that $\Psi''' < 0$. Note also that the functions Ψ for these regularizers can be found in Table 1.

We first focus on the log-barrier regularizer⁴. Figure 4 reveals an unusual bifurcation phenomenon, which, to our knowledge, is not known in other natural interval maps. We observe simultaneous evolution of two attractors in the opposite directions: one attractor, generated by the trajectory of the left critical point, is shrinking, while the other one, generated by the trajectory of the right critical point, is growing. Figure 5 shows another unusual bifurcation phenomenon: a chaotic attractor arises via period-doubling bifurcations and then collapses. After that, the trajectories of the critical points, one after the other, jump, and then they together follow a period-doubling route to chaos once more.

Finally we study the bifurcation diagrams generated from Havrda-Charvát-Tsallis regularizer with $q = 0.5$. In Figure 6 we observe a finite number of period-doubling (and period-halving) bifurcations, a behavior that does not lead to chaos. Nevertheless, as a increases from 39.915 to 39.93, the trajectory of the right critical point leaves the attractor which it shared with the trajectory of the left critical point, and builds a separate chaotic attractor. When chaos arises, however, we observe that the induced dynamics of the log-barrier regularizer and of the Havrda-Charvát-Tsallis regularizer with $q = 0.5$ both exhibit period-doubling routes to chaos though the regularizers are starkly different, see Figure 7 and Figure 8 respectively.

9. Conclusion

Recently, there has been intense interest in understanding unstable, chaotic behaviors of learning dynamics in games. For example, the inability of continuous-time FoReL dynamics to converge to mixed Nash equilibria in normal form games has recently been established in (Vlatakis-Gkaragkounis et al., 2020); however, no results about chaos were shown. Moreover, (Galla & Farmer, 2013; Pangallo et al., 2019) showcase empirically that numerous learning dynamics exhibit instability in different games. However, once again, no formal proofs of Li-Yorke chaotic behaviors are established.

In this work, we study FoReL dynamics in non-atomic congestion games with arbitrarily small but fixed step-sizes, rather than with decreasing and regret-optimizing step-sizes. We show that, under sufficiently large demand, dynamics will unavoidably become chaotic and unpredictable. Our work generalizes previous results that hold in the special case of Multiplicative Weights Update (Chotibut et al., 2020;

⁴From the regularity of the map $f_{a,b}$ (see Appendix B), we know that every limit cycle of the dynamics generated by $f_{a,b}$ can be found by studying the behavior of the critical points of $f_{a,b}$. Therefore, all attractors of this dynamics can be revealed by following the trajectories of these two critical points (as in Figure 4).

2021; Palaiopoulos et al., 2017). We also provide a variety of undocumented complex behaviors such as the co-existence of a locally attracting Nash equilibrium and of chaos in the same game. Despite this complexity of the day-to-day behavior, the time-average system behavior is *always* perfectly regular, converging to an exact equilibrium. Our analysis showcases that local stability in games should not be considered as a foregone conclusion and paves the way toward further investigations at the intersection of optimization theory, (behavioral) game theory, and dynamical systems.

Acknowledgements

Georgios Piliouras acknowledges supports 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). Fryderyk Falniowski acknowledges the support of the National Science Centre, Poland, grant 2016/21/D/HS4/01798 and COST Action CA16228 “European Network for Game Theory.” Research of Michał Misiurewicz was partially supported by grant number 426602 from the Simons Foundation. Jakub Bielawski, Fryderyk Falniowski, and Grzegorz Kosiorowski acknowledge support from a subsidy granted to Cracow University of Economics. Thiparat Chotibut acknowledges a fruitful discussion with Tanapat Deesuwan on generalized entropies, and was partially supported by grants for development of new faculty staff, Ratchadaphiseksomphot endowment fund, and Sci-Super VI fund, Chulalongkorn University.

References

- Adler, R. L., Konheim, A. G., and McAndrew, M. H. Topological entropy. *Transactions of American Mathematical Society*, 114:309–319, 1965.
- Arora, S., Hazan, E., and Kale, S. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- Bailey, J. P. and Piliouras, G. Fast and furious learning in zero-sum games: Vanishing regret with non-vanishing step sizes. In *Advances in Neural Information Processing Systems*, volume 32, pp. 12977–12987, 2019.
- Bailey, J. P., Gidel, G., and Piliouras, G. Finite regret and cycles with fixed step-size via alternating gradient descent-ascent. *CoRR*, abs/1907.04392, 2019.
- Berenbrink, P., Hoefer, M., and Sauerwald, T. Distributed selfish load balancing on networks. In *ACM Transactions on Algorithms (TALG)*, 2014.
- Block, L. and Coppel, W. A. *Dynamics in one dimension*, volume 513 of *Lecture Notes in Mathematics*. Springer, Berlin New York, 2006.
- Blum, A., Even-Dar, E., and Ligett, K. Routing without regret: On convergence to nash equilibria of regret-minimizing algorithms in routing games. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing*, pp. 45–52. ACM, 2006.
- Bomze, I. M., Mertikopoulos, P., Schachinger, W., and Staudigl, M. Hessian barrier algorithms for linearly constrained optimization problems. *SIAM Journal on Optimization*, 29(3):2100–2127, 2019.
- Camerer, C. F. *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, 2011.
- Chotibut, T., Falniowski, F., Misiurewicz, M., and Piliouras, G. The route to chaos in routing games: When is price of anarchy too optimistic? *Advances in Neural Information Processing Systems*, 33, 2020.
- Chotibut, T., Falniowski, F., Misiurewicz, M., and Piliouras, G. Family of chaotic maps from game theory. *Dynamical Systems*, 36(1):48–63, 2021. doi: 10.1080/14689367.2020.1795624. URL <https://doi.org/10.1080/14689367.2020.1795624>.
- Coucheney, P., Gaujal, B., and Mertikopoulos, P. Penalty-regulated dynamics and robust learning procedures in games. *Mathematics of Operations Research*, 40(3):611–633, 2015.
- Csiszár, I. Axiomatic characterizations of information measures. *Entropy*, 10(3):261–273, 2008.
- Even-Dar, E. and Mansour, Y. Fast convergence of selfish rerouting. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '05*, pp. 772–781, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics. ISBN 0-89871-585-7.
- Fischer, S., Räcke, H., and Vöcking, B. Fast convergence to wardrop equilibria by adaptive sampling methods. In *Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing, STOC '06*, pp. 653–662, New York, NY, USA, 2006. ACM. ISBN 1-59593-134-1.
- Fotakis, D., Kaporis, A. C., and Spirakis, P. G. Atomic congestion games: Fast, myopic and concurrent. In Monien, B. and Schroeder, U.-P. (eds.), *Algorithmic Game Theory*, volume 4997 of *Lecture Notes in Computer Science*, pp. 121–132. Springer Berlin Heidelberg, 2008. ISBN 978-3-540-79308-3.

- Galla, T. and Farmer, J. D. Complex dynamics in learning complicated games. *Proceedings of the National Academy of Sciences*, 110(4):1232–1236, 2013. ISSN 0027-8424.
- Glasner, E. and Weiss, B. Sensitive dependence on initial conditions. *Nonlinearity*, 6(6):1067–1085, 1993.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.
- Grünwald, P. D. and Dawid, A. P. Game theory, maximum entropy, minimum discrepancy and robust bayesian decision theory. *The Annals of Statistics*, 32(4):1367–1433, 2004.
- Harper, M. Escort evolutionary game theory. *Physica D: Nonlinear Phenomena*, 240(18):1411–1415, 2011.
- Havrda, J. and Charvát, F. Quantification method of classification processes. concept of structural α -entropy. *Kybernetika*, 3(1):30–35, 1967.
- Hazan, E. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.
- Ho, T.-H. and Camerer, C. Experience-weighted attraction learning in coordination games: Probability rules, heterogeneity, and time-variation. *Journal of Mathematical Psychology*, 42:305–326, 1998.
- Ho, T.-H. and Camerer, C. Experience-weighted attraction learning in normal form games. *Econometrica*, 67:827–874, 1999.
- Ho, T.-H., Camerer, C. F., and Chong, J.-K. Self-tuning experience weighted attraction learning in games. *Journal of Economic Theory*, 133:177–198, 2007.
- Karev, G. P. and Koonin, E. V. Parabolic replicator dynamics and the principle of minimum Tsallis information gain. *Biology Direct*, 8(1):19, 2013.
- Kianercy, A. and Galstyan, A. Dynamics of Boltzmann Q learning in two-player two-action games. *Phys. Rev. E*, 85:041145, Apr 2012. doi: 10.1103/PhysRevE.85.041145. URL <https://link.aps.org/doi/10.1103/PhysRevE.85.041145>.
- Kleinberg, R., Piliouras, G., and Tardos, É. Multiplicative updates outperform generic no-regret learning in congestion games. In *ACM Symposium on Theory of Computing (STOC)*, 2009.
- Kleinberg, R., Piliouras, G., and Tardos, É. Load balancing without regret in the bulletin board model. *Distributed Computing*, 24(1):21–29, 2011.
- Koutsoupias, E. and Papadimitriou, C. Worst-case equilibria. In Meinel, C. and Tison, S. (eds.), *STACS 99*, pp. 404–413, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg. ISBN 978-3-540-49116-3.
- Li, T. Y. and Yorke, J. A. Period three implies chaos. *The American Mathematical Monthly*, 82:985–992, 1975.
- Li, T. Y., Misiurewicz, M., Pianigiani, G., and Yorke, J. A. Odd chaos. *Physics Letters A*, 87(6):271–273, 1982.
- Liang, T. and Stokes, J. Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 907–915. PMLR, 2019.
- Mertikopoulos, P. and Sandholm, W. H. Learning in games via reinforcement and regularization. *Mathematics of Operations Research*, 41(4):1297–1324, 2016.
- Mertikopoulos, P. and Sandholm, W. H. Riemannian game dynamics. *Journal of Economic Theory*, 177:315–364, 2018.
- Mertikopoulos, P. and Zhou, Z. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173(1-2):465–507, 2019.
- Mescheder, L., Geiger, A., and Nowozin, S. Which training methods for gans do actually converge? *arXiv preprint arXiv:1801.04406*, 2018.
- Misiurewicz, M. Horseshoes for mapping of the interval. *Bull. Acad. Polon. Sci. Sér. Sci.*, 27:167–169, 1979.
- Misiurewicz, M. and Szlenk, W. Entropy of piecewise monotone mappings. *Studia Mathematica*, 67(1):45–63, 1980.
- Monderer, D. and Shapley, L. S. Fictitious play property for games with identical interests. *Journal of Economic Theory*, 68(1):258–265, 1996.
- Nagarajan, V. and Kolter, J. Z. Gradient descent gan optimization is locally stable. In *Advances in Neural Information Processing Systems*, pp. 5585–5595, 2017.
- Palaiopanos, G., Panageas, I., and Piliouras, G. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. In *Advances in Neural Information Processing Systems*, pp. 5872–5882, 2017.

- Panageas, I., Piliouras, G., and Wang, X. Multiplicative weights updates as a distributed constrained optimization algorithm: Convergence to second-order stationary points almost always. In *International Conference on Machine Learning*, pp. 4961–4969. PMLR, 2019.
- Pangallo, M., Heinrich, T., and Farmer, J. D. Best reply structure and equilibrium convergence in generic games. *Science advances*, 5(2):eaat1328, 2019.
- Rényi, A. On measures of entropy and information. *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, 1961.
- Rosenthal, R. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- Sato, Y. and Crutchfield, J. P. Coupled replicator equations for the dynamics of learning in multiagent systems. *Phys. Rev. E*, 67:015206, Jan 2003. doi: 10.1103/PhysRevE.67.015206. URL <https://link.aps.org/doi/10.1103/PhysRevE.67.015206>.
- Sato, Y., Akiyama, E., and Crutchfield, J. P. Stability and diversity in collective adaptation. *Physica D: Nonlinear Phenomena*, 210(1):21 – 57, 2005. ISSN 0167-2789. doi: <https://doi.org/10.1016/j.physd.2005.06.031>. URL <http://www.sciencedirect.com/science/article/pii/S0167278905002708>.
- Shalev-Shwartz, S. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.
- Sharkovsky, A. N. Coexistence of the cycles of a continuous mapping of the line into itself. *Ukrain. Math. Zh.*, 16: 61–71, 1964.
- Słomczyński, W., Kwapien, J., and Życzkowski, K. Entropy computing via integration over fractal measures. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 10(1): 180–188, 2000.
- Tsallis, C. Possible generalization of Boltzmann-Gibbs statistics. *Journal of Statistical Physics*, 52(1-2):479–487, 1988.
- Vlatakis-Gkaragkounis, E.-V., Flokas, L., Mertikopoulos, P., and Piliouras, G. No-regret learning and mixed Nash equilibria: They do not mix. In *Annual Conference on Neural Information Processing Systems*, 2020.
- Weiss, B. *Single orbit dynamics*, volume 95 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, RI, 2000.
- Yazici, Y., Foo, C.-S., Winkler, S., Yap, K.-H., Piliouras, G., Chandrasekhar, V., et al. The unusual effectiveness of averaging in gan training. In *International Conference on Learning Representations*, 2018.
- Zhao, Q. *Multi-Armed Bandits: Theory and Applications to Online Learning in Networks*. Morgan and Claypool Publishers, 2019.