



HAL
open science

On the discrete-time origins of the replicator dynamics: From convergence to instability and chaos

Fryderyk Falniowski, Panayotis Mertikopoulos

► To cite this version:

Fryderyk Falniowski, Panayotis Mertikopoulos. On the discrete-time origins of the replicator dynamics: From convergence to instability and chaos. 2024. hal-04629311

HAL Id: hal-04629311

<https://hal.science/hal-04629311v1>

Preprint submitted on 29 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

ON THE DISCRETE-TIME ORIGINS OF THE REPLICATOR DYNAMICS: FROM CONVERGENCE TO INSTABILITY AND CHAOS

FRYDERYK FALNIOWSKI* AND PANAYOTIS MERTIKOPOULOS[‡]

ABSTRACT. We consider three distinct discrete-time models of learning and evolution in games: a biological model based on intra-species selective pressure, the dynamics induced by pairwise proportional imitation, and the exponential / multiplicative weights algorithm for online learning. Even though these models share the same continuous-time limit – the replicator dynamics – we show that second-order effects play a crucial role and may lead to drastically different behaviors in each model, even in very simple, symmetric 2×2 games. Specifically, we study the resulting discrete-time dynamics in a class of parametrized congestion games, and we show that (i) in the biological model of intra-species competition, the dynamics remain convergent for any parameter value; (ii) the dynamics of pairwise proportional imitation exhibit an entire range of behaviors for larger time steps and different equilibrium configurations (stability, instability, and even Li-Yorke chaos); while (iii) in the exponential / multiplicative weights algorithm, increasing the time step (almost) inevitably leads to chaos (again, in the formal, Li-Yorke sense). This divergence of behaviors comes in stark contrast to the globally convergent behavior of the replicator dynamics, and serves to delineate the extent to which the replicator dynamics provide a useful predictor for the long-run behavior of their discrete-time origins.

1. INTRODUCTION

Ever since it was introduced by Nash [46], the notion of a Nash equilibrium and its refinements have remained among the most prominent solution concepts of noncooperative game theory. As such, one of the most fundamental questions in the field has been to specify whether – and under what conditions – players eventually end up emulating an equilibrium (or equilibrium-like) behavior through repeated interactions; put differently, whether a dynamic process driven by the agents’ individual interests converges to a rational outcome, in which (classes of) games, etc.

Historically, this question fueled the intense interest in game dynamics brought about by the inception of evolutionary game theory in the mid-1970’s, then the surge of activity that followed in the field of economic theory in the 1990’s, and, more recently, through various connections to machine learning and artificial intelligence, in theoretical computer science. Accordingly, depending on the context, game dynamics are usually derived in one of the following ways: (i) from a biological model of population evolution, typically phrased in terms of the reproductive fitness of the species involved; (ii) from a set of economic microfoundations that express the growth rate of a type (or strategy) within a population via a *revision protocol* (an economic model prescribing an agent’s propensity to switch to a better-performing strategy); or (iii) from some learning algorithm designed to optimize a myopic performance criterion (such as the minimization of an agent’s regret),

* DEPARTMENT OF MATHEMATICS, KRAKOW UNIVERSITY OF ECONOMICS, RAKOWICKA 27, 31-510 KRAKÓW, POLAND.

[‡] UNIV. GRENOBLE ALPES, CNRS, INRIA, GRENOBLE INP, LIG, 38000 GRENOBLE, FRANCE.

E-mail addresses: falniowf@uek.krakow.pl, panayotis.mertikopoulos@imag.fr.

in an otherwise agnostic setting where the players do not know the game being played. This has in turn generated an immense body of literature, see e.g., Hofbauer and Sigmund [30] for the biological viewpoint, Fudenberg and Levine [21], Weibull [76] and Sandholm [60] for a more economic-oriented approach, and Cesa-Bianchi and Lugosi [14] for the algorithmic / information-theoretic aspects of the theory.

Now, depending on the precise context, the question of whether the players' behavior converges to equilibrium or not may admit a wide range of answers, from positive to negative. Starting with the positive, a folk result states that if the players of a finite game follow a no-regret learning process, the players' empirical frequency of play converges in the long run to the set of coarse correlated equilibria (CCE) – also known as the game's *Hannan set* [26]. This result has been pivotal for the development of the field because no-regret play can be achieved through fairly simple myopic processes like the exponential / multiplicative weights (EW) update scheme [3, 4, 37, 75] and its many variants [52, 65, 68]. On the downside however (a) this convergence result does not concern the actual strategies employed by the players on a day-by-day basis; and (b) in many games, CCE may violate even the weakest axioms of rationalizability. For example, as was shown by Viossat and Zapechelnuk [74], it is possible for players to have *negative regret* for all time, but nonetheless play *only strictly dominated strategies* for the entire horizon of play.

This takes us to the negative end of the spectrum. If we focus on the evolution of the players' mixed strategies, a series of well-known impossibility results by Hart and Mas-Colell [27, 28] have shown that there are no uncoupled learning dynamics – deterministic or stochastic, in either continuous or discrete time – that converge to Nash equilibrium (NE) in *all* games from any initial condition.¹ In turn, this leads further weight to examining the question of equilibrium convergence within a specific class of games, and for a specific (class of) game dynamics.

In this regard, one of the most – if not *the* most – widely studied game dynamics are the *replicator dynamics* (RD) of Taylor and Jonker [70], arguably the *spiritus movens* of evolutionary game theory. Originally derived as a model for the evolution of biological populations under selective pressure in the spirit of Moran [43], the replicator dynamics were subsequently rederived in economic theory via a mechanism known as *pairwise proportional imitation* (PPI), originally due to Helbing [29],² and, at around the same time, as the mean dynamics of a stimulus-response model known as the *exponential / multiplicative weights* (EW) algorithm, cf. Auer et al. [3], Littlestone and Warmuth [37], Vovk [75] and Rustichini [56, 57].

This convergence of viewpoints is quite remarkable: even though the starting point of these considerations is a set of conceptually very different and mathematically disparate discrete-time models, they all share the replicator dynamics as a continuous-time limit. In this way, by studying the replicator dynamics, one can hope to obtain plausible predictions for the long-run behavior of the above models, at least when the time step δ of the underlying discrete-time model is sufficiently small to justify the descent to continuous time. However, since real-life modeling considerations often involve larger values of δ (e.g., as in the case of species with longer evolutionary cycles or agents with faster revision rates), we are led to the following natural question:

Do the discrete-time models underlying the replicator dynamics lead to qualitatively different outcomes? And, if so, to what extent?

¹The adjective “uncoupled” means here that a player's update rule does not explicitly depend on the other players' strategies (except implicitly, through the player's payoff function).

²See also Binmore and Samuelson [10] for a derivation via a related mechanism known as “imitation driven by dissatisfaction”, complementing the “imitation of success”; for a comprehensive account, cf. Sandholm [60].

Our contributions. One could plausibly expect that the answer to this question is most likely positive in large, complicated games, with a wide range of different behaviors emerging in the long run; on the other hand, in smaller, simpler games, the range of behaviors that arise would probably be qualitatively similar, and only differ at a quantitative level (such as the rate of convergence to an equilibrium or the like).

Somewhat surprisingly, we show that this expectation is too optimistic, even in the class of potential games [42] (which arguably possess the strongest convergence guarantees under the replicator dynamics, [48, 59–61]), and even for cases where agents are symmetric and only have two actions at their disposal (the smallest meaningful game).³ In particular, we consider the case of symmetric random matching in a 2×2 congestion game, and we show that the different discrete-time origins of the replicator dynamics exhibit the following qualitatively different behaviors:

- (1) In the biological model of intra-species competition, the dynamics converge to Nash equilibrium for any value of the time step $\delta > 0$.
- (2) In the economic model of pairwise proportional imitation, there exist certain equilibrium configurations that are globally attracting for *any* value of δ , others for which the game’s equilibrium is repelling for a range of values of δ , and yet others that lead, through the loss of equilibrium stability and period doubling, to the emergence of Li-Yorke chaos (for a different range of values of δ).
- (3) Finally, in the case of the EW algorithm, all equilibrium configurations become unstable and, unless gains from abandoning the most congested choice are equal, Li-Yorke chaos emerges whenever the time-step exceeds a certain threshold depending on the exact position of the game’s equilibrium.

In the above, the notion of *Li-Yorke chaos* – as introduced in the seminal paper of Li and Yorke [35] – means that there exists an uncountable set of initial conditions that is *scrambled*, i.e., every pair of points in this set eventually comes arbitrarily close and then drifts apart again infinitely often. In the type of systems that we consider here, Li-Yorke chaos implies other features of chaotic behavior like positive topological entropy or the existence of a set on which one can detect sensitive dependence on initial conditions à la Devaney [55]. In this sense, the system is truly unpredictable, which comes in stark contrast to the universally convergent landscape that arises in the continuous-time limit of the process (and which is only shared by the biological model above).

In this regard, the discrete-time origins of the replicator dynamics exemplify the mantra “*discretization matters*” to the extreme: The result of descending from discrete to continuous time and back is radically different, even in cases where the underlying continuous-time dynamics exhibit a universally convergent landscape that would make all asymptotic pseudo-trajectories of the process (stochastic or deterministic) converge [6]. We find this outcome particularly intriguing, as it provides a concrete, quantitative cautionary tale for the extent to which the replicator dynamics can serve as a meaningful predictor for the long-run behavior of their discrete-time origins.

Related work. There is a significant corpus of recent works suggesting that complex, non-equilibrium behaviors of boundedly rational agents (employing learning rules) seems to be common rather than exceptional. In this aspect, the seminal work of Sato et al. [63] showed analytically that even in a simple two-player zero-sum game of Rock-Paper-Scissors, the (symmetric) replicator dynamics exhibit Hamiltonian chaos. Sato and Crutchfield [62]

³Moreover, potential games are one of the few classes of games in which a variety of learning dynamics are known to converge to Nash equilibria [19, 20, 33].

subsequently extended this result to more general multiagent systems, opening the door to detecting chaos in many other games (always in the continuous-time regime).

More recently, Becker et al. [5] and Geller et al. [23] exhibited chaotic behavior for Nash maps in games like matching pennies, while Sparrow et al. [69] and van Strien and Sparrow [72] showed that fictitious play also possesses rich periodic and chaotic behavior in a class of 3x3 games, including Shapley’s game and zero-sum dynamics. In a similar vein, Piliouras and Shamma [51] showed that the replicator dynamics are Poincaré recurrent in zero-sum games, a result which was subsequently generalized to the so-called “follow-the-regularized-leader” (FTRL) dynamics [40], even in more general classes of games [38]; see also [12, 31, 39, 41] for a range of results exhibiting convergence to limit cycles and other non-trivial attractors.

It is also known that when FTRL is applied with constant step-size in zero-sum games it becomes unstable and in fact Lyapunov chaotic [15], while there is growing evidence that a class of algorithms from behavioral game theory known as experience-weighted attraction (EWA) also exhibits chaotic behavior for two-agent games with many strategies in a large parameter space [22], or in games with many agents [58]. In particular, Pangallo et al. [50] showed experimentally that EWA leads to limit cycles and high-dimensional chaos in two-agent games with negatively correlated payoffs. Chaotic behavior has been also detected for anti-coordination games for discretizations similar in vein to one pairwise proportional imitation model discussed in this article. Vilone et al. [73] showed numerically Lyapunov chaos for some values of parameters in Snowdrift game, while other works detected complex behavior in Battle of sexes and Leader game [44, 45, 49]. All in all, careful examination suggests a complex behavioral landscape in many games (small or large) for which no single theoretical framework currently applies.

However, none of the above results implies chaos in the formal sense of Li-Yorke. The first formal proof of Li-Yorke chaos was shown for the EW algorithm in a single instance of two-agent two-strategy congestion game by Palaiopanos et al. [47]. This result was generalized and strengthened (in the sense of positive topological entropy) for all two-agent two-strategy congestion games [17]. In arguably the main precursor of our work [16] topological chaos in nonatomic congestion game where agents use EW was established. This result was then extended to FTRL with steep regularizers [8] and EWA algorithms [9], but the resulting framework does not apply to the range of models from biology and economic theory considered here (species competition and revision protocols respectively).

Finally, the question regarding different qualitative behaviors for discrete-time dynamics featuring the same continuous-time limit is of substantial interest in other fields like theoretical ecology as well. In particular, three discrete-time models are often used to describe analogues of logistic growth or the logistic ODE in discrete-time ecological models: the Beverton-Holt model [7] (which is actually a special case Model I in the case of frequency-independent payoffs for each strategy), the logistic map, and the Ricker model [53]. Substantial work has been done to show that while the Beverton-Holt model always features stable equilibrium behavior, both the logistic map and the Ricker model can produce periodic and chaotic behavior for different time-steps and ecological parameter values, even though all three models are designed to represent an analogue of the same continuous-time model (see e.g. [1, 13, 32, 34, 64, 71]).

Notation. In what follows, we will write $\mathbf{v} \cdot \mathbf{w} = \sum_{i=1}^m v_i w_i$ for the (Euclidean) inner product between two real m -dimensional vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^m$, and $\mathbf{v} \odot \mathbf{w} = (v_1 w_1, \dots, v_m w_m)$ for their Hadamard product. Finally, to simplify notation, given a function $f: \mathbb{R} \rightarrow \mathbb{R}$, we will thread it over all elements of $\mathbf{v} \in \mathbb{R}^m$ by writing $f(\mathbf{v}) := (f(v_1), \dots, f(v_m))$.

2. PRELIMINARIES

Throughout our paper, we will focus on games with a continuum of nonatomic players modeled by the unit interval $\mathcal{I} = [0, 1]$, with each player choosing (in a measurable way) an *action* – or *pure strategy* – from a finite set indexed by $i \in \mathcal{A} \equiv \{1, \dots, m\}$. Letting $x_i \in [0, 1]$ denote the mass of agents playing $i \in \mathcal{A}$, the overall distribution of actions at any point in time will be specified by the *state of the population* $\mathbf{x} = (x_1, \dots, x_m)$, viewed here as a point in the unit simplex $\mathcal{X} := \Delta(\mathcal{A}) = \{\mathbf{x} \in \mathbb{R}_+^m : \sum_{i \in \mathcal{A}} x_i = 1\}$ of \mathbb{R}^m . The players' payoffs in this setting are described by an ensemble of payoff functions $u_i : \mathcal{X} \rightarrow \mathbb{R}_+$ (assumed Lipschitz), with $u_i(\mathbf{x})$ denoting the payoff to agents playing $i \in \mathcal{A}$ when the population is at state \mathbf{x} . Putting everything together, we will write $u(\mathbf{x}) = \sum_{i \in \mathcal{A}} x_i u_i(\mathbf{x})$ for the population's *mean payoff* at state $\mathbf{x} \in \mathcal{X}$, $v(\mathbf{x}) = (u_1(\mathbf{x}), \dots, u_m(\mathbf{x}))$ for the associated *payoff vector* at state \mathbf{x} , and we will refer to the tuple $\mathcal{G} \equiv \mathcal{G}(\mathcal{A}, v)$ as a *population game*.⁴ Finally, a state $\mathbf{x} \in \mathcal{X}$ is a Nash equilibrium of the game \mathcal{G} if every strategy in use earns a maximal payoff (equivalently, each agent in population chooses an optimal strategy with respect to the choices of others).

In the general context of population games, the most widely studied evolutionary game dynamics are the *replicator dynamics* (RD) of Taylor and Jonker [70]. These are succinctly described by the continuous-time system

$$\dot{x}_i = x_i[u_i(\mathbf{x}) - u(\mathbf{x})] \quad (\text{RD})$$

which specifies that the per capita growth rate of the population share of a given strategy $i \in \mathcal{A}$ is proportional to the difference between the payoff $u_i(\mathbf{x})$ of said strategy and the mean population payoff $u(\mathbf{x}) = \mathbf{x} \cdot v(\mathbf{x}) = \sum_{i=1}^m x_i u_i(\mathbf{x})$. For an introduction to the vast literature surrounding the replicator dynamics, see [30, 60, 76] and references therein.

A specific class of population games – and, indeed, the class that will be of most interest to us – is obtained when two individuals are selected randomly from the population and are matched to play a symmetric two-player game with payoff matrix $M = (M_{ij})_{i,j \in \mathcal{A}}$. In this case, the payoff to agents playing $i \in \mathcal{A}$ at state \mathbf{x} is $u_i(\mathbf{x}) = \sum_{j \in \mathcal{A}} M_{ij} x_j$, so the game's payoff field can be written in concise form as $v(\mathbf{x}) = M\mathbf{x}$. Following standard conventions in the field, we will refer to this scenario as *symmetric random matching* [25, 30, 60, 76].

3. DYNAMICS

In this section, we discuss and derive three established models for the evolution of large populations in discrete time. All three models share the same continuous-time limit, namely the replicator equation (RD); however, as we show in Section 4, the behavior of each model is drastically different in discrete time, even in the simplest of games.

Most of the material presented in this section is not new, but we chose to include all relevant details for self-completeness and uniformity of notation.

Model I: Intra-species competition. We begin with the biological microfoundations of an evolutionary model for intra-species competition in the spirit of Moran [43].⁵ Here, each pure strategy $i \in \mathcal{A} = \{1, \dots, m\}$ represents a genotype in a population that reproduces asexually and $u_i(\mathbf{x})$ represents the reproductive fitness of the i -th genotype when the population is at state $\mathbf{x} \in \mathcal{X}$ – i.e., the net number of offspring per capita in the unit of time. Then, if

⁴We focus here on games played by a single population of nonatomic agents. The extension of our analysis to multi-population settings requires more elaborate notation, but is otherwise straightforward.

⁵The case of inter-species competition is similar, but the notation is heavier so we do not treat it here.

$z_i(t)$ denotes the *absolute* size of the i -th genotype at time t and δ is the interval between generations, the evolution of the population will be governed by the discrete-time dynamics

$$z_i(t + \delta) = z_i(t) + z_i(t)u_i(\mathbf{x}(t))\delta \quad \text{with} \quad x_i(t) = \frac{z_i(t)}{\sum_{j \in \mathcal{A}} z_j(t)}. \quad (1)$$

Accordingly, expressing everything in terms of population states – that is, as a function of the relative frequency $x_i(t)$ of each genotype – we obtain the autonomous dynamics

$$\begin{aligned} x_i(t + \delta) &= \frac{z_i(t + \delta)}{\sum_{j \in \mathcal{A}} z_j(t + \delta)} \\ &= \frac{z_i(t) \cdot [1 + u_i(\mathbf{x}(t))\delta]}{\sum_{j \in \mathcal{A}} z_j(t) \cdot [1 + u_j(\mathbf{x}(t))\delta]} \\ &= \frac{x_i(t) \cdot [1 + u_i(\mathbf{x}(t))\delta]}{1 + u(\mathbf{x}(t))\delta} \end{aligned} \quad (1)$$

where, in the last step, we have used the fact that $\sum_{j \in \mathcal{A}} x_j(t) = 1$.

Formally, **Model I** is mathematically equivalent to the so-called “linear multiplicative weights” algorithm in computer science and learning theory, cf. Arora et al. [2], Littlestone and Warmuth [37], Vovk [75] and references therein. In addition, for small δ , a simple first-order Taylor expansion yields

$$\begin{aligned} x_i(t + \delta) - x_i(t) &= x_i(t) \cdot [1 + u_i(\mathbf{x}(t))\delta] \cdot [1 - \delta u(\mathbf{x}(t)) + \mathcal{O}(\delta^2)] - x_i(t) \\ &= \delta x_i(t)[u_i(\mathbf{x}(t)) - u(\mathbf{x}(t))] + \mathcal{O}(\delta^2) \end{aligned} \quad (2)$$

so, in the continuous-time limit $\delta \rightarrow 0$, we get

$$\dot{x}_i(t) \sim \frac{x_i(t + \delta) - x_i(t)}{\delta} = x_i(t)[u_i(\mathbf{x}(t)) - u(\mathbf{x}(t))] + \mathcal{O}(\delta). \quad (3)$$

In the above, the asymptotic equality sign “ \sim ” is to be interpreted loosely and is only meant to suggest that **Model I** represents an Euler discretization of (RD) up to a higher-order $\mathcal{O}(\delta^2)$ correction term. Because this term is negligible in the continuous-time limit $\delta \rightarrow 0$, (RD) is commonly regarded in the literature as the mean dynamics of **Model I** [30, 60].

Model II: Pairwise proportional imitation. The second model that we consider has its roots in the mass-action interpretation of game theory and, more precisely, the theory of revision protocols [60, 76]. Referring to the classical textbook of Sandholm [60] for the details, suppose that each agent occasionally receives an opportunity to switch actions – say, based on the rings of a Poisson alarm clock – and, at such moments, they reconsider their choice of action by comparing its payoff to that of a randomly chosen individual in the population. A *revision protocol* of this kind is typically defined by specifying the *conditional switch rate* $\rho_{ij}(\mathbf{x})$ at which a revising i -strategist switches to strategy j when the population is at state $\mathbf{x} \in \mathcal{X}$. In this case, the population share of agents switching from strategy i to strategy j over a small interval of time δ will be $x_i \rho_{ij} \delta$, leading to the inflow-outflow equation

$$x_i(t + \delta) = x_i(t) + \delta \left[\sum_{j \neq i} x_j(t) \rho_{ji}(\mathbf{x}(t)) - x_i(t) \sum_{j \neq i} \rho_{ij}(\mathbf{x}(t)) \right] \quad (4)$$

One of the most widely studied revision protocols of this type is the *pairwise proportional imitation* (PPI) of Helbing [29], as described by the switch rate functions

$$\rho_{ij}(\mathbf{x}) = x_j [u_j(\mathbf{x}) - u_i(\mathbf{x})]_+. \quad (\text{PPI})$$

Under this protocol, a revising agent first observes the action of a randomly selected opponent, so a j -strategist is observed with probability x_j when the population is at state $\mathbf{x} \in \mathcal{X}$. Then,

if the payoff of the incumbent strategy $i \in \mathcal{A}$ is lower than that of the benchmark strategy j , the agent imitates the selected agent with probability proportional to the payoff difference $[u_j(\mathbf{x}) - u_i(\mathbf{x})]_+$; otherwise, the revising agent skips the revision opportunity and sticks to their current action.

Now, substituting the protocol (PPI) into (4), a straightforward calculation yields the autonomous, discrete-time dynamics⁶

$$\begin{aligned} x_i(t + \delta) &= x_i(t) + \delta x_i(t) \left[\sum_{j \neq i} x_j(t) [u_i(\mathbf{x}(t)) - u_j(\mathbf{x}(t))]_+ - \sum_{j \neq i} x_j(t) [u_j(\mathbf{x}(t)) - u_i(\mathbf{x}(t))]_+ \right] \\ &= x_i(t) + \delta x_i(t) \left[\sum_{j \neq i} x_j(t) [u_i(\mathbf{x}(t)) - u_j(\mathbf{x}(t))] \right] \\ &= x_i(t) + \delta x_i(t) [u_i(\mathbf{x}(t)) - u(\mathbf{x}(t))] \end{aligned} \quad (\text{II})$$

where, in the second-to-last line, we used the fact that $\sum_{j \neq i} x_j(t) = 1 - x_i(t)$. Just like [Model I](#), [Model II](#) can be seen as an Euler discretization of (RD); however, in contrast to its biological counterpart, there is *no* $\mathcal{O}(\delta^2)$ correction term in [Model II](#). Albeit negligible in the limit $\delta \rightarrow 0$, we will see in [Section 4](#) that the residual $\mathcal{O}(\delta^2)$ term that appears in [Model I](#) plays a major role in the long-run behavior of the dynamics.

Model III: Learning with exponential weights. The last model we consider has its origins in learning theory and, more specifically, the so-called multi-armed bandit problem, cf. Auer et al. [3, 4], Cesa-Bianchi and Lugosi [14] and references therein. Following Hadikhanloo et al. [25], this model can be described in our population setting as follows: at time t , the performance of each strategy $i \in \mathcal{A}$ is scored by measuring its cumulative payoff over an interval of time δ ; subsequently, at time $t + \delta$, the population is redistributed with each strategy receiving a population share that is exponentially proportional to its cumulative payoff up to time $t + \delta$ (i.e., agents select with exponentially higher probability those strategies that perform better over time).

Formally, this simple stimulus-response model amounts to the *exponential/multiplicative weights* update

$$y_i(t + \delta) = y_i(t) + \delta u_i(\mathbf{x}(t)) \quad \text{with} \quad x_i(t) = \frac{\exp(y_i(t))}{\sum_{j \in \mathcal{A}} \exp(y_j(t))} \quad (\text{EW})$$

which, in the context of learning theory, forms the basis of the *exponential-weights algorithm for exploration and exploitation* (EXP3) [3, 4].⁷ Thus, under (EW), the associated population shares will be governed by the autonomous dynamics

$$\begin{aligned} x_i(t + \delta) &= \frac{\exp(y_i(t + \delta))}{\sum_{j \in \mathcal{A}} \exp(y_j(t + \delta))} \\ &= \frac{\exp(y_i(t)) \exp(\delta u_i(\mathbf{x}(t)))}{\sum_{j \in \mathcal{A}} \exp(y_j(t)) \exp(\delta u_j(\mathbf{x}(t)))} \\ &= \frac{x_i(t) \exp(\delta u_i(\mathbf{x}(t)))}{\sum_{j \in \mathcal{A}} x_j(t) \exp(\delta u_j(\mathbf{x}(t)))} \end{aligned} \quad (\text{III})$$

⁶Instead of [Model I](#) where δ can be arbitrarily large, in [Model II](#) the step size is bounded as otherwise its iterations may fail to lie in the simplex.

⁷This particular instantiation of the algorithm is known as Hedge [3]; in some other threads of the literature, (EW) is referred to as the *multiplicative weights update* (MWU) [2]. We employ the original terminology of [3, 4] to distinguish it from the linearized version of [37].

where, in the last line, we used the fact that $x_i(t) \propto \exp(y_i(t))$, as per (EW). A first-order Taylor expansion then yields

$$\begin{aligned}
x_i(t + \delta) - x_i(t) &= x_i(t) \left[\frac{\exp(\delta u_i(\mathbf{x}(t)))}{\sum_{j \in \mathcal{A}} \mathbf{x}_j(t) \exp(\delta u_j(\mathbf{x}(t)))} - 1 \right] \\
&= x_i(t) \left[\frac{1 + \delta u_i(\mathbf{x}(t)) + \mathcal{O}(\delta^2)}{1 + \delta u(\mathbf{x}(t)) + \mathcal{O}(\delta^2)} - 1 \right] \\
&= x_i(t) \left[(1 + \delta u_i(\mathbf{x}(t)))(1 - \delta u(\mathbf{x}(t))) + \mathcal{O}(\delta^2) - 1 \right] \\
&= \delta x_i(t) [u_i(\mathbf{x}(t)) - u(\mathbf{x}(t))] + \mathcal{O}(\delta^2)
\end{aligned} \tag{5}$$

so [Model III](#) can also be seen as an Euler discretization of (RD) up to an $\mathcal{O}(\delta^2)$ correction term. Conceptually, this is similar to [Model I](#), though the two corrections are, in general, different; we will see in [Section 4](#) that this difference plays a major role in the qualitative behavior of dynamics when δ is not infinitesimally small.

4. ANALYSIS AND RESULTS

As we saw, [Models I–III](#) admit the same continuous-time limit – the replicator dynamics (RD) – so it would be natural to expect that they behave similarly in the long run, especially when the time unit δ is infinitesimally small. However, when real-life modeling considerations call for larger values of δ (e.g., as in the case of species with longer evolutionary cycles or agents with faster revision rates), it is not clear if this heuristic holds true (and to what extent), even when (RD) would suggest a unique long-run outcome. In view of this, our goal in the sequel will be to examine in detail the asymptotic behavior of [Models I–III](#), and to see whether any qualitative differences arise between these models and/or the underlying dynamics (RD).

For concreteness, we will focus on the two extremes of the spectrum of possible asymptotic behaviors: (global) convergence on the one hand, and (deterministic) chaos on the other. The reason for this is straightforward: if the dynamics are globally convergent, the population’s initial state is eventually forgotten, and all initializations ultimately settle down to the same state; instead, if the dynamics are chaotic, even arbitrarily small differences in the population’s initial state would lead to drastically different behavior. As such, convergence and chaos can be seen as antithetical to each other – and hence, as opposites in terms of long-run predictions.

Now, to put all this on a solid footing, it will be convenient to recast [Models I–III](#) in abstract recursive form as

$$\mathbf{x}((n+1)\delta) = f(\mathbf{x}(n\delta)) \tag{6}$$

where $\mathbf{x}(n\delta)$ denotes the state of the population at time $t = n\delta$, $n = 1, 2, \dots$, and the dynamics’ *update map* $f: \mathcal{X} \rightarrow \mathcal{X}$ is defined as follows:

- Under [Model I](#):

$$f(\mathbf{x}) \equiv \frac{\mathbf{x} + \delta v(\mathbf{x}) \odot \mathbf{x}}{1 + \delta u(\mathbf{x})} \tag{7.I}$$

- Under [Model II](#):

$$f(\mathbf{x}) \equiv \mathbf{x} + \delta[v(\mathbf{x}) \odot \mathbf{x} - u(\mathbf{x}) \odot \mathbf{x}] \tag{7.II}$$

- Under [Model III](#):

$$f(\mathbf{x}) \equiv \frac{\mathbf{x} \odot \exp(\delta v(\mathbf{x}))}{\mathbf{x} \cdot \exp(\delta v(\mathbf{x}))} \tag{7.III}$$

Definitions. Let (X, d) be a nonempty metric space and let $f: X \rightarrow X$ be a continuous map (that is (X, f) is a discrete dynamical system). In this setting we call the fixed point x of the map f .

- *attracting*, if there is an open neighborhood $U \subset X$ of x such that for every $y \in U$ we have $\lim_{n \rightarrow \infty} f^n(y) = x$, where f^n is a composition of the map f with itself n -times.
- *repelling*, if there is an open neighborhood $U \subset X$ of x such that for every $y \in U$, $y \neq x$ there exists $n \in \mathbb{N}$ such that $f^n(x) \in X \setminus U$.

As we will study differentiable maps on the unit interval, the fixed point $x \in I = [0, 1]$ is attracting if $|f'(x)| < 1$, and x is repelling when if $|f'(x)| > 1$. If $|f'(x)| = 1$ we need more information. An orbit $\{f^n(x)\}$ is called *periodic* of period T if $f^{n+T}(x) = f^n(x)$ for any $n \in \mathbb{N}$. The smallest such T is called the period of x . The periodic orbit is called attracting, if x is an attracting fixed point of (X, f^T) , and repelling, if x is a repelling fixed point of (X, f^T) .

On the antipodes of convergence to the fixed point lays chaotic behavior. The most widely used definition of chaos is due to Li and Yorke [35]: A pair of points $x, x' \in X$ is said to be *scrambled* – or a *Li–Yorke pair* – if

$$\liminf_{n \rightarrow \infty} \text{dist}(f^n(x), f^n(x')) = 0 \quad \text{and} \quad \limsup_{n \rightarrow \infty} \text{dist}(f^n(x), f^n(x')) > 0. \quad (8)$$

We then say that (6) is *chaotic* (in the Li–Yorke sense) if it admits an uncountable *scrambled set*, i.e., a set $\mathcal{S} \subseteq X$ such that every pair of distinct points $x, x' \in \mathcal{S}$ is scrambled.

Remark 1. The origins of the definition of Li-Yorke chaos can be traced back to the seminal paper of Li and Yorke [35], where the notion of a scrambled set was introduced as a surrogate for mixing: intuitively, any two (distinct) orbits starting in a scrambled set will come arbitrarily close to each other and subsequently spring aside infinitely many times. In one-dimensional systems, which we study here, Li-Yorke chaos implies other features of chaotic behavior like positive topological entropy or existence of the set on which one can detect sensitive dependence on initial conditions in the sense of Devaney [55].

Game. In general, showing that a system exhibits chaotic behavior is a task of considerable difficulty, and a satisfactory theory exists only for low-dimensional systems. On that account, we will focus on a random matching scenario induced by the 2×2 symmetric game with actions $\mathcal{A} = \{A, B\}$ and payoff bimatrix

$$\begin{array}{c|cc} & A & B \\ \hline A & (0, 0) & (\gamma_A, \gamma_B) \\ B & (\gamma_B, \gamma_A) & (0, 0) \end{array} \quad (9)$$

This is a normalized anti-coordination / congestion game where the parameters $\gamma_A, \gamma_B > 0$ reflect the benefit of deviating from the most congested choice: if the entire population plays A , an agent would gain γ_B by deviating to B ; and, likewise, if the entire population plays B , an agent would gain γ_A by deviating to A .⁸

For this game, the replicator dynamics (RD) boil down to the one-dimensional system

$$\dot{x} = x(1-x)[\gamma_A(1-x) - \gamma_B x] \quad (10)$$

⁸Congestion games, introduced by Rosenthal [54] are isomorphic to potential games [42].

where, in a slight abuse of notation, we write $x \equiv x_A$ for the population share of A -strategists. This system admits three fixed points, 0, 1, and $p = \gamma_A/(\gamma_A + \gamma_B)$, with the following stability properties:⁹

- *For the fixed points at 0 and 1:* letting $RD(x) := x(1-x)[\gamma_A(1-x) - \gamma_Bx]$ denote the RHS of (10), we trivially get $RD'(0) = \gamma_A > 0$ and $RD'(1) = \gamma_B > 0$ so, by standard results in dynamical systems theory [67], they are both linearly unstable under (10).
- *For the fixed point at p :* working as above, we get $RD'(p) = -\gamma_A\gamma_B/(\gamma_A + \gamma_B) < 0$, so p is linearly stable under (10).

In fact, it is easy to see that p is the unique symmetric Nash equilibrium of (9) and, in fact, it is a global evolutionarily stable state (ESS) thereof; as a result, the dynamics (10) converge to p from every interior initialization $x(0) \in (0, 1)$.

Given this robust, global convergence landscape and the existence of a (global) evolutionarily stable state, one might expect that Models I–III enjoy similar convergence properties. However, as we show below, the long-run behavior of Models I–III can be drastically different, ranging from fully convergent to fully chaotic.

Theorem 1. *With notation as above, the dynamics (6) with f given by Eqs. (7.I)–(7.III) exhibit for the game (9) the following asymptotic behavior:*

- *Under Model I:* trajectories of all points of $(0, 1)$ converge to Nash equilibrium p for any $\delta > 0$.
- *Under Model II, which is well defined for $\delta \leq \delta^* = \min\{4/p^2, 4/(1-p)^2\}$:*
 - (1) Trajectories of all points of $(0, 1)$ converge to Nash equilibrium p for $\delta \leq 2/\gamma_A\gamma_B$.
 - (2) If $p \in (0, 1/3) \cup (2/3, 1)$, then trajectories of all points of $(0, 1)$ converge to Nash equilibrium p for any δ .
 - (3) If $p \in (1/3, 2/3)$, then for $\delta \in (2/\gamma_A\gamma_B, \delta^*)$ the (unique) Nash equilibrium p is repelling and so, except for a countable set of initial conditions, trajectories of all points from $(0, 1)$ do not converge to equilibrium.
 - (4) If $p \in (\frac{1}{23}(31 - 12\sqrt{3}), \frac{1}{23}(12\sqrt{3} - 8))$, then there exists a unique δ_p^{II} such that f has periodic orbits of all periods and is Li-Yorke chaotic for any $\delta \in (\delta_p^{II}, \delta^*)$.
- *Under Model III:*
 - (1) Trajectories of all points of $(0, 1)$ converge to Nash equilibrium p for $\delta \leq 2/\gamma_A\gamma_B$.
 - (2) If $\gamma_A \neq \gamma_B$, then there exists δ_p^{III} such that if $\delta > \delta_p^{III}$ then f has periodic orbits of all periods and is Li-Yorke chaotic.
 - (3) If $\gamma_A = \gamma_B$ and $\delta > 2/\gamma_A^2$, then f has a periodic attracting orbit $\{\sigma_\delta, 1 - \sigma_\delta\}$, where $0 < \sigma_\delta < 1/2$. This orbit attracts trajectories of all points of $(0, 1)$, except countably many points whose trajectories eventually fall into the repelling fixed point at $1/2$.

Theorem 1 shows that the dynamics given by the models we analyze, described by (RD) in the continuous case, exhibit qualitatively different behaviors in the discrete case. This points to a distinct nature of micro foundations of applications, and interpretations, of evolutionary game theory in different scientific contexts. First, it puts its biological foundations and foundations in economics and computer science at odds. In the biological model we see convergence to the unique Nash equilibrium regardless of the length of the evolutionary cycle. On the contrary, for Model II and Model III convergence to Nash equilibrium depends

⁹Recall here that a point is *stable* if all trajectories that start sufficiently close to it remain close for all time; otherwise, the point is called *unstable*. Moreover, if all nearby orbits converge to the point in question, it is called *attracting*; and if a point is both stable and attracting, it is called *asymptotically stable*.

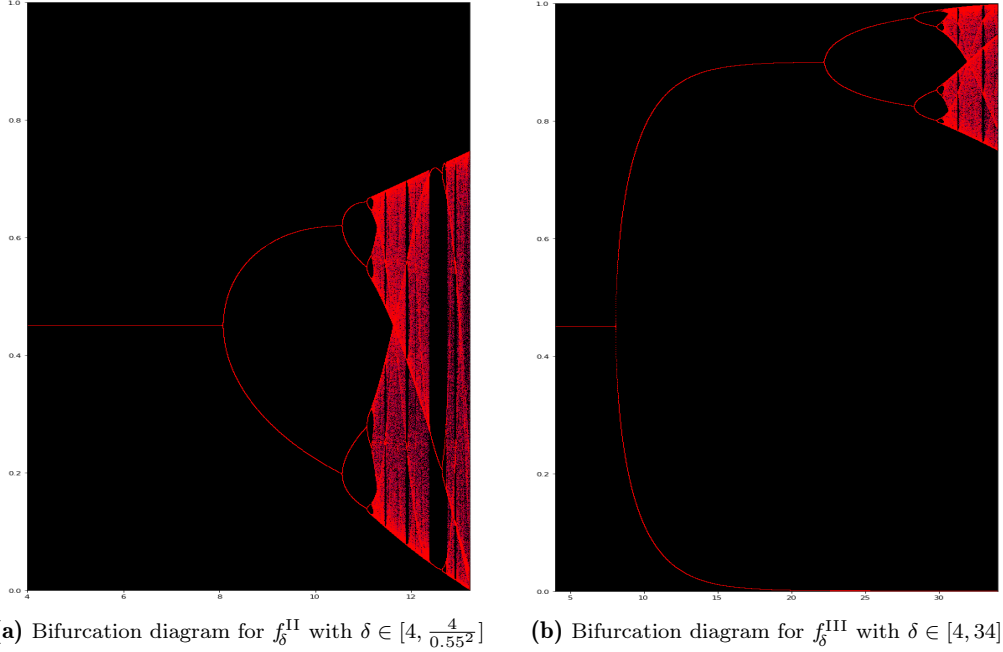


Figure 1: Bifurcation diagrams for f_δ^{II} and f_δ^{III} with the equilibrium $p = 0.45$. At horizontal axes we have values of δ , while on vertical axes ratio of A -strategists in the population. The first bifurcation on both diagrams is at $\delta_0 = 2/0.45 \cdot 0.55 = \frac{22}{9}$. But then the second bifurcation for f_δ^{II} is much faster (around $\delta = 10.5$ while for f_δ^{III} it is around $\delta = 25$). It is worth pointing out that period 3 can be detected for f_δ^{II} at $\delta \approx 12.5$, while for f_δ^{III} it is $\delta \approx 55.5$.

on the revision rate and learning rate, respectively. The convergence result fails for these models after crossing the value $\delta_0 = 2/\gamma_A\gamma_B$, when both systems become unpredictable. Nevertheless, [Model II](#) and [Model III](#) differ in other aspects. For [Model II](#) once the Nash equilibrium loses stability (which happens only if $\gamma_A + \gamma_B < 3\gamma_A < 2(\gamma_A + \gamma_B)$), the system is going through period doubling bifurcation relatively fast, and for values of γ_A and γ_B close to each other (p close to $1/2$), leads to periodic orbits of any period and chaotic behavior (see [Figures 1, 2](#)). On the other hand, increasing learning rate in [Model III](#) will inevitably lead to chaotic behavior only if $\gamma_a \neq \gamma_B$, but the speed of period doubling road to chaos depends on the asymmetry of gains γ_A and γ_B . In particular, for values of γ_A, γ_B for which in [Model II](#) we detect chaos, in [Model III](#) the period doubling road to chaos can be (extremely) slow (compare bifurcation diagrams in [Figure 1](#)).

4.1. Skeleton of the proof of [Theorem 1](#). For the game [\(9\)](#) we can concentrate on the population share of A -strategists. Then

- for [Model I](#) we get the map $f_\delta^{\text{I}}: I \mapsto I$, of the unit interval $I = [0, 1]$, given by

$$f_\delta^{\text{I}}(x) = x \frac{1 + (\gamma_A + \gamma_B)\delta p(1-x)}{1 + (\gamma_A + \gamma_B)\delta x(1-x)}. \quad (11)$$

- for [Model II](#) we are interested in the dynamics of the map $f_\delta^{\text{II}}: I \mapsto I$ given by

$$f_\delta^{\text{II}}(x) = x(1 - (\gamma_A + \gamma_B)\delta(1-x)(x-p)). \quad (12)$$

- for **Model III** we get $f_\delta^{\text{III}}: I \mapsto I$ given by

$$f_\delta^{\text{III}}(x) = \frac{x}{x + (1-x)\exp[(\gamma_A + \gamma_B)\delta(x-p)]}. \quad (13)$$

Thus, in the proof of Theorem 1 we will focus on dynamics of these maps and how it depends on the choice of the step size δ . Without loss of generality we can assume that $\gamma_A + \gamma_B = 1$, as otherwise we can proceed for the step size $\delta' = (\gamma_A + \gamma_B)\delta$. We begin in Section 4.2 with the discussion on the conditions for global convergence for these maps. To this aim we introduce Schwarzian derivative [of an interval map](#)

$$Sf \equiv \frac{f'''}{f'} - \frac{3}{2} \left(\frac{f''}{f'} \right)^2,$$

as negative Schwarzian derivative guarantees *good* behavior of the interval map. We will show that under conditions, which all studied maps fulfill, existence of an interior attracting fixed point will imply global convergence of dynamics to this point (see Proposition 1). In Section 4.3 we use Proposition 1 to show global convergence of game dynamics introduced by the map (11) for any δ . Dynamics of f_δ^{II} is studied in Section 4.4. By similar argument as for **Model I** we show that as long as p is attracting it attracts all points. Nevertheless, once the step size crosses the value of $2/p(1-p)$ the system becomes unstable and for a range of values of p close to $1/2$ will eventually (for large δ) be chaotic. To show that we carefully choose a point x_0 , which fulfills assumptions of Li-Misiurewicz-Panigiani-Yorke theorem [36]. Finally, in Section 4.5 we discuss dynamics of (13). Once more Proposition 1 guarantees that p attracts all points from $(0, 1)$ as long as the step size is smaller than $2/p(1-p)$. The map f_δ^{III} is already known in the literature, see [16, 17, 47], and the complete proof of the rest of the theorem comes from Theorems 3.10 and 3.11 from [17]. Here, for the completeness of the presentation, we sketch the proof of Li-Yorke chaos for $p \neq 1/2$.

4.2. Auxiliary result. Let

$$\mathcal{F} = \{f: I \mapsto I, \text{Fix}f = \{0, p, 1\}, 0, 1 \text{ repelling fixed points}, p \in (0, 1)\},$$

where $\text{Fix}f$ denotes the set of fixed points of f .

Lemma 1. *Let $f \in \mathcal{F}$. If the trajectories of all points $0 < x < p$ are attracted to p , then the trajectories of all points from $(0, 1)$ are attracted to p . Similarly, if the trajectories of all points $1 > x > p$ are attracted to p , then the trajectories of all points from $(0, 1)$ are attracted to p .*

Proof. Assume that there is a point in $(0, 1)$, whose trajectory is not attracted to p . Since both 0 and 1 are repelling, by [66], f has a periodic orbit of period 2. If the trajectories of all points $x < p$ (respectively, $x > p$) are attracted to p , this periodic orbit has to lie entirely to the right (respectively, left) of p . Thus, there is a fixed point to the right (respectively, left) of p , a contradiction. \square

Proposition 1. *Let $f \in \mathcal{F}$. Let f fulfill one of the following conditions*

- (1) f is increasing
- (2) f is bimodal and its Schwarzian derivative is negative.

If p is attracting, then it is globally attracting.

Proof. If f is strictly increasing, then it does not have a periodic orbit of period 2, so p is globally attracting.

Assume that f is bimodal. If p belongs to the left or right lap, then, by Lemma 1, p is globally attracting. Assume that p belongs to the interior of the middle lap. Because the Schwarzian derivative of f is negative, then, by Singer's theorem [18], the interval joining

p with one of the critical points of f is in the basin of attraction A of p . We may assume that this critical point is the left one, c_l . There is a unique point $y < c_l$ such that $f(y) = p$. Then $f([y, c_l]) = f([c_l, p]) \subset A$, so $[y, p] \subset A$. For every point $x < y$ we have $x < f(x) < p$. Therefore, the trajectory of x increases as long as it stays to the left of y . Since there are no fixed points to the left of y , the trajectory has to enter $[y, p]$ sooner or later. This proves that $(0, p] \subset A$, so by Lemma 1, p is globally attracting. \square

Remark 2. If the differentiable map $f \in \mathcal{F}$ fulfill one of the assumptions of Proposition 1, then p is attracting if and only if $|f'(p)| < 1$.

4.3. Game dynamics under Model I. We study the dynamics introduced by (11). First, f_δ^I is well-defined for every $\delta > 0$, that is $0 \leq f_\delta^I(x) \leq 1$ for every $x \in [0, 1]$. Obviously values $f_\delta^I(x)$ are always nonnegative and the condition $f_\delta^I(x) \leq 1$ is equivalent to $1 + \delta x(1 - p) \geq 0$, which is always satisfied. Thus, $f_\delta^I: [0, 1] \mapsto [0, 1]$ for any $\delta > 0$.

To study dynamics introduced by f_δ^I we look at fixed points and the derivative of f_δ^I . So, $f_\delta^I(x) = x$ if and only if $x = 0$ or $\delta p(1 - x) = \delta x(1 - x)$. Thus, our map has three fixed points: 0, 1 and p . Now, we look at the derivative of f_δ^I , which is equal to

$$(f_\delta^I)'(x) = \frac{\delta x^2 - 2p\delta x + \delta p + 1}{(1 + \delta x(1 - x))^2}. \quad (14)$$

We check stability of fixed points:

$$(f_\delta^I)'(0) = 1 + \delta\gamma_A > 0, \quad (f_\delta^I)'(1) = 1 + \delta\gamma_B > 0$$

for any $\delta > 0$. So, both 0 and 1 are repelling. Thus, $f \in \mathcal{F}$. Moreover,

$$f_\delta^I'(p) = \frac{1}{1 + \delta\gamma_A\gamma_B},$$

so $|(f_\delta^I)'(p)| < 1$ and p is attracting for any value of $\delta > 0$. Finally, we look at the monotonicity of f_δ^I . The sign of the derivative of f_δ^I depends only on the sign of

$$Q(x) = \delta x^2 - 2p\delta x + \delta p + 1.$$

Nevertheless, [the discriminant](#)

$$\Delta = 4\delta(\delta p(p - 1) - 1)$$

is always negative as $p \in (0, 1)$ and $\delta > 0$. So, f_δ^I has no extrema and, as $(f_\delta^I)'(0) > 0$, f_δ^I is an increasing map for any δ .

Therefore, by Proposition 1 we obtain the part of Theorem 1 on game dynamics under [Model I](#).¹⁰

4.4. Game dynamics under Model II. Game dynamics for the map f_δ^{II} is well-defined only when we cannot leave the simplex. Thus, we have to assume that

$$\delta \leq \min_{x \in (p, 1)} \frac{1}{(1 - x)(x - p)} = \frac{4}{(1 - p)^2} \quad \text{and} \quad \delta \leq \min_{x \in (0, p)} \frac{1}{x(p - x)} = \frac{4}{p^2}.$$

Thus, f_δ^{II} is well-defined when

$$\delta \leq \delta^* = \min \left\{ \frac{4}{p^2}, \frac{4}{(1 - p)^2} \right\}.$$

¹⁰Thus, if only the initial state of the population is polymorphic, then the system converges to the evolutionary stable state p . On the other hand, as 0 and 1 are repelling, monomorphic populations are sensitive to small perturbations.

Fixed points of f_δ^{II} are 0 and solutions of the equation

$$\delta(1-x)(x-p) = 0.$$

So the map f_δ^{II} has three fixed points: 0, 1 and $p \in (0, 1)$.

As

$$(f_\delta^{\text{II}})'(x) = 3\delta x^2 - 2\delta(1+p)x + p\delta + 1 \quad (15)$$

we see that $(f_\delta^{\text{II}})'(0) = 1 + p\delta$ and $(f_\delta^{\text{II}})'(1) = 1 + (1-p)\delta$ so both 0 and 1 are repelling for every δ . Thus, $f \in \mathcal{F}$. As well as for Model I, crucial in our considerations will be stability of the interior fixed point p . This point will be attracting as long as $|(f_\delta^{\text{II}})'(p)| < 1$. We have

$$(f_\delta^{\text{II}})'(p) = 1 - \delta p(1-p),$$

so p is attracting as long as $\delta < \frac{2}{p(1-p)}$ and repelling otherwise.

From (15) our map is increasing as long as $\delta < \frac{3}{1-p+p^2}$. Then, for $\delta > \frac{3}{1-p+p^2}$ the map f_δ^{II} is bimodal with critical points

$$c_l = \frac{1+p}{3} - \frac{1}{3} \sqrt{\frac{\delta(1+p)^2 - 3p\delta - 3}{\delta}}, \quad \text{and} \quad c_r = \frac{1+p}{3} + \frac{1}{3} \sqrt{\frac{\delta(1+p)^2 - 3p\delta - 3}{\delta}}.$$

For interval maps with *good* behavior are those with negative Schwarzian derivative. We should not expect that $Sf_\delta^{\text{II}} < 0$ for $\delta < \frac{3}{1-p+p^2}$ as then f_δ^{II} is a homeomorphism. For $\delta > \frac{3}{1-p+p^2}$ we have the following fact.

Lemma 2. *If $\delta > \frac{3}{1-p+p^2}$, then $Sf_\delta^{\text{II}} < 0$.*

Proof. Elementary calculations give us

$$(f_\delta^{\text{II}})'(x) = 3\delta x^2 - 2\delta(1+p)x + 1 + p\delta,$$

$$(f_\delta^{\text{II}})''(x) = 6\delta x - 2\delta(1+p),$$

$$(f_\delta^{\text{II}})'''(x) = 6\delta.$$

Schwarzian derivative is negative if and only if $2(f_\delta^{\text{II}})'(f_\delta^{\text{II}})''' - 3((f_\delta^{\text{II}})'')^2 < 0$, thus we want to show that

$$12\delta[3\delta x^2 - 2\delta(1+p)x + 1 + p\delta] - 3[6\delta x - 2\delta(1+p)]^2 < 0.$$

Thus,

$$\delta(1+p+p^2 - 4(1+p)x + 6x^2) > 1$$

for any x . The minimum of $g(x) = 1+p+p^2 - 4(1+p)x + 6x^2$ is at $x = \frac{1+p}{3}$ and is equal to $\frac{p^2-p+1}{3}$. Therefore, $Sf_\delta^{\text{II}} < 0$ for $\delta > \frac{3}{1-p+p^2}$. \square

Now we are able to describe what happens when p is attracting, that is when $\delta < 2/p(1-p)$.

By Lemma 2 and Proposition 1 we get that p attracts all trajectories as long as $\delta \leq 2/p(1-p)$. In particular, as long as $\delta^* < 2/p(1-p)$ Nash equilibrium always attracts all points. This implies first two results for Model II.

Our aim is to investigate the long-term behavior of the orbits of f_δ^{II} . As for $\delta > 2/p(1-p)$ all fixed points are repelling there is no convergence to the equilibrium, and we need to study this case more thoroughly. In particular, we should look for periodic orbits, their stability, or chaos. Figure 2 suggests existence of chaotic behavior for p near-uniform and large values of δ . When speaking of chaos, we will use its most popular kind, *Li-Yorke chaos*.¹¹

¹¹For description of Li-Yorke chaos and its connections with other definitions of chaos we refer the reader to [11, 55]

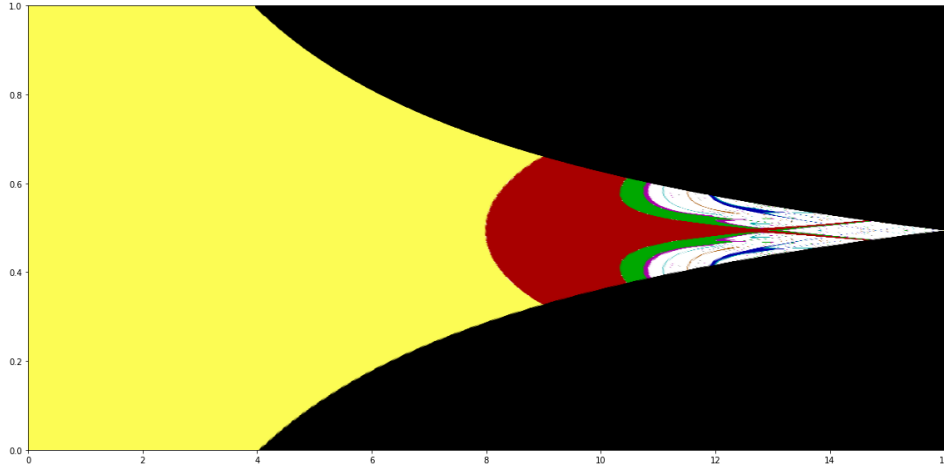


Figure 2: Period diagrams of the small-period attracting periodic orbits associated with the map f_δ^{II} (drawn on the black background). The horizontal axes are $\delta \in [4, 16]$ and the vertical axes are the asymmetry of gains $p \in [0, 1]$. The colors encode the periods of attracting periodic orbits as follows: period 1 (fixed point, which is Nash equilibrium p) = yellow, period 2 = red, period 3 = blue, period 4 = green, period 5 = brown, period 6 = cyan, period 7 = darkgray, period 8 = magenta, and period larger than 8 = white. The picture is generated from the following algorithm: 20000 preliminary iterations are discarded. Then a point is considered periodic of period n if $|(f_\delta^{\text{II}})^n(x) - x| < 10^{-10}$ and it is not periodic of any period smaller than n . As long as we are in the yellow region we have convergence to Nash equilibrium, once we get out of this region almost all trajectories will never converge to the fixed point.

As we deal with the interval map to prove Li-Yorke chaos we can use odd period argument, that is, show existence of periodic orbit of odd period, which implies chaotic behavior in the sense of Li and Yorke [35]. Thus, we will show the following proposition.

Proposition 2. *If $p \in (\frac{1}{23}(31 - 12\sqrt{3}), \frac{1}{23}(12\sqrt{3} - 8))$, then there exists a unique δ_p^{II} such that f_δ^{II} has periodic point of period 3 for any $\delta \in (\delta_p^{\text{II}}, \delta^*]$.*

Proof. We assume that $p \in (0, 1/2]$. We will show that for p sufficiently close to $1/2$ the following conditions hold:

$$f_\delta^{\text{II}}\left(\frac{1+p}{6}\right) > \frac{1+p}{2} \quad (16)$$

and

$$(f_\delta^{\text{II}})^2\left(\frac{1+p}{2}\right) < \frac{1+p}{6}. \quad (17)$$

Before justifying these inequalities let us show how (16) and (17) guarantee existence of the periodic point of period 3. From (16), continuity of f_δ^{II} and the fact that $f_\delta^{\text{II}}(0) = 0$, $f_\delta^{\text{II}}(p) = p$ we obtain existence of $x_0 \in (c_l, p)$ such that $f_\delta^{\text{II}}(x_0) = \frac{1+p}{2}$. From (17) and the fact that $f_\delta^{\text{II}}(x) > x$ when $x < p$, we have that

$$(f_\delta^{\text{II}})^3(x_0) = (f_\delta^{\text{II}})^2\left(\frac{1+p}{2}\right) < \frac{1+p}{6} < c_l < x_0 < f_\delta^{\text{II}}(x_0).$$

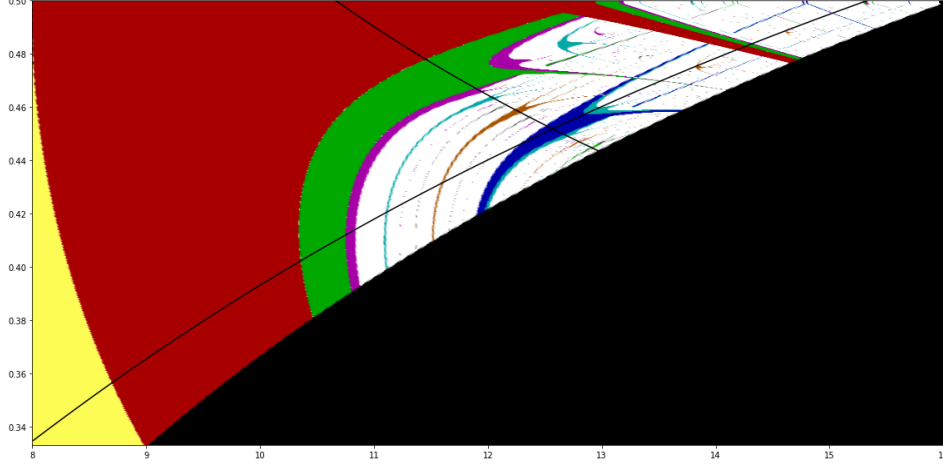


Figure 3: Period diagrams of the small-period attracting periodic orbits associated with the map f_δ^{II} (drawn on the black background). The horizontal axes are $\delta \in [8, 16]$ and the vertical axes are the asymmetry of gains $p \in [1/3, 1/2]$. The colors encode the periods of attracting periodic orbits as follows: period 1 (fixed point) = yellow, period 2 = red, period 3 = blue, period 4 = green, period 5 = brown, period 6 = cyan, period 7 = darkgray, period 8 = magenta, and period larger than 8 = white. The picture is generated from the following algorithm: 20000 preliminary iterations are discarded. Then a point is considered periodic of period n if $|(f_\delta^{\text{II}})^n(x) - x| < 10^{-10}$ and it is not periodic of any period smaller than n . On the picture we also draw the black curves for conditions (16) and (17) (these are curves crossing colored area, intersect at blue area).

Therefore, from theorem by Li, Misiurewicz, Panigiani and Yorke [36], f_δ^{II} has periodic orbit of period 3. By $f_{\delta,p}^{\text{II}}(x) = 1 - f_{\delta,1-p}^{\text{II}}(x)$ we will conclude similar result for $p > 1/2$.¹²

Now we will show conditions (16) and (17). We begin with the first inequality, which is equivalent to

$$-\frac{5}{216}\delta(1+p)^3 + \frac{1}{6}(1+p)(1+\delta p) > \frac{1+p}{2}.$$

Thus,

$$\delta \left(-\frac{5}{36}(1+p)^2 + p \right) > 2.$$

There is no positive δ which fulfills this condition as long as $p \leq \frac{1}{5}$. If $p > \frac{1}{5}$, then $\delta > \frac{72}{-5p^2+26p-5}$. As $\delta \leq \delta^*$ we obtain that (16) holds when

$$p \in \left(\frac{1}{23}(31 - 12\sqrt{3}), \frac{1}{2} \right]. \quad (18)$$

Next we turn our attention to condition (17), which is equivalent to

$$\frac{3}{256}(4 - \delta(1-p)^2)(64 - 16\delta(1-p)^2 + \delta^3(1-p)^4(1+p)^2) < 1. \quad (19)$$

¹²Although from numerical experiments we may conclude that periodic orbits of period 3 may arise for δ smaller than those for which conditions (16) and (17) hold, it seems that it estimates presence of period 3 quite well, see Figure 3.

Obviously if $\delta = \delta^*$, then the left hand side of (19) is equal to zero. Define

$$F(\delta) = \frac{3}{256}(4 - \delta(1 - p)^2)(64 - 16\delta(1 - p)^2 + \delta^3(1 - p)^4(1 + p)^2). \quad (20)$$

We want to show that there is a unique δ_p^{II} such that for $\delta > \delta_p^{II}$ values of F are smaller than 1.

We know that $F(0) = 3$ and $F(\delta^*) = 0$. The derivative of F is equal to

$$F'(\delta) = \frac{3}{256}(-68(1 - p)^2 + 2\delta(1 - p)^4 + 12\delta^2(1 - p)^4(1 + p)^2 - 4\delta^3(1 - p)^6(1 + p)^2).$$

As F' is a polynomial of degree 3 and

$$F'(0) = -\frac{51}{64(1 - p)^2} < 0,$$

with

$$F'(\delta^*) = \frac{3}{256}(-60(1 - p)^2 - 64(1 + p)^2) < 0,$$

we get that inside the interval $[0, \delta^*]$ the map F' can have 0 or 2 roots. As

$$F'\left(\frac{\delta^*}{2}\right) = -\frac{3}{16}(3 - 10p + 3p^2) > 0$$

when $p \in (\frac{1}{3}, \frac{1}{2}]$, we exclude possibility of no roots. Thus, inside the interval $[0, \delta^*]$ the map F has local minimum δ_{\min} and local maximum δ_{\max} , $\delta_{\min} < \delta_{\max}$. As $F(\delta^*) = 0$ we know that there exists δ_p^{II} such that (17) holds for $\delta > \delta_p^{II}$. To show uniqueness of δ_p^{II} we need to show that $F(\delta_{\min}) > 1$.

We have

$$F''(\delta) = \frac{3}{128}(1 - p)^4(1 + 12\delta(1 + p)^2 - 6\delta^2(1 + p)^2(1 - p)^2),$$

and simple calculations show that F'' has one positive root δ_1 . Thus,

$$0 < \delta_{\min} < \delta_1 < \delta_{\max}$$

and F is convex in $(0, \delta_1)$.

Now take $\delta = \frac{\delta^*}{4}$. Then $F'(\frac{\delta^*}{4}) = -\frac{3}{32}(11p^2 - 26p + 11) < 0$ as $p < \frac{1}{2} < \frac{1}{11}(13 - 4\sqrt{3})$. Define an affine map

$$G(\delta) = F'\left(\frac{\delta^*}{4}\right) \cdot \delta + F\left(\frac{\delta^*}{4}\right) - F'\left(\frac{\delta^*}{4}\right) \cdot \frac{\delta^*}{4}.$$

Because $F'(\frac{\delta^*}{4}) < 0$, the map G is decreasing. Convexity of F guarantees that $F(\delta_{\min}) \geq G(\delta_{\min})$. Moreover, $\frac{\delta^*}{4} < \delta_{\min} < \frac{\delta^*}{2}$. Thus, we have

$$F(\delta_{\min}) \geq G(\delta_{\min}) > G\left(\frac{\delta^*}{2}\right),$$

where

$$G\left(\frac{\delta^*}{2}\right) = F\left(\frac{\delta^*}{4}\right) + \frac{\delta^*}{4}F'\left(\frac{\delta^*}{4}\right).$$

Let

$$H(p) = G\left(\frac{\delta^*}{2}\right) - 1 = \frac{-79p^2 + 290p - 79}{256(1 - p)^2}.$$

We have $H(p) > 0$ when $p \in (\frac{1}{79}(145 - 8\sqrt{231}), \frac{1}{79}(145 + 8\sqrt{231}))$. Combining all restrictions we obtain that

$$F(\delta_{\min}) \geq G(\delta_{\min}) > G\left(\frac{\delta^*}{2}\right) > 1$$

for $p \in \left(\frac{1}{23}(31 - 12\sqrt{3}), \frac{1}{2}\right]$. This completes the proof for $p \in \left(\frac{1}{23}(31 - 12\sqrt{3}), \frac{1}{2}\right]$.

Finally,

$$1 - f_{\delta, 1-p}^{II}(1-x) = 1 - (1-x)(1 - \delta x(p-x)) = x(1 + \delta(p-x) - \delta x(p-x)) = f_{\delta, p}^{II}(x), \quad (21)$$

which implies the assertion of the theorem for $p \in \left(\frac{1}{23}(31 - 12\sqrt{3}), \frac{1}{23}(12\sqrt{3} - 8)\right)$. \square

We complete the proof of the case of Model II by applying Sharkovsky's theorem [66].

4.5. Dynamics of f_{δ}^{III} . The map (13) is already known and its dynamics was studied in various papers [16, 17, 47]. For the completeness of the exposition we describe here crucial properties of this map. Fixed points of f_{δ}^{III} are 0 and the roots of the equation

$$(1-x)(1 - \exp(\delta(x-p))) = 0.$$

So fixed points of f_{δ}^{III} are 0, 1 and p . The derivative of f_{δ}^{III} is given by

$$(f_{\delta}^{\text{III}})'(x) = \frac{(\delta x^2 - \delta x + 1) \exp(\delta(x-p))}{(x + (1-x) \exp(\delta(x-p)))^2}. \quad (22)$$

Thus,

$$(f_{\delta}^{\text{III}})'(0) = \exp(\delta p), \quad (f_{\delta}^{\text{III}})'(1) = \exp(\delta(1-p)), \quad (f_{\delta}^{\text{III}})'(p) = \delta p^2 - \delta p + 1.$$

We see that the fixed points 0 and 1 are always repelling, while p is repelling if $\delta > \frac{2}{p(1-p)}$.

The critical points of f_{δ}^{III} are solutions to $\delta x^2 - \delta x + 1 = 0$. Thus, if $0 < \delta \leq 4$, then f_{δ}^{III} is strictly increasing. If $\delta > 4$, it has two critical points

$$\kappa_l = \frac{1}{2} - \sqrt{\frac{1}{4} - \frac{1}{\delta}}, \quad \kappa_r = \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{1}{\delta}}, \quad (23)$$

so the map f_{δ}^{III} is bimodal.

Let us investigate regularity of f_{δ}^{III} . By Proposition 3.2 from [17] the map f_{δ}^{III} has negative Schwarzian derivative for $\delta > 4$.

For maps with negative Schwarzian derivative each attracting or neutral periodic orbit has a critical point in its immediate basin of attraction. Thus, we know that if $\delta > 4$ then f_{δ}^{III} can have at most two attracting or neutral periodic orbits.

Thus, $f_{\delta}^{\text{III}} \in \mathcal{F}$, so from Proposition 1 we get that as long as $\delta < 2p/(1-p)$, Nash equilibrium attracts all trajectories from $(0, 1)$.

Now, we sketch the proof of Li-Yorke chaos for f_{δ}^{III} when $\gamma_A \neq \gamma_B$. Although this fact was already shown in [17], we present here this proof to give a proper comparison with the proof of chaotic behavior for pairwise proportional imitation. To show chaotic behavior one can use period 3 arguments once more. To do that we notice that $f_{\delta}^{\text{III}}(x) > x$ when $x < p$. Moreover,

$$(f_{\delta}^{\text{III}})^n(x) = \frac{x}{x + (1-x) \exp(\delta \sum_{k=0}^{n-1} ((f_{\delta}^{\text{III}})^k(x) - p))}.$$

So, $(f_{\delta}^{\text{III}})^3(x) < x$ if and only if

$$x + f_{\delta}^{\text{III}}(x) + (f_{\delta}^{\text{III}})^2(x) > 3p. \quad (24)$$

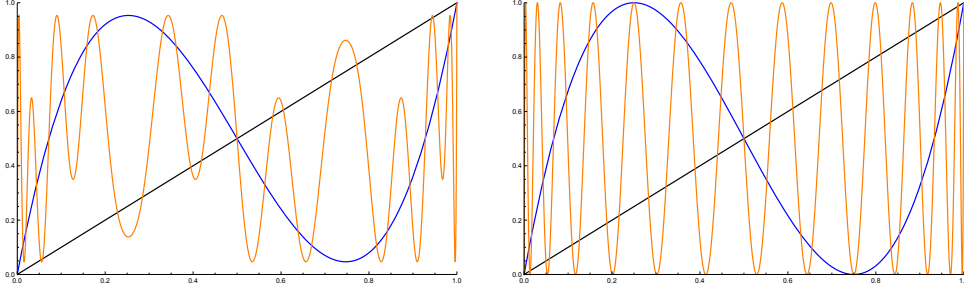


Figure 4: Map f_δ^{II} and its third iterate when $p = 0.5$ with $\delta = 15$ (left) and $\delta = \delta^* = 16$ (right).

For $0 < p < 1/2$ we take $3p - 1 < x < p$. As $\lim_{\delta \rightarrow \infty} f_\delta^{\text{III}}(x) = 1$ and $(f_\delta^{\text{III}})^2(x) > 0$ we see that (24) will be fulfilled for sufficiently large δ . Thus, we can use Misiurewicz-Li-Panigiani-Yorke theorem and Sharkovsky theorem and obtain Li-Yorke chaos for $p < 1/2$. As

$$\varphi \circ f_{\delta,p}^{\text{III}} = f_{\delta,1-p}^{\text{III}} \circ \varphi, \quad \text{where } \varphi(x) = 1 - x. \quad (25)$$

we have the same result for $p > 1/2$.

The result for $\gamma_A = \gamma_B$ follows from Theorem 3.10 from [17]. This completes the proof of Theorem 1.

Although Theorem 1 highlights differences in the dynamics introduced by different models we should comment on the convergence part of the theorem.

Remark 3. For small step size all three models converge to Nash equilibrium. Moreover, the thresholds for Model II and Model III are the same.

Remark 3 is the consequence of Remark 2 and the fact that $f_\delta^{\text{II}'}(p) = f_\delta^{\text{III}'}(p)$. This dependence suggests that the connection between Model II and Model III is closer than with Model I where the derivative $f_\delta^{\text{I}'}(p)$ is different.

Finally, let us comment on the technical differences in the proof for Model II and Model III. Both are done by using Li-Misiurewicz-Panigiani-Yorke theorem [36] and period 3 argument. To show this we need to find a point $x_0 \in (0, 1)$ such that $f^3(x_0) < x_0 < f(x_0)$ or $f^3(x_0) > x_0 > f(x_0)$. Nevertheless, in Model III the trick is to take very large step size (depending on b). Moreover, when δ is sufficiently large, the choice of a point which fulfills the condition is easy as we can take it from the large range of values, see Figure 5. This can't be done in Model II, where the step size is bounded. In addition, the condition for the third iterate is met only on small interval which varies with the step size, see Figure 4. Thus, we carefully choose then a point which first iterate lays close to the (right) critical point, but values can be estimated analytically. This cannot be done for the f_δ^{III} map, where we can get only numerical estimations.

5. DISCUSSION

Our work shows that three distinct game dynamics in discrete time – a biological model of intra-species competition in the spirit of Moran [43], evolution under the pairwise proportional imitation protocol of Helbing [29] in economic theory, and learning with the exponential / multiplicative weights algorithm from the theory of adversarial online learning [3] – exhibit qualitatively different long-run properties, despite the fact that they all share the *same* continuous-time limit – the replicator dynamics. This disconnection occurs even in the simplest of games – a 2×2 symmetric random matching congestion game – and

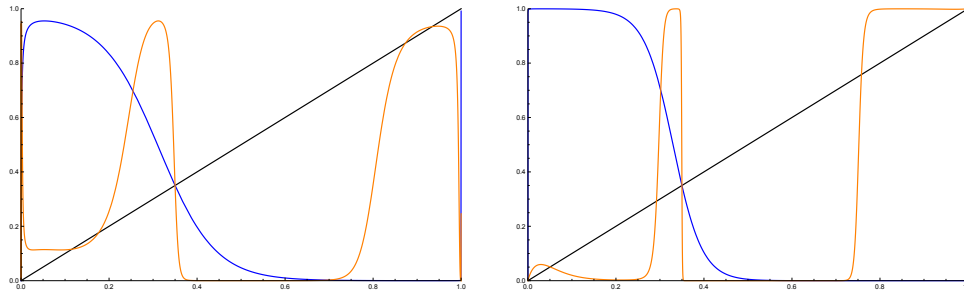


Figure 5: Map f_δ^{III} and its third iterate when $p = 0.35$ with $\delta = 20$ (left) and $\delta = 35$ (right).

leads to drastically different predictions (or lack thereof): *a*) the biological model guarantees universal convergence to Nash equilibrium for all initial conditions and all equilibrium and (hyper)parameter configurations; *b*) the economic model demonstrates the entire range of possible behaviors for large δ (convergence, instability, periodic and chaotic behavior); and, finally, *c*) fully chaotic behavior under the EW algorithm (for large δ). This divergence of behaviors provides a crisp cautionary tale to the effect that “*discretization matters*”, and serves to highlight the extent to which concrete conclusions can be drawn from the behavior of continuous-time models – or, rather, the failure thereof.

Finally, we expect that the discrepancy between discrete dynamics of discussed models is robust. It should occur for any anti-coordination game (9) in which strategy profiles (A, A) and (B, B) would give different payoffs. Moreover, we know that if one considers more strategies, then for Model III one can detect chaotic behavior [16]. Eventually, we showed that the expectation that in simple games qualitative behavior of models convolved by replicator equation in the continuous limit, fails. And thus different qualitative behavior in more complicated games should be expected. Moreover, discretization in higher dimensions may substantially change the dynamics in both directions [24].

REFERENCES

- [1] Z. AlSharawi and M. B. Rhouma. The beverton–holt model with periodic and conditional harvesting. *Journal of Biological Dynamics*, 3(5):463–478, 2009.
- [2] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: A meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- [3] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, 1995.
- [4] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- [5] R. A. Becker, S. K. Chakrabarti, W. Geller, B. Kitchens, and M. Misiurewicz. Dynamics of the Nash map in the game of Matching Pennies. *Journal of Difference Equations and Applications*, 13(2-3): 223–235, 2007.
- [6] M. Benaïm. Dynamics of stochastic approximation algorithms. In J. Azéma, M. Émery, M. Ledoux, and M. Yor, editors, *Séminaire de Probabilités XXXIII*, volume 1709 of *Lecture Notes in Mathematics*, pages 1–68. Springer Berlin Heidelberg, 1999.
- [7] R. J. Beverton and S. J. Holt. *On the dynamics of exploited fish populations*, volume 11. Springer Science & Business Media, 2012.
- [8] J. Bielawski, T. Chotibut, F. Falniowski, G. Kosiorowski, M. Misiurewicz, and G. Piliouras. Follow-the-regularized-leader routes to chaos in routing games. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *PMLR*, pages 925–935, 2021.

- [9] J. Bielawski, T. Chotibut, F. Falniowski, M. Misiurewicz, and G. Piliouras. Memory loss can prevent chaos in games dynamics. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 34(1):013146, 2024. doi: 10.1063/5.0184318.
- [10] K. Binmore and L. Samuelson. Muddling through: Noisy equilibrium selection. *Journal of Economic Theory*, 74(2):235–265, June 1997.
- [11] F. Blanchard. Topological chaos: what may this mean? *Journal of Difference Equations and Applications*, 15(1):23–46, 2009.
- [12] V. Boone and P. Mertikopoulos. The equivalence of dynamic and strategic stability under regularized learning in games. In *NeurIPS '23: Proceedings of the 37th International Conference on Neural Information Processing Systems*, 2023.
- [13] E. Braverman and D. Kinzebulatov. On linear perturbations of the ricker model. *Mathematical Biosciences*, 202(2):323–339, 2006.
- [14] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [15] 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 *Conference on Learning Theory*, PMLR, pages 807–834, 2019.
- [16] T. Chotibut, F. Falniowski, M. Misiurewicz, and G. Piliouras. The route to chaos in routing games: When is price of anarchy too optimistic? *Advances in Neural Information Processing Systems*, 33: 766–777, 2020.
- [17] T. Chotibut, F. Falniowski, M. Misiurewicz, and G. Piliouras. Family of chaotic maps from game theory. *Dynamical Systems*, 36(1):48–63, 2021. <https://doi.org/10.1080/14689367.2020.1795624>.
- [18] W. De Melo and S. Van Strien. *One-dimensional dynamics*, volume 25. Springer Science & Business Media, 2012.
- [19] S. Fischer, H. Räcke, and B. Vöcking. Fast convergence to wardrop equilibria by adaptive sampling methods. In *Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing*, STOC '06, pages 653–662, New York, NY, USA, 2006. ACM. ISBN 1-59593-134-1.
- [20] D. Fotakis, A. C. Kaporis, and P. G. Spirakis. Atomic congestion games: Fast, myopic and concurrent. In B. Monien and U.-P. Schroeder, editors, *Algorithmic Game Theory*, volume 4997 of *Lecture Notes in Computer Science*, pages 121–132. Springer Berlin Heidelberg, 2008. ISBN 978-3-540-79308-3.
- [21] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*, volume 2 of *Economic learning and social evolution*. MIT Press, Cambridge, MA, 1998.
- [22] T. Galla and J. D. Farmer. Complex dynamics in learning complicated games. *Proceedings of the National Academy of Sciences*, 110(4):1232–1236, 2013. ISSN 0027-8424.
- [23] W. Geller, B. Kitchens, and M. Misiurewicz. Microdynamics for Nash maps. *Discrete & Continuous Dynamical Systems*, 27(3):1007–1024, 2010.
- [24] R. T. Guy and M. Misiurewicz. Euler approximations can destroy unbounded solutions. *Fasciculi Mathematici*, 44:43–52, 2010.
- [25] S. Hadikhanloo, R. Laraki, P. Mertikopoulos, and S. Sorin. Learning in nonatomic games, Part I: Finite action spaces and population games. *Journal of Dynamics and Games*, 9(4, William H. Sandholm memorial issue):433–460, October 2022.
- [26] J. Hannan. Approximation to Bayes risk in repeated play. In M. Dresher, A. W. Tucker, and P. Wolfe, editors, *Contributions to the Theory of Games, Volume III*, volume 39 of *Annals of Mathematics Studies*, pages 97–139. Princeton University Press, Princeton, NJ, 1957.
- [27] S. Hart and A. Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93(5):1830–1836, 2003.
- [28] S. Hart and A. Mas-Colell. Stochastic uncoupled dynamics and Nash equilibrium. *Games and Economic Behavior*, 57:286–303, 2006.
- [29] D. Helbing. A mathematical model for behavioral changes by pair interactions. In G. Haag, U. Mueller, and K. G. Troitzsch, editors, *Economic Evolution and Demographic Change: Formal Models in Social Sciences*, pages 330–348. Springer, Berlin, 1992.
- [30] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK, 1998.
- [31] Y.-P. Hsieh, P. Mertikopoulos, and V. Cevher. The limits of min-max optimization algorithms: Convergence to spurious non-critical sets. In *ICML '21: Proceedings of the 38th International Conference on Machine Learning*, 2021.

- [32] V. Jiménez López and E. Liz. Destabilization and chaos induced by harvesting: Insights from one-dimensional discrete-time models. *Journal of Mathematical Biology*, 82:1–28, 2021.
- [33] R. Kleinberg, G. Piliouras, and É. Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 533–542, 2009.
- [34] M. Kot. *Elements of mathematical ecology*. Cambridge University Press, 2001.
- [35] T.-Y. Li and J. A. Yorke. Period three implies chaos. *The American Mathematical Monthly*, 82(10):985–992, 1975.
- [36] T.-Y. Li, M. Misiurewicz, G. Pianigiani, and J. A. Yorke. Odd chaos. *Physics Letters A*, 87(6):271–273, 1982.
- [37] N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.
- [38] K. Lotidis, P. Mertikopoulos, and N. Bambos. The stability of matrix multiplicative weights dynamics in quantum games. In *CDC '23: Proceedings of the 62nd IEEE Annual Conference on Decision and Control*, 2023.
- [39] P. Mertikopoulos and Y. Viossat. Survival of dominated strategies under imitation dynamics. *Journal of Dynamics and Games*, 9(4, William H. Sandholm memorial issue):499–528, October 2022.
- [40] P. Mertikopoulos, C. Papadimitriou, and G. Piliouras. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2703–2717. SIAM, 2018.
- [41] P. Mertikopoulos, Y.-P. Hsieh, and V. Cevher. A unified stochastic approximation framework for learning in games. *Mathematical Programming*, 203:559–609, January 2024.
- [42] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- [43] P. A. P. Moran. *The Statistical Processes of Evolutionary Theory*. Clarendon Press, Oxford, 1962.
- [44] A. Mukhopadhyay and S. Chakraborty. Periodic orbit can be evolutionarily stable: Case study of discrete replicator dynamics. *Journal of theoretical biology*, 497:110288, 2020.
- [45] A. Mukhopadhyay and S. Chakraborty. Replicator equations induced by microscopic processes in nonoverlapping population playing bimatrix games. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31(2):023123, 2021.
- [46] J. F. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences of the USA*, 36:48–49, 1950.
- [47] G. Palaioupanos, I. Panageas, and G. Piliouras. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. In *Advances in Neural Information Processing Systems*, pages 5872–5882, 2017.
- [48] I. Panageas and G. Piliouras. Average case performance of replicator dynamics in potential games via computing regions of attraction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 703–720, 2016.
- [49] V. Pandit, A. Mukhopadhyay, and S. Chakraborty. Weight of fitness deviation governs strict physical chaos in replicator dynamics. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28(3):033104, 2018.
- [50] M. Pangallo, J. B. Sanders, T. Galla, and J. D. Farmer. Towards a taxonomy of learning dynamics in 2×2 games. *Games and Economic Behavior*, 132:1–21, 2022.
- [51] G. Piliouras and J. S. Shamma. Optimization despite chaos: Convex relaxations to complex limit sets via poincaré recurrence. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 861–873. SIAM, 2014.
- [52] A. Rakhlin and K. Sridharan. Optimization, learning, and games with predictable sequences. In *NIPS '13: Proceedings of the 27th International Conference on Neural Information Processing Systems*, 2013.
- [53] W. E. Ricker. Stock and recruitment. *Journal of the Fisheries Board of Canada*, 11(5):559–623, 1954.
- [54] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [55] S. Ruelle. *Chaos on the interval*, volume 67 of *University Lecture Series*. American Mathematical Society, 2017.
- [56] A. Rustichini. Optimal properties of stimulus-response learning models. *Games and Economic Behavior*, 29(1-2):244–273, 1999.

- [57] A. Rustichini. Minimizing regret: The general case. *Games and Economic Behavior*, 29(1-2):224–243, October 1999.
- [58] J. B. T. Sanders, J. D. Farmer, and T. Galla. The prevalence of chaotic dynamics in games with many players. *Scientific Reports*, 8(1):1–13, 2018.
- [59] W. H. Sandholm. Potential games with continuous player sets. *Journal of Economic theory*, 97(1): 81–108, 2001.
- [60] W. H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, MA, 2010.
- [61] W. H. Sandholm, E. Dokumacı, and R. Lahkar. The projection dynamic and the replicator dynamic. *Games and Economic Behavior*, 64(2):666–683, 2008.
- [62] Y. Sato and J. P. Crutchfield. Coupled replicator equations for the dynamics of learning in multiagent systems. *Physical Review E*, 67(1):015206, 2003.
- [63] Y. Sato, E. Akiyama, and J. D. Farmer. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*, 99(7):4748–4751, 2002. ISSN 0027-8424.
- [64] S. J. Schreiber. Chaos and population disappearances in simple ecological models. *Journal of Mathematical Biology*, 42:239–260, 2001.
- [65] S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.
- [66] A. N. Sharkovsky, S. F. Kolyada, A. G. Sivak, and V. V. Fedorenko. Dynamics of one-dimensional maps. *Mathematics and its Applications*, 407, 1997.
- [67] M. Shub. *Global Stability of Dynamical Systems*. Springer-Verlag, Berlin, 1987.
- [68] S. Sorin. Exponential weight algorithm in continuous time. *Mathematical Programming*, 116(1):513–528, 2009.
- [69] C. Sparrow, S. van Strien, and C. Harris. Fictitious play in 3x3 games: The transition between periodic and chaotic behaviour. *Games and Economic Behavior*, 63(1):259 – 291, 2008. ISSN 0899-8256.
- [70] P. D. Taylor and L. B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1-2):145–156, 1978.
- [71] H. Thunberg. Periodicity versus chaos in one-dimensional dynamics. *SIAM review*, 43(1):3–30, 2001.
- [72] S. van Strien and C. Sparrow. Fictitious play in 3x3 games: Chaos and dithering behaviour. *Games and Economic Behavior*, 73(1):262 – 286, 2011. ISSN 0899-8256.
- [73] D. Vilone, A. Robledo, and A. Sánchez. Chaos and unpredictability in evolutionary dynamics in discrete time. *Physical review letters*, 107(3):038101, 2011.
- [74] Y. Viossat and A. Zapechelnuk. No-regret dynamics and fictitious play. *Journal of Economic Theory*, 148(2):825–842, March 2013.
- [75] V. G. Vovk. Aggregating strategies. In *COLT '90: Proceedings of the 3rd Workshop on Computational Learning Theory*, pages 371–383, 1990.
- [76] J. W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.